PETER SCOTT
GADDE A. SWARUP

Regular neighbourhoods and canonical decompositions for groups

Astérisque, tome 289 (2003)
<http://www.numdam.org/item?id=AST_2003__289__R1_0>

ASTÉRISQUE 289

REGULAR NEIGHBOURHOODS AND CANONICAL DECOMPOSITIONS FOR GROUPS

Peter Scott

Gadde A. Swarup

Société Mathématique de France 2003

Publié avec le concours du Centre National de la Recherche Scientifique

P. Scott
Mathematics Department, University of Michigan, Ann Arbor, Michigan 48109, USA.
E-mail : pscott@umich.edu
Url : http://www.math.lsa.umich.edu/~pscott

G.A. Swarup
Department of Mathematics and Statistics, University of Melbourne, Victoria 3010, Australia.
E-mail : gadde@ms.unimelb.edu.au
Url : http://u3nma05.u-3mrs.fr/~asgadde

2000 *Mathematics Subject Classification.* — 20E06, 20E08, 20E34, 20F65, 57M07, 57M50, 57N12, 57N16.

Key words and phrases. — Canonical decomposition, splitting, trees, ends, regular neighbourhood, characteristic submanifold.

The first author was partially supported by NSF grants DMS 034681 and 9626537. The second author was partially supported by ARC grant DP 0209769.

REGULAR NEIGHBOURHOODS AND CANONICAL DECOMPOSITIONS FOR GROUPS

Peter Scott, Gadde A. Swarup

Abstract. — We find canonical decompositions for (almost) finitely presented groups which essentially specialise to the classical JSJ-decomposition when restricted to the fundamental groups of Haken manifolds. The decompositions that we obtain are invariant under automorphisms of the group. A crucial new ingredient is the concept of a regular neighbourhood of a family of almost invariant subsets of a group. An almost invariant set is an analogue of an immersion.

Résumé **(Voisinages réguliers et décompositions canoniques pour les groupes)**
Nous définissons une décomposition canonique pour les groupes presque finiment présentés qui correspond à la décomposition JSJ classique dans le cas du groupe fondamental d'une variété de Haken. Les automorphismes du groupe laissent invariante cette décomposition. Un élément crucial et nouveau est le concept de voisinage régulier d'une famille de sous-ensembles du groupe qui sont presque invariants. Un ensemble presque invariant est un analogue d'une immersion.

CONTENTS

Appendix

INTRODUCTION

This article is devoted to the study of analogues for groups of the classical JSJ-decomposition (see Jaco and Shalen [**25**], Johannson [**26**] and Waldhausen [**56**]) for orientable Haken 3–manifolds. The orientability restriction is not essential but it will simplify our discussions. An announcement of our results is in [**46**]. This field was initiated by Kropholler [**27**] who studied analogous decompositions for Poincaré duality groups of any dimension greater than 2. But the current interest in this kind of decomposition started with the work of Sela [**49**] on one-ended torsion-free hyperbolic groups. His results were generalised by Rips and Sela [**36**], Bowditch [**5**][**8**], Dunwoody and Sageev [**14**], Dunwoody and Swenson [**15**], and Fujiwara and Papasoglu [**20**], but none of these results yields the classical JSJ-decomposition when restricted to the fundamental group of an orientable Haken manifold. In this paper, we give a new approach to this subject, and we give decompositions for finitely presented groups which essentially specialise to the classical JSJ-decomposition when restricted to the fundamental groups of Haken manifolds. An important feature of our approach is that the decompositions we obtain are unique and are invariant under automorphisms of the group. In previous work such strong uniqueness results were only found for decompositions of word hyperbolic groups. Most of the results of these previous authors for virtually polycyclic groups can be deduced from our work. But our arguments use some of the results of these authors, particularly those of Bowditch. In addition, we use the important work of Dunwoody and Roller in [**13**]. Our work also yields some extensions of the results on the Algebraic Annulus and Torus Theorems in [**43**], [**5**] and [**15**]. It should be remarked that even though we obtain canonical decompositions for all finitely presented groups, these decompositions are often trivial. This is analogous to the fact that any finitely generated group possesses a free product decomposition, but this decomposition is trivial whenever the given group is freely indecomposable. We should also remark that many of the ideas in this paper and the above mentioned papers can be traced back to the groundbreaking work of Stallings on groups with infinitely many ends [**52**][**53**].

We focus on what we consider to be the most important aspects of the topological JSJ-decomposition. Our choice of the crucial property of this decomposition is the Enclosing Property of the characteristic submanifold, and we use an algebraic generalisation of this property. The topological Enclosing Property can be described briefly as follows. See chapter 1 for a more detailed discussion. For an orientable Haken 3–manifold M, Jaco and Shalen [**25**] and Johannson [**26**] proved that there is a family $\mathcal{T}$ of disjoint essential annuli and tori embedded in M, unique up to isotopy, and with the following properties. The manifolds obtained by cutting M along $\mathcal{T}$ are simple or are Seifert fibre spaces or I–bundles over surfaces. The Seifert and I–bundle pieces of M are said to be *characteristic*, and any essential map of the annulus or torus into M can be properly homotoped to lie in a characteristic piece. This is called the Enclosing Property of $\mathcal{T}$. The characteristic submanifold $V(M)$ of M consists essentially (see chapter 1 for details) of the union of the characteristic pieces of the manifold obtained from M by cutting along $\mathcal{T}$. The fundamental group G of M is the fundamental group of a graph Γ of groups, whose underlying graph is dual to the frontier of $V(M)$. Thus the edge groups of Γ are all isomorphic to $\mathbb{Z}$ or $\mathbb{Z} \times \mathbb{Z}$, and the vertex groups are the fundamental groups of simple manifolds or of Seifert fibre spaces or of surfaces. The uniqueness up to isotopy of the splitting family $\mathcal{T}$ implies that Γ is unique. Further, the Enclosing Property implies that any subgroup of G which is represented by an essential annulus or torus in M is conjugate into a characteristic vertex group.

All the previous algebraic analogues of the topological JSJ-decomposition consist of producing a graph of groups structure Γ for a given group G with the edge groups of Γ being of some specified type and with some "characteristic" vertices. The algebraic analogue of the topological Enclosing Property which was used is the property that certain "essential" subgroups of G must be conjugate into one of the characteristic vertex groups of Γ. Note that the word "essential" was not used by any of these authors, and they considered several different classes of subgroups. We use the term as a convenience to allow us to compare their differing results.

Our results also yield a graph of groups structure Γ for a given group G with some "characteristic" vertices, but our algebraic generalisation of the topological Enclosing Property corresponds more closely to the topological situation.

Here is a more detailed discussion of the previous algebraic analogues of the topological JSJ-decomposition. In all of these cases, G is a finitely presented one-ended group, and an essential subgroup of G is of the same abstract type as the edge groups of Γ. For example, when trying to describe all splittings of a group G over infinite cyclic subgroups, previous authors produced a decomposition with infinite cyclic edge groups, such that if G splits over an infinite cyclic subgroup H, then H is conjugate into a characteristic vertex group. The first such result was by Kropholler [**27**], who considered the special case when G is a Poincaré duality group of dimension n and the edge groups of Γ are virtually polycyclic (VPC) groups of Hirsch length $n-1$. For brevity, we will refer to the length rather than the Hirsch length of a VPC group throughout this paper. We will also say that a VPC group of length n is $VPCn$. In his case, any $VPC(n-1)$ subgroup H is essential. Such a subgroup H will have a subgroup K of index at most 2 such that $e(G,K) \geqslant 2$. This corresponds to considering all π_1–injective maps of closed $(n-1)$–dimensional manifolds into a n–manifold rather than considering just embeddings of such manifolds. Note that a VPC group of length at most 2 is virtually abelian, so that when $n=3$ his result is closely related to the topology of 3–manifolds. In fact, Kropholler [**28**] used his results in [**27**] to give a new proof of the existence of the JSJ-decomposition for closed 3–manifolds.

In most of the papers which came after [**27**], a subgroup H of G is essential if G possesses a splitting over H. Such subgroups correspond to embedded codimension–1 manifolds in a manifold. Sela [**48**] considered the case when G is a torsion-free word hyperbolic group, the essential subgroups are infinite cyclic and the edge groups of Γ are also infinite cyclic. Rips and Sela [**36**] generalised this to the case where G is a torsion-free finitely presented group. The essential subgroups and the edge groups are again infinite cyclic. Dunwoody and Sageev [**14**] considered the case when G is a finitely presented group and the essential subgroups and the edge groups of Γ are slender groups (i.e. groups in which every subgroup is finitely generated), subject to the constraint that if H is an edge group, then G admits no splitting over a subgroup of infinite index in H. Fujiwara and Papasoglu [**20**] considered the case when G is a finitely presented group and the essential subgroups and the edge groups of Γ are finitely generated small groups (i.e. groups which do not admit a hyperbolic action on a tree), subject to the weaker constraint that if H is an edge group, then no splitting of G over H can cross strongly a splitting of G over a subgroup of H of infinite index. (See chapter 2 for a discussion of crossing and strong crossing.) Dunwoody and Swenson [**15**] considered the case when G is a finitely presented group and the essential subgroups and the edge groups of Γ are VPC groups of a fixed length n, subject to the constraint that G admits no splitting over a VPC subgroup of length less than n. In their work, a $VPCn$ subgroup H of G is essential if $e(G,H) \geqslant 2$. This corresponds to considering singular codimension–1 manifolds in a manifold rather than just embedded ones. Finally, Bowditch [**5**] considered the case when G is a word hyperbolic group and the essential subgroups of G and the edge groups of Γ are two-ended (which is equivalent to being virtually infinite cyclic). In his work, a two-ended subgroup H of G is essential if $e(G,H) \geqslant 2$. This corresponds to considering all essential annuli in a 3–manifold rather than just embedded ones. In this case, Bowditch proved an existence and uniqueness result, precisely analogous to the 3–manifold theory in the atoroidal case.

The above results are often referred to vaguely but collectively as the JSJ-decomposition of a finitely presented group. While these results are commonly regarded as being an algebraic analogue of the topological JSJ theory, none of them recovers the topological result when applied to the fundamental group of an orientable Haken 3–manifold. We list some reasons for this.

– None of them has as strong a uniqueness property as the topological JSJ-decomposition, apart from Bowditch's work [**5**] for word hyperbolic groups. In particular, there is no invariance under automorphisms of the group G except when G is word hyperbolic.

– The topological JSJ-decomposition involves both annuli and tori. Apart from [**20**], none of the algebraic theories can simultaneously handle splittings over free abelian groups H and K of different ranks, and [**20**] can only handle this in certain cases.

– In all the previous work apart from that of Bowditch [**5**], of Dunwoody and Swenson [**15**], and of Kropholler [**27**] only splittings are considered, whereas in the topological JSJ theory, singular annuli and tori play a crucial role.

– The algebraic theories only consider strong crossing of splittings, whereas weak crossing is a key ingredient in the topology.

– The Enclosing Property of the characteristic submanifold $V(M)$ is stronger than the condition stated earlier that any subgroup of $G = \pi_1(M)$ which is represented by an essential annulus or torus in M is conjugate into a characteristic vertex group. For tori there is no difference, but there are many examples where a Haken 3–manifold M contains two essential annuli A and A' which carry the same subgroup of G, and $V(M)$ has a component W such that A is properly homotopic into W but A' is not.

Now we will discuss the above points in more detail.

A feature one would expect from any canonical decomposition of a group G is some sort of invariance under automorphisms of G, but in previous work this was present only in the case when G is word hyperbolic. This invariance has been exploited by Sela in the case of word hyperbolic groups [**49**], and by Johannson in the classical case [**26**], for several striking applications.

The characteristic pieces of a 3–manifold M are of three types, namely I–bundles (which must meet ∂M), Seifert fibre spaces which meet ∂M, and Seifert fibre spaces which are in the interior of M. We refer to the two types which meet ∂M as peripheral. For the discussion in this paragraph, we restrict our attention to essential annuli in M. Then the collection of all peripheral characteristic pieces of M has the Enclosing Property for all such annuli and, except for a few special cases, no subcollection has this property. But if one applies any of the above algebraic results to splittings of $G = \pi_1(M)$ over infinite cyclic subgroups, one obtains only analogues of the I–bundle pieces, and the peripheral Seifert fibre spaces are split up in an arbitrary fashion. Note that by splitting a peripheral Seifert fibre space, one loses the structure of this piece and hence loses the topological Enclosing Property. A further point to note is that the above algebraic splittings may not even yield all the I–bundle pieces of M. For suppose that M has an I–bundle piece homeomorphic to $F \times I$, where F is homeomorphic to the thrice punctured sphere. Thus the only embedded incompressible annuli in $F \times I$ are the three boundary components. Now all the relevant algebraic results, apart from those of Dunwoody and Swenson [**15**] and of Bowditch [**5**], consider only infinite cyclic subgroups of G over which G splits, and this corresponds to restricting attention to embedded annuli. Thus these results will yield a graph of groups structure for G with three edges corresponding to the three boundary components of F, but the vertex with associated group $\pi_1(F)$ will not be regarded as characteristic. One defect of this is that the structure of the characteristic vertex groups of such decompositions cannot be a quasi-isometry invariant. More precisely, the property of having a characteristic vertex group which is not two-ended is not an invariant of the quasi-isometry type of G. For suppose that the characteristic submanifold of a 3–manifold M consists of an I–bundle X homeomorphic to $F \times I$, where F is homeomorphic to the thrice punctured sphere. Let M' be a finite cover of M in which a component X' of the pre-image of X is homeomorphic to $F' \times I$, where F' is not homeomorphic to F, and let G' denote $\pi_1(M')$. Then the decomposition of G' will have a characteristic vertex group which is not two-ended, but the corresponding decomposition of G will not.

There are two important special cases of the topological JSJ-decomposition of an orientable Haken 3–manifold M. One occurs when M is closed, in which case $\mathcal{T}$ consists of tori only and the Enclosing Property applies only to maps of tori into M. The other occurs when M is atoroidal, meaning that $\pi_1(M)$ contains no non-peripheral $(\mathbb{Z} \times \mathbb{Z})$–subgroup, in which case $\mathcal{T}$ consists of annuli only and the Enclosing Property applies only to maps of annuli into M. These two special cases seem to have guided the development of all previous algebraic analogues of the topological JSJ-decomposition. In particular, when trying to describe all splittings of a group over subgroups of a given type, for example infinite cyclic, previous authors looked for a decomposition described by splittings over subgroups of the same type. However, if we return to the general topological situation and consider only essential annuli, we observe that the collection of peripheral characteristic pieces of M may well have some frontier tori. (See the end of chapter 1 for some relevant examples.) This means that even if one wishes initially to consider only splittings of a group over infinite cyclic subgroups, one is naturally led to consider splittings over more complicated subgroups as well. Surprisingly, we will see that, in general, these more complicated groups need not even be finitely generated.

We believe that our ideas in this paper handle all the above problems. We obtain decompositions of all finitely presented groups which are unique, and hence invariant under automorphisms, and which essentially specialise to the classical JSJ-decomposition. In particular, our ideas can handle simultaneously splittings over free abelian groups of many different ranks. Our decompositions arise in a simple and natural way, whereas the previous constructions were all rather indirect. We obtain more than one such decomposition and there seems to be a number of further questions about finding refinements and properties of these decompositions. We are preparing two further papers on such questions. In one, we consider a relative version of much of the theory in this paper. This involves developing a theory of relative almost invariant subsets. In another paper, we consider the case of Poincaré duality pairs in more detail, and discuss an analogue of the full topological JSJ decomposition.

Here is an introduction to our ideas. As mentioned before, our choice of the crucial feature of the classical JSJ-decomposition is the Enclosing Property for immersions. This property implies that the characteristic submanifold $V(M)$ of an orientable Haken 3–manifold M contains a representative of every homotopy class of an essential annulus or torus in M. We will say that it *encloses* every essential annulus and torus in M. In this paper, we introduce a natural algebraic analogue of enclosing. If we restrict attention to embedded surfaces, this analogue is simple to explain. First recall the graph of groups structure Γ for $G = \pi_1(M)$, whose underlying graph is dual to the frontier $\mathrm{fr}(V(M))$ of $V(M)$. Let F be an essential annulus or torus embedded in a component W of $V(M)$, so that F determines a splitting of G, and let Γ_F denote the graph of groups structure for G, whose underlying graph is dual to $\mathrm{fr}(V(M)) \cup F$. Thus Γ_F has one more edge than Γ, and this extra edge f corresponds to F. Collapsing f yields Γ again, and the image of f in Γ is the vertex which corresponds to W. Now let G be any group, let Γ be a graph of groups structure for G, and let σ be a splitting of G. We say that σ is *enclosed* by a vertex v of Γ, if there is a graph of groups structure Γ_σ for G, with an edge e which determines the splitting σ, such that collapsing the edge e yields Γ, and v is the image of e. We emphasise that the condition that σ is enclosed by the vertex v is in general stronger than the condition that the edge group of σ is conjugate into the vertex group of v. This is particularly clear if σ is a free product decomposition of G, as then the edge group of σ is trivial.

An important observation is that $V(M)$ is closely related to a regular neighbourhood of some (finite) union of essential annuli and tori in M. In some cases, one can choose a finite family of essential annuli and tori in $V(M)$ so that $V(M)$ is a regular neighbourhood of their union. More usually, $V(M)$ can be obtained from such a regular neighbourhood by adding solid tori to compressible torus boundary components. In particular, except for a few special cases, $V(M)$ is minimal (up to isotopy) among incompressible submanifolds of M which enclose every essential annulus and torus in M. Thus it seems natural to think of $V(M)$ as a regular neighbourhood of all the essential annuli and tori in M. The peripheral pieces of the characteristic submanifold can be thought of as a regular neighbourhood of all the essential annuli only. Our main results can be thought of as algebraic versions of these statements. In order to explain our ideas further, we need to discuss the algebraic analogues of immersed annuli and tori and the algebraic analogue of a regular neighbourhood.

An analogue of a two-sided π_1–injective immersion in codimension 1 has been studied by group theorists for some time. If H denotes the image in G of the fundamental group of the codimension–1 manifold, this analogue is a subset of G called a H–almost invariant set or an almost invariant set over H. Any two-sided π_1–injective immersion in codimension 1 has a H–almost invariant set associated to it in a natural way. In particular, this applies to any splitting of G. Further, there is a natural idea of what it means for an almost invariant subset of a group G to be enclosed by a vertex of a graph of groups decomposition for G, which generalises the idea of enclosing a splitting. We should mention that almost invariant subsets of a group G can appear disconnected in the Cayley graph of G so that the analogy with immersions may seem a little forced. However, this is an artifact of the particular choice of generators made when constructing the Cayley graph. So long as H is finitely generated, one can always change generators to make any given H–almost invariant set connected. The appropriate notions of intersection and disjointness for almost invariant sets were introduced by Scott in [**42**]. These notions were further developed in [**44**] and the necessary definitions and results will be recalled in chapter 2. For the convenience of the reader, the full texts of [**42**] and [**44**] are included as appendices in this paper. In [**42**], Scott defined the intersection number of two nontrivial almost invariant subsets of a group and showed it was symmetric. Further his definition generalises the natural idea of intersection number of curves on a surface, and the intersection number of closed surfaces in a 3–manifold introduced in [**19**]. In [**44**], the main results, Theorems 2.5 and 2.8, were algebraic analogues of the facts that curves on a surface with intersection number zero can be homotoped to be disjoint, and that a curve with self-intersection number zero can be homotoped to cover an embedding. These results do strongly suggest that an almost invariant set is the appropriate analogue of an immersion.

The key new idea of this paper is an algebraic version of regular neighbourhood theory. We describe an algebraic regular neighbourhood of a family of almost invariant subsets of a group G. This is a graph of groups structure for G, with the property that certain vertices enclose the given almost invariant sets. As splittings have almost invariant sets naturally associated, this also yields an idea of an algebraic regular neighbourhood of a family of splittings. In [**20**], Fujiwara and Papasoglu developed an idea of a regular neighbourhood of two splittings in special cases. However, their idea is not the same as ours as they concentrate on enclosing subgroups rather than splittings. In our algebraic construction of regular neighbourhoods, as well as several other techniques, we have greatly benefited from the two papers of Bowditch [**5**][**8**]. Bowditch's use of pretrees showed us how to enclose almost invariant sets under very general conditions. See our construction in chapter 3. In the case of word hyperbolic groups, Bowditch [**5**] was effectively the first to enclose such sets although he does not use this terminology. He also showed that the characteristic vertex groups of his graph of groups have the structure of finite-by-Fuchsian groups. Bowditch's techniques were further developed by Dunwoody and Swenson in [**15**] and Swenson [**54**]. Finally in [**8**], Bowditch extended these techniques still further to study simultaneously splittings of finitely presented groups over two-ended subgroups and enclosing groups for such subgroups. This paper is closest to our approach but there are some important differences. It turns out that the decompositions obtained by Bowditch in [**8**] differ from those which we obtain in this paper in chapters 9 and 10. We really want to enclose almost invariant sets (which are the algebraic analogue of immersions of codimension–1 manifolds) whereas Bowditch does not have a clear analogue of an immersion. Bowditch uses what he calls an 'axis' and this can be taken as an analogue of an immersion in some cases (see chapter 2 for more details). Secondly, the intersection number we studied in [**42**] and [**44**] seems to involve the right notion of crossing for almost invariant sets, whereas the notion of crossing used by Bowditch does not recognise what we call weak crossings. Taking care of these difficulties seems to make our decompositions more canonical and also corresponds better with the topological situation.

In order to explain the idea of an algebraic regular neighbourhood, we return to the characteristic submanifold $V(M)$ of a 3–manifold M and the graph of groups decomposition Γ of $G = \pi_1(M)$, whose underlying graph is dual to the frontier $\mathrm{fr}(V(M))$ of $V(M)$. Note that Γ is naturally a bipartite graph, because its vertices correspond to components of $V(M)$ or of $M - V(M)$, and each edge of Γ joins vertices of distinct types. The vertices which correspond to $V(M)$ will be called V_0–vertices and the vertices which correspond to $M - V(M)$ will be called V_1–vertices. Here are two properties of $V(M)$, which have algebraic analogues. The first is the Enclosing Property, which says that any essential annulus or torus in M is enclosed by $V(M)$. The second is that if F is any embedded essential closed surface in M, not necessarily a torus, and if F has intersection number zero with every essential annulus and torus in M, then F is homotopic into $M - V(M)$. These conditions are not sufficient to characterise $V(M)$ up to isotopy, but they do contain much of the information needed for such a characterisation. The algebraic analogue of the Enclosing Property is that the almost invariant subsets of G which correspond to essential annuli or tori are enclosed by the V_0–vertices of Γ. The algebraic analogue of the second property is that the splitting associated to F is enclosed by a V_1–vertex of Γ.

Now let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. Then our algebraic regular neighbourhood of the X_λ's in G is a bipartite graph of groups structure Γ for G such that the V_0–vertices of Γ enclose the X_λ's, and splittings of G which have intersection number zero with each X_λ are enclosed by the V_1–vertices of Γ. In addition, we need to insist that Γ is minimal in order to have any uniqueness results. There are two further technical conditions which we need to impose, and we discuss the details in chapter 6.

This paper is organised as follows. In chapter 1, we recall the basic properties of the characteristic submanifold of a Haken 3–manifold. In chapter 2, we recall some of the algebraic concepts and results that we need from our paper [**44**]. In chapters 3, 4, 5 and 6, we develop our general theory of regular neighbourhoods of families of almost invariant subsets of a group. We give a precise definition of a regular neighbourhood, and prove that if one exists then it is unique. We show that regular neighbourhoods exist for any finite family of almost invariant subsets each of which is over a finitely generated group, although they need not exist in general. This theory seems to be useful for studying splittings under more general conditions than those we consider for JSJ-decompositions. In chapter 6, we note a strengthening of a theorem of Niblo. The results of chapters 2-6 are very general and apply to almost invariant subsets of any finitely generated group.

In chapters 7 up to 10, we construct our first canonical decomposition. We restrict our attention to a one-ended, finitely presented group G and almost invariant subsets over two-ended subgroups of G. Our decomposition is a graph of groups structure for G which is a regular neighbourhood of all such almost invariant subsets of G. The restriction to finitely presented groups is natural because we use certain accessibility results. However it is now standard that the accessibility results that we use extend to almost finitely presented groups. Thus all the results of this paper which are stated under the assumption that G is finitely presented are valid when G is almost finitely presented (see [**4**]). In particular, the results are valid for Poincaré duality groups as these are known to be almost finitely presented. These results for almost finitely presented groups will be used in subsequent papers on the analogues of JSJ-decompositions for Poincaré duality pairs.

In chapter 7, we consider two-ended subgroups of G whose commensuriser in G is "small". We discuss the properties of a regular neighbourhood of finitely many almost invariant subsets over such subgroups of G. This means that we consider a graph of groups structure for G with certain vertices which enclose all these almost invariant subsets. When such almost invariant subsets cross strongly, the enclosing groups can be identified as finite-by-Fuchsian groups by the work of Bowditch and others. For weak crossings, it follows from our general theory of regular neighbourhoods that the enclosing groups themselves are two-ended. We formulate a slightly nonstandard accessibility result that we need for this argument.

In chapter 8, we consider two-ended subgroups of G whose commensuriser in G may be "large", and prove the following technical result. Let H be such a subgroup of G, and let $B(H)$ denote the Boolean algebra of all almost invariant subsets of G which are over subgroups commensurable with H. We show that $B(H)$ is finitely generated over the commensuriser of H in G. The proof depends on standard accessibility results and on techniques of Dunwoody and Roller [**13**].

Using the results of chapters 7 and 8, we obtain, in chapters 9 and 10, a regular neighbourhood of all the nontrivial almost invariant subsets of G which are over two-ended subgroups, i.e. a natural decomposition of any one-ended, finitely presented group G which encloses all such almost invariant subsets. Our uniqueness result for regular neighbourhoods implies that this decomposition is unique and also is invariant under automorphisms of G. The most remarkable point about this decomposition is that although the graph of groups structure obtained is finite, in general not all the vertex and edge groups will be finitely generated. We end chapter 10 by deducing the existence of a regular neighbourhood of all the splittings of G which are over two-ended subgroups. This may seem a more natural object, but it only seems possible to prove its existence by first considering the above regular neighbourhood of all almost invariant subsets of G which are over two-ended subgroups.

In chapter 11, we discuss several examples of this decomposition and compare it with the topological JSJ-decomposition for a 3–manifold and with Bowditch's decomposition in [**8**]. We show that when G is the fundamental group of a Haken manifold M, the enclosing groups obtained in chapter 10 essentially correspond to the peripheral characteristic submanifold. In particular, the decomposition of G essentially corresponds to the full characteristic submanifold when M is atoroidal.

In chapter 12, we generalise all the preceding results as follows. If G is a one-ended, finitely presented group which does not split over any VPC subgroup of length less than n, we construct a regular neighbourhood Γ_n of all the nontrivial almost invariant subsets of G which are over $VPCn$ subgroups of G, i.e. a natural decomposition of G which encloses all such almost invariant subsets. As a group is $VPC1$ if and only if it is two-ended, Γ_1 is exactly the decomposition of G constructed in chapter 10. For $n = 2$, if G is the fundamental group of a closed Haken 3–manifold M, then Γ_2 essentially yields the JSJ-decomposition of M. For $n \geqslant 2$, if G is a Poincaré duality group of dimension $n+1$, our construction of Γ_n recovers the results of Kropholler in [**27**]. Our results are slightly more general because they apply to all Poincaré duality groups, whereas Kropholler's results apply only to Poincaré duality groups such that any VPC subgroup has finitely generated centraliser. An example due to Mess [**33**] shows that this condition is not always satisfied. The results of this chapter also imply the existence of a regular neighbourhood of all the splittings of G which are over $VPCn$ subgroups.

In chapter 13, we use all the preceding ideas to construct an algebraic analogue of the whole characteristic submanifold of a 3–manifold. Given a one-ended, finitely presented group G, let E_k denote the collection of all the nontrivial almost invariant subsets of G over $VPCk$ subgroups of G. We start with the graph of groups structure Γ_1 for G which we gave in chapter 10. This is a regular neighbourhood of all the elements of E_1. Next we consider how to enclose elements of E_2. The closest analogue to the topology is obtained by considering only those elements of E_2 which do not cross any element of E_1. These are called 1–canonical. We show that Γ_1 can be refined to a graph of groups structure $\Gamma_{1,2}$ for G by adding new vertices which enclose all the 1–canonical elements of E_2. In the case when G is the fundamental group of an orientable Haken 3–manifold M, our work in [**45**] shows that this essentially corresponds to the topological decomposition of M. It may seem unsatisfactory that we do not find a decomposition of G with vertex groups which enclose all elements of E_1 and E_2, but our work in [**45**] shows that if such a decomposition exists it cannot be a refinement of Γ_1. We discuss this at the end of chapter 11. If G does not split over any VPC subgroups of length less than n, as in chapter 12, then we start with the graph of groups structure Γ_n for G which we gave in that chapter. This is a regular neighbourhood of all the elements of E_n. We say that an element of E_k is *n-canonical* if it crosses no element of E_i, for $i \leqslant n$. As in the case when $n = 1$, we show that Γ_n can be refined to a graph of groups structure $\Gamma_{n,n+1}$ for G by adding new vertices which enclose all the n–canonical elements of E_{n+1}. As in previous chapters, these results also imply the existence of a regular neighbourhood of all the splittings of G which are over $VPCn$ subgroups, and of all the n–canonical splittings of G which are over $VPC(n+1)$ subgroups.

In chapter 14, we discuss the natural question of whether one can continue in the same way. We want to refine $\Gamma_{1,2}$ to $\Gamma_{1,2,3}$ by adding new vertices which enclose all the 2–canonical elements of E_3. We show that this can be done if we restrict attention to virtually abelian subgroups of G. Letting A_k denote the corresponding subset of E_k, we show that this procedure can be repeated to obtain, for every n, a decomposition $\Gamma_{1,2,\dots,n}$ of G with vertices which enclose all the $(k-1)$–canonical elements of A_k, $1 \leqslant k \leqslant n$. In some cases, this sequence stabilises at some finite stage, so that we obtain a decomposition Γ_∞ of G with vertices which enclose all the $(k-1)$–canonical elements of A_k, for all $k \geqslant 1$. Surprisingly, this procedure does not work for VPC subgroups. We give a simple example to show that in general there is no decomposition analogous to $\Gamma_{1,2,3}$ for VPC subgroups of G. As in previous chapters, for each n, these results also imply the existence of a regular neighbourhood of all the $(k-1)$–canonical splittings of G which are over virtually abelian subgroups of length $k \leqslant n$.

In chapter 15, we discuss how the algebraic JSJ-decompositions of previous authors can be related to ours, and in chapter 16, we briefly discuss possible extensions of our results to more general classes of groups.

In Appendices A and B, at the suggestion of the referee, we include the complete text of our papers [**42**] and [**44**]. For the convenience of the reader, references to results in these papers will be made directly to the appropriate appendix.

Acknowledgments. — The first author acknowledges the support of the Hebrew University, the University of Melbourne, and the University of Warwick. His visits to Melbourne were supported by the ARC, and his visits to Warwick were supported by the EPSRC. The second author acknowledges the support of the University of Michigan. His visits to Michigan were supported by the NSF. Both authors acknowledge the support of the Université Paul Sabatier in Toulouse. They were supported by the CNRS while in Toulouse.

CHAPTER 1

THE CHARACTERISTIC SUBMANIFOLD

In this chapter, we will give a brief summary of the theory, emphasising those points which are closely related to the algebraic theory in this paper.

Let M be an orientable Haken 3–manifold, let F be a compact orientable surface, not S^2 or D^2, and let $f : F \to M$ be a map. Then f is *proper* if it sends ∂F into ∂M, and is *incompressible* if it is π_1–injective. A proper and incompressible map f is *essential* if it is not properly homotopic into ∂M. A codimension–0 submanifold of M is *incompressible* if each frontier component is incompressible in M. Finally, an incompressible codimension–0 submanifold W of M is *simple* in M if any essential map of the annulus or torus into M which has image in W can be properly homotoped into the frontier of W.

From now on we need to assume that M has incompressible boundary. Jaco and Shalen [**25**] and Johannson [**26**] proved that there is a family $\mathcal{T}$ of disjoint essential annuli and tori embedded in M, unique up to isotopy, and with the following properties. The manifolds obtained by cutting M along $\mathcal{T}$ are simple in M or are Seifert fibre spaces or I–bundles over surfaces. In fact, $\mathcal{T}$ can be characterised as the minimal family of annuli and tori with this property. The Seifert and I–bundle pieces of M are said to be *characteristic*, and any essential map of the annulus or torus into M can be properly homotoped to lie in a characteristic piece. This is called the Enclosing Property of $\mathcal{T}$. The characteristic submanifold $V(M)$ of M consists essentially of the union of the characteristic pieces of the manifold obtained from M by cutting along $\mathcal{T}$. However, if two characteristic pieces of M have a component S of $\mathcal{T}$ in common, we add a second copy of S to the family $\mathcal{T}$, thus separating the two characteristic pieces by a copy of $S \times I$, which we regard as a non-characteristic piece of M. Similarly, if two non-characteristic pieces of M have a component S of $\mathcal{T}$ in common, we add a second copy of S to the family $\mathcal{T}$, thus separating the two non-characteristic pieces by a copy of $S \times I$, which we regard as a characteristic piece of M. This is clearly needed if $V(M)$ is to have the Enclosing Property. Thus the frontier of $V(M)$

is usually not equal to $\mathcal{T}$. Some annuli or tori in $\mathcal{T}$ may appear twice in the frontier of $V(M)$. This discussion brings out a somewhat confusing fact about the characteristic submanifold, which is that both $V(M)$ and its complement can have components which are homeomorphic to $S \times I$, where S is an annulus or torus. One other basic point to note is that it is quite possible that $\mathcal{T}$ is empty. In this case, either $V(M)$ is equal to M or it is empty, so that either M is a Seifert fibre space or an I–bundle over a closed surface, or M admits no essential annuli and tori.

In order to complete this description of $V(M)$, we need to say a little more about its frontier. If W is a component of $V(M)$ which is an I–bundle over a compact surface F, then the frontier of W in M is the restriction of the bundle to ∂F, which is homeomorphic to $\partial F \times I$. If W is a Seifert fibre space component of $V(M)$, then there is a Seifert fibration on W such that the frontier of W in M consists of vertical annuli and tori.

An important point to make is that any such manifold can occur as a characteristic submanifold. More precisely, let W denote a compact 3–manifold such that each component is either an I–bundle or a Seifert fibre space over a compact surface with non-empty boundary. If V is a component of W which is an I–bundle over a surface F, let Σ_V denote the restriction of the bundle to ∂F. If V is a component of W which is a Seifert fibre space, let Σ_V denote a finite collection of disjoint vertical annuli and tori in ∂V. Let Σ denote the union of all the Σ_V's, and assume that Σ is incompressible in W. Then there is a Haken 3–manifold M with incompressible boundary whose characteristic submanifold $V(M)$ is homeomorphic to W and has frontier in M equal to Σ. Such a manifold M can be constructed by starting with W and gluing manifolds to Σ, but it is not trivial to prove that this can be done in the required way.

We emphasise that the Enclosing Property applies to any essential map of the annulus or torus into M and not just to essential embeddings. A related concept which is important for our approach to JSJ-decompositions is the idea of a canonical surface in M. The concept of canonical surface is not discussed in the original memoirs [**25**] and [**26**]. It emerged from [**32**] and [**34**] and is further developed in [**45**]. In [**45**], an embedded essential annulus or torus S in M is called *canonical* if any essential map of the annulus or torus into M can be properly homotoped to be disjoint from S. The Enclosing Property clearly implies that any annulus or torus in $\mathcal{T}$ is canonical. In [**45**], we showed that the family of isotopy classes of all canonical annuli and tori in M is equal to $\mathcal{T}$. In [**34**], Neumann and Swarup considered a slightly different version of this idea. They defined an embedded essential annulus or torus in M to be canonical if every *embedded* essential annulus or torus in M can be properly homotoped to be disjoint from it. We let $\mathcal{T}_e$ denote the family of isotopy classes of essential annuli and tori in M which are canonical in this sense. Clearly $\mathcal{T}$ is contained in $\mathcal{T}_e$. They showed that their family $\mathcal{T}_e$ is not, in general, the same as the family $\mathcal{T}$, but they were able to describe the differences and thereby give a new derivation of the classical JSJ-decomposition. Thus $\mathcal{T}_e$ determines a canonical decomposition of M which is finer than that determined by $\mathcal{T}$. Neumann and Swarup showed that $\mathcal{T}_e - \mathcal{T}$ consists of annuli only, which they call matched annuli. They list the possibilities for such annuli in Lemma 3.4 of [**34**]. However, their list is not quite correct. One case which they give occurs when M is the Seifert fibre space $W_{p,q}$ which is constructed by gluing two solid tori together along an annulus A which has degree p in one solid torus and degree q in the other solid torus, where $p, q \geqslant 2$. Thus $V(M)$ equals M in this case. They assert that A lies in $\mathcal{T}_e$. But this is not true when $p = q = 2$, although it is true for all other values of p and q. To see this, we start by showing that $W_{2,2}$ can also be viewed as the twisted I–bundle over the Klein bottle K. (Note that this is the unique I–bundle over K with orientable total space.) For recall that K contains a circle C which cuts K into two Moebius bands. The restriction of the I–bundle to each Moebius band is a solid torus, and the restriction of the I–bundle to C is an annulus which has degree 2 in each solid torus. Thus the twisted I–bundle over K is homeomorphic to $W_{2,2}$. Now K also has a non-separating two-sided simple closed curve D, and C and D cannot be homotoped apart. The restriction of the I–bundle to D is an annulus B in $W_{2,2}$, and it follows that A and B cannot be homotoped apart. Thus A does not lie in $\mathcal{T}_e$ in the case $p = q = 2$.

As discussed in the introduction, the guiding idea behind this paper is that $V(M)$ should be thought of as a regular neighbourhood of the family of all essential annuli and tori in M. By this we mean that every such map is properly homotopic into $V(M)$ and that $V(M)$ is minimal, up to isotopy, among all incompressible submanifolds of M with this property. It will be convenient to say that the collection of all such maps *fills* $V(M)$ when $V(M)$ has this minimality property. The word "fill" is used in the same way to describe certain curves on a surface. A subtle point which arises here is that there are exceptional cases where $V(M)$, as defined by Jaco-Shalen and Johannson, is not filled by the collection of all essential annuli and tori in M. In these cases our algebraic decomposition does not quite correspond to the topological JSJ-decomposition.

Let $V'(M)$ denote the incompressible submanifold of M which encloses every essential annulus and torus in M and is filled by them. We will see that $V'(M)$ is only slightly different from $V(M)$. The algebraic decomposition which we produce in this paper corresponds to $V'(M)$ rather than to $V(M)$. Clearly $V'(M)$ is a submanifold of $V(M)$, so that its frontier in M must also consist of essential annuli and tori in M. It follows that, as for $V(M)$, the isotopy classes of the frontier components of $V'(M)$ are precisely those of the canonical annuli and tori in M. Hence the difference between $V'(M)$ and $V(M)$ is essentially that certain exceptional components of $V(M)$ are discarded. Here is a description of the exceptional components, most of which are solid tori. A solid torus component W of $V(M)$ will fail to be filled by annuli essential in M when its frontier consists of 3 annuli each of degree 1 in W, or when its frontier consists of 1 annulus of degree 2 in W, or when its frontier consists of 1 annulus of degree 3 in W. Another exceptional case occurs when a component W of $V(M)$ lies in the interior of M and is homeomorphic to the manifold $W_{2,2}$ which, as discussed above, can also be viewed as the twisted I–bundle over the Klein bottle. Then any incompressible torus in W is homotopic into the boundary, so that W is not filled by tori which are essential in M. In all these cases, one obtains $V'(M)$ from $V(M)$ by first replacing W by a regular neighbourhood of its frontier in M, and then removing any redundant product components from the resulting submanifold.

We end this chapter by giving two examples, the first of which was mentioned in the introduction. These examples show that when enclosing one type of surface, the frontier of the enclosing submanifold may be of a different type. Let W be a connected Seifert fibre space with at least two boundary components, and choose Σ to consist of a single boundary component of W. Let M be a 3–manifold with characteristic submanifold W such that W has frontier in M equal to Σ. As W meets ∂M, it is the minimal submanifold of M which encloses all the essential annuli in M, but it has a frontier torus.

A different, but related, example can be constructed if we allow M to be non-orientable. Let F denote the compact surface obtained from a disc by removing the interiors of two disjoint discs. Let a and b denote the generators of $\pi_1(F)$ carried by two boundary components of F, oriented so that the third component carries ab. Thus ab^{-1} is carried by a figure eight loop γ in F. Let W denote the S^1–bundle over F with fundamental group the extension of $\mathbb{Z}$ by $\pi_1(F)$ in which a and b act on $\mathbb{Z}$ by inversion, and let Σ be the restriction of this bundle to ∂F. Then there is a 3–manifold M with characteristic submanifold W such that W has frontier in M equal to Σ. As F is a (topological) regular neighbourhood of the figure eight loop γ, it follows that W is a (topological) regular neighbourhood of the torus which lies above γ. Thus W is the minimal submanifold of M which encloses all the essential tori in M, but it has two frontier Klein bottles.

CHAPTER 2

PRELIMINARIES

We start by introducing the idea of an almost invariant subset of a finitely generated group G. Throughout this paper, we will always assume that G is finitely generated, but we will sometimes need to consider subgroups which are not finitely generated. We emphasise here that all the results of this chapter apply to the case when subgroups are not finitely generated, unless it is specifically stated that subgroups must be finitely generated. We will need several definitions which we take from [**44**]. See Appendix B, and see Appendix A for a discussion.

Definition 2.1. — Two sets P and Q are *almost equal* if their symmetric difference $P - Q \cup Q - P$ is finite. We write $P \overset{a}{=} Q$.

Definition 2.2. — If a group G acts on the right on a set Z, a subset P of Z is *almost invariant* if $Pg \overset{a}{=} P$ for all g in G. An almost invariant subset P of Z is *nontrivial* if P and its complement $Z - P$ are both infinite. The complement $Z - P$ will be denoted simply by P^*, when Z is clear from the context.

This idea is connected with the theory of ends of groups via the Cayley graph Γ of G with respect to some finite generating set of G. (Note that in this paper groups act on the left on covering spaces and, in particular, G acts on its Cayley graph on the left.) Using $\mathbb{Z}_2$ as coefficients, we can identify 0–cochains and 1–cochains on Γ with sets of vertices or edges. A subset P of G represents a set of vertices of Γ which we also denote by P, and it is a beautiful fact, due to Cohen [**9**], that P is an almost invariant subset of G if and only if δP is finite, where δ is the coboundary operator in Γ. Thus G has a nontrivial almost invariant subset if and only if the number of ends $e(G)$ of G is at least 2. Further $e(G)$ can be identified with the number of nontrivial almost invariant subsets of G, when this count is made correctly. If H is a subgroup of G, we let $H\backslash G$ denote the set of cosets Hg of H in G, i.e. the quotient of G by the left action of H. Of course, G will no longer act on the left on this quotient, but it will still act on the right. Thus we have the idea of an almost invariant subset of $H\backslash G$. Further, P is an almost invariant subset of $H\backslash G$ if and only if δP is finite, where δ is the coboundary operator in the graph $H\backslash\Gamma$. Thus $H\backslash G$ has a nontrivial almost invariant subset if and only if the number of ends $e(G,H)$ of the pair (G,H) is at least 2. Considering the pre-image X in G of an almost invariant subset P of $H\backslash G$ leads to the following definitions.

Definition 2.3. — If G is a finitely generated group and H is a subgroup, then a subset X of G is H*–almost invariant* if X is invariant under the left action of H, and simultaneously $H\backslash X$ is an almost invariant subset of $H\backslash G$. We may also say that X is *almost invariant over* H. In addition, X is a *nontrivial* H–almost invariant subset of G, if the quotient sets $H\backslash X$ and $H\backslash X^*$ are both infinite.

Remark 2.4. — Note that if X is a nontrivial H–almost invariant subset of G, then $e(G,H)$ is at least 2, as $H\backslash X$ is a nontrivial almost invariant subset of $H\backslash G$. In fact $e(G,H)$ can be identified with the number of nontrivial H–almost invariant subsets of G, when this count is made correctly. See [**47**] for details.

In [**30**], Kropholler and Roller gave a different definition of a H–almost invariant subset of a group, which they used in several later papers. We will discuss this at the end of this chapter.

Definition 2.5. — If G is a group and H is a subgroup, then a subset W of G is H*–finite* if it is contained in the union of finitely many left cosets Hg of H in G.

Definition 2.6. — If G is a group and H is a subgroup, then two subsets V and W of G are H*–almost equal* if their symmetric difference is H–finite.

It will also be convenient to avoid this rather clumsy terminology sometimes, particularly when the group H is not fixed, so we make the following definitions.

Definition 2.7. — Two subgroups H and K of a group G are *commensurable* if $H\cap K$ has finite index in both H and K.

Definition 2.8. — If X is a H–almost invariant subset of G and Y is a K–almost invariant subset of G, and if X and Y are H–almost equal, then we will say that X and Y are *equivalent* and write $X\sim Y$.

Remark 2.9. — Note that H and K must be commensurable, so that X and Y are also K–almost equal and $(H\cap K)$–almost equal.

A more elegant and equivalent formulation is that X is equivalent to Y if and only if each is contained in a bounded neighbourhood of the other.

Equivalence is important because usually one is interested in an equivalence class of almost invariant subsets of a group rather than a specific such subset.

A *splitting* of a group G is an expression of G as an amalgamated free product $A*_C B$, where $A\neq C\neq B$, or as a HNN extension $A*_C$. Thus a splitting of G always describes a nontrivial decomposition. If one thinks of a splitting of a group as an algebraic analogue of the topological notion of an embedded π_1–injective, codimension–1 and two-sided submanifold, then almost invariant sets should be thought of as analogues of immersions of such manifolds. We can describe the connection between these ideas as follows. Let M be a closed manifold with fundamental group G and consider a codimension–1 two-sided manifold S immersed in M such that the induced map of fundamental groups is injective. Let H denote the image of $\pi_1(S)$ in G, and let M_H denote the cover of M such that $\pi_1(M_H) = H$. For simplicity, we will assume that the lift of S to M_H is an embedding, whose image we will also denote by S. Then S must separate M_H, and we let A denote the closure of one side of S in M_H. Let $\widetilde{M}$ denote the universal cover of M, and let $\widetilde{S}$ and $\widetilde{A}$ denote the pre-images in $\widetilde{M}$ of the submanifolds S and A of M_S. Thus $\widetilde{S}$ is a copy of the universal cover of S. Next pick a generating set for G and represent it by a bouquet of circles embedded in M. We will assume that the wedge point of the bouquet does not lie on the image of S. The pre-image of this bouquet in $\widetilde{M}$ will be a copy of the Cayley graph Γ of G with respect to the chosen generating set. The pre-image in M_H of the bouquet will be a copy of the graph $H\backslash\Gamma$. Let P denote the set of all vertices of $H\backslash\Gamma$ which lie in A. Then P has finite coboundary, as δP equals exactly those edges of $H\backslash\Gamma$ which cross S. Hence P is an almost invariant subset of $H\backslash G$. If X denotes the set of vertices in Γ which lie in $\widetilde{A}$, then X is the pre-image of P and so is a H–almost invariant subset of G. Note that if we replace A by its complement, then P is replaced by P^*. If we choose a different generating set for G or a different embedding of the bouquet of circles in M, the H–almost invariant subsets P and P^* will change, but the new sets will be equivalent to P or to P^*. Thus we have associated a H–almost invariant set to the given immersion S and this set is unique up to equivalence and complementation.

If one has a splitting of a group G over a subgroup H, the above discussion leads to a natural way to associate to this splitting a standard H–almost invariant subset of G which is essentially unique (up to complementation). But it is simpler and clearer to work in a more combinatorial setting. For this we recall the basic result of Bass-Serre theory [**50**]. (See [**51**] for an English translation.) This tells us that an expression of a group G as $A *_H B$ or as $A*_H$ is equivalent to an action of G on a tree T without inversions so that $G\backslash T$ has a single edge and the edge stabilisers are conjugates of H. (Throughout this paper, we will only consider G–trees on which G acts without inversions, i.e. if an element of G preserves an edge, it also fixes that edge pointwise. This is only a very minor restriction. For if G acts on a tree T with inversions, let T' denote the tree obtained from T by dividing each edge into two edges. There is a natural action of G induced on T' which clearly has no inversions.)

First we consider a general action without inversions of a group G on a tree T. Recall that there is a natural partial order on the oriented edges of a tree T, given by saying that if s and t are oriented edges of T, then $s \leqslant t$ if and only if there is an oriented path in T which starts with s and ends with t. For any action without inversions of a group G on a tree T, and any edge s of T, we have a natural partition of G into two sets $X_s = \{g : gs \geqslant s \text{ or } g\overline{s} \geqslant s\}$ and $X_s^* = \{g : gs < s \text{ or } g\overline{s} < s\}$. We will show below that if S denotes the stabiliser of s, then X_s and X_s^* are both S–almost invariant. Although this partition of G is natural, it is not quite right for our purposes, because it is not equivariant under the action of G on T. In fact, if t denotes the edge ks of T, then X_t is equal to kX_sk^{-1}, whereas we would like it to be equal to kX_s. We resolve this problem in the following way. We choose a vertex w of T to be a basepoint. Next we define $\varphi : G \to V(T)$ by the formula $\varphi(g) = gw$. Clearly φ is G–equivariant, by which we mean that $k\varphi(g) = \varphi(kg)$, for all elements g and k of G. Note that $\varphi(e) = w$. Now an oriented edge s of T determines a natural partition of $V(T)$ into two sets, namely the vertices of the two subtrees obtained by removing the interior of s from T. Let Y_s denote the collection of all the vertices of the subtree which contains the terminal vertex v of s, and let Y_s^* denote the complementary collection of vertices. Then s determines a natural (in terms of our choice of the basepoint w) partition of G into two sets, namely $Z_s = \varphi^{-1}(Y_s)$ and $Z_s^* = \varphi^{-1}(Y_s^*)$. Clearly, these sets are equivariant, i.e. if t denotes the edge ks of T, then Z_t is equal to kZ_s. We will show that if S denotes the stabiliser of s, then Z_s and Z_s^* are both S–almost invariant. Further, Z_s is S–almost equal to the set X_s defined above. It follows that although Z_s depends on our choice of basepoint, the equivalence class of Z_s is independent of this choice.

Lemma 2.10. — *Let T be a G–tree, and let s be an oriented edge of T with stabiliser S.*

(1) *Then the subset $X_s = \{g : gs \geqslant s \text{ or } g\overline{s} \geqslant s\}$ of G is S–almost invariant.*

(2) *Let w be a basepoint for T, and define Z_s as above. Then Z_s is S–almost invariant and is S–almost equal to X_s. If w equals the terminal vertex v of s, then $Z_s = X_s$.*

Remark 2.11. — Note that this result does not require that G be finitely generated.

Proof. — 1) We need to show that $hX_s = X_s$, for all h in S, and that X_sh and X_s are S–almost equal for all h in G.

If $gs \geqslant s$, then $hgs \geqslant hs$ which equals s, for all h in S. Similarly if $g\overline{s} \geqslant s$, then $hg\overline{s} \geqslant hs = s$, for all h in S. Thus $hX_s \subset X_s$, for all h in S. Hence $h^{-1}X_s \subset X_s$, for all h in S, so that $hX_s = X_s$, for all h in S, as required.

Now consider an element k of $X_sh - X_s$. Thus there is g in X_s such that $k = gh$ does not lie in X_s. This means that the edge s lies between the edges gs and $ks = ghs$. Applying g^{-1}, we see that the edge $g^{-1}s$ lies between s and hs. Thus $g^{-1}s$ is one of the finitely many edges of T between s and hs, so that g^{-1} lies in some finite union of cosets g_iS of S in G. Hence g lies in the union of cosets Sg_i^{-1} and so k also lies in some finite union of cosets Sk_i.

Next consider an element k of $X_s - X_s h$. Thus k lies in X_s and kh^{-1} does not. Hence s lies between ks and $kh^{-1}s$ so that $k^{-1}s$ lies between s and $h^{-1}s$. As in the preceding paragraph, it follows that k lies in some finite union of cosets Sk_j of S.

It follows from the previous two paragraphs that $X_s h$ and X_s are S–almost equal for all h in G, as required.

2) Recall that $\varphi : G \to V(T)$ is defined by the formula $\varphi(g) = gw$, and that Z_s denotes $\varphi^{-1}(Y_s)$. It follows that $Z_s = \{g \in G : gw \in Y_s\}$. Now recall that v denotes the terminal vertex of s. It is easy to check that $X_s = \{g \in G : gv \in Y_s\}$, so that $Z_s = X_s$ when w equals v.

Let k be an element of $Z_s - X_s$. Thus $kw \in Y_s$ and $kv \notin Y_s$. Hence s lies between kv and kw, so that $k^{-1}s$ lies between v and w. As in part 1), it follows that k lies in some finite union of cosets Sk_i of S. Similarly, $X_s - Z_s$ is contained in a finite union of cosets of S, so that Z_s is S–almost equal to X_s, as required. $\square$

In terms of the above discussion, there is now an easy and natural way to associate a H–almost invariant subset of G to a splitting σ of G over H. Given σ, let T be the associated G–tree, and let s denote an oriented edge of T with stabiliser H. Choose the basepoint w of T to be an endpoint of s, and then take the H–almost invariant subset Z_s or its complement. This description involves three choices, namely the choices of the edge s, its orientation and the choice of an endpoint of s. The choice of s will alter the almost invariant sets obtained by a translation, which corresponds to a conjugation of the splitting. Once this choice is made, we end up with precisely four H–almost invariant subsets of G which are naturally associated to the given splitting over H. In [**44**], we gave a different, but equivalent, description of these four sets and called them the standard H–almost invariant subsets of G associated to σ. Each of these sets has a particularly nice property, which will play an important role later on. If X denotes any one of these four sets, then $X = \{g \in G : gX^{(*)} \subset X\}$. Here we use $X^{(*)}$ to denote a set which might be X or X^*, so that $gX^{(*)} \subset X$ is shorthand for $gX \subset X$ or $gX^* \subset X$. To see why this is true, consider the case when w is the terminal vertex of s and $X = Z_s$. Note that it is never possible to have the equation $g\overline{s} = s$, as G acts on T without inversions. Thus

$$\begin{aligned} X = Z_s = \varphi^{-1}(Y_s) &= \{g \in G : gw \in Y_s\} \\ &= \{g \in G : gs \geqslant s \text{ or } g\overline{s} \geqslant s\} = \{g \in G : gX^{(*)} \subset X\}. \end{aligned}$$

An interesting special case occurs when the given splitting of G is not a HNN extension, and one of the vertex groups, say B, contains H with index 2. This implies that the vertex of s with stabiliser B has valence 2. If we choose the other vertex of s as our basepoint, and let b denote any element of $B - H$, then the resulting H–almost invariant subset Z_s satisfies the equation $bZ_s = Z_s^*$. This is a situation which we will want to avoid for reasons which will become clear later. We will show that it is a fairly unusual phenomenon. The following introduces the terminology we want.

Definition 2.12. — If X is an H–almost invariant subset of a group G associated to a splitting of G over H, then X is *invertible* if there is an element g in G such that $gX = X^*$.

In order to understand when such sets can be invertible, we first need the following result.

Lemma 2.13. — *Let G be a group with subgroups H and K. Suppose that Xg is K–almost equal to X for all g in G, and that X is H–finite. Then either X is K–finite or H has finite index in G.*

Remark 2.14. — We do not assume that $KX = X$, so that X need not be K–almost invariant.

Proof. — Suppose that X is K–infinite. As X is H–finite, it is contained in a finite union of left cosets Hg. Hence there must be a coset Hl such that X meets Hl in a K–infinite set. Now, for any element g of G, it follows that $Xl^{-1}g$ meets $Hl(l^{-1}g) = Hg$ in a K–infinite set. As $Xl^{-1}g$ is K–almost equal to X, it follows that X itself meets every coset Hg in a K–infinite set. As we are assuming that X is contained in a finite union of such cosets, it follows that this finite collection of cosets contains every coset Hg, and so H has finite index in G, as required. □

Using this result, we can now show that it is very unusual for almost invariant sets to be invertible. In fact, we show the stronger result that if X is a H–almost invariant set associated to a splitting of a group G over H, then it is very unusual for there to be an element g in G such that gX is equivalent to X^*.

Lemma 2.15. — *Let σ denote a splitting of a group G over a subgroup H. Let X denote a H–almost invariant subset of G associated to σ and suppose there is an element g in G such that gX is equivalent to X^*. Then σ is not a HNN extension, g must lie in a conjugate of a vertex group B of σ, and B must contain H with index 2.*

Remark 2.16. — A proof of this result could be disentangled from the proof of Lemma B.2.3, but the argument we give here seems clearer.

Proof. — The splitting σ of G determines an action of G on a tree T with quotient a graph Γ with a single edge. Let s be an edge of T with stabiliser H. As gX is equivalent to X^*, it follows that gZ_s is equivalent to Z_s^*. In particular, as gZ_s is not equal to Z_s, it follows that g cannot preserve s.

We claim that every interior vertex of the path λ which joins s and gs must have valence 2. For suppose that v is a vertex of the interior of λ of valence at least 3, let l and m denote the two edges of λ which are incident to v and let n denote a third edge incident to v. We will obtain a contradiction. By reversing the orientation of s if needed, we can assume that s is oriented into λ. It follows that gs is also oriented into λ, for otherwise we would have $gZ_s \subset Z_s$ which contradicts the fact that gZ_s is equivalent to Z_s^*. We now choose l, m and n to be oriented towards v, and choose our notation for l and m so that the path from s to v contains l. Thus $Z_s^* \subset Z_l^*$, $Z_m \subset Z_{gs}$, and $Z_m \supset Z_l^* \cup Z_n^*$. Hence $Z_s^* \subset Z_l^* \subset Z_l^* \cup Z_n^* \subset Z_m \subset Z_{gs} = gZ_s$. As gZ_s is equivalent to Z_s^*, each of the preceding inclusions is an equivalence. As Z_l^* and Z_n^* are disjoint, it follows that Z_n^* is H–finite. But n is a translate of s, so that Z_n^* is K–infinite where K denotes the stabiliser of n. Now Lemma 2.13 implies that H must have finite index in G, which is the required contradiction.

If every vertex of T has valence 2, then T is a line which must be fixed by H. Thus H is normal in G and the quotient G/H acts on T by translations and reflections. As gZ_s is equivalent to Z_s^*, the action of g cannot be a translation. It follows that G/H is isomorphic to $\mathbb{Z}_2 * \mathbb{Z}_2$, and that σ is of the form $G = A *_H B$, where both A and B contain H with index 2. Further, as g acts on T by a reflection, it must lie in a conjugate of A or B.

If T has vertices of valence greater than 2, then every edge of T must have such a vertex. Let v denote the vertex of s of valence greater than 2, and let w denote the vertex of s of valence 2. Then gs must intersect s in w, and g must stabilise w. As the valences of v and w are distinct, it follows that Γ has two vertices, so that σ is not a HNN extension. As $\mathrm{Stab}(w)$ contains H with index at most 2, and as σ is a splitting of G, it follows that the index must be exactly 2. The result follows. □

This lemma shows that invertible almost invariant sets can only occur in the situation when $G = A *_H B$, and B contains H with index 2, which is the case we discussed just before Definition 2.12. In particular, invertible almost invariant sets do exist for any splitting of this type. However, even in this situation it is clear that "most" of the H–almost invariant subsets of G associated to the splitting are not invertible. For if X is such a set which is invertible, and if y denotes any element of X^*, then the set $X \cup Hy$ is H–almost invariant, is equivalent to X and is clearly not invertible. Note that if H has index greater than 2 in A, the other pair of standard H–almost invariant subsets of G associated to the given splitting are not invertible. But if H has index 2 in both A and B, then all four of the standard H–almost invariant subsets of G are invertible. In topological terms, the non-invertible case corresponds to considering a two-sided simple closed curve C on a surface, and the invertible case can arise only when C bounds a Moebius band and it corresponds to replacing C by the core of that Moebius band.

The next definition makes precise the notion of crossing of almost invariant sets. This is an algebraic analogue of crossing of codimension–1 manifolds, but it ignores "inessential" crossings.

Definition 2.17. — Let X be a H–almost invariant subset of G and let Y be a K–almost invariant subset of G. We will say that Y *crosses* X if each of the four sets $X \cap Y$, $X^* \cap Y$, $X \cap Y^*$ and $X^* \cap Y^*$ is H–infinite. Thus each of the four sets projects to an infinite subset of $H \backslash G$.

The motivation for the above definition is that when one of the four sets is empty, we clearly have no crossing, and if one of the four sets is "small" (see Definition 2.19 below), then we have "inessential crossing". Note that Y may be a translate of X in which case such crossing corresponds to the self-intersection of a single immersion.

Remark 2.18. — It is shown in Lemma A.2.3 that if X and Y are nontrivial, then $X \cap Y$ is H–finite if and only if it is K–finite. It follows that crossing of nontrivial almost invariant subsets of G is symmetric, i.e. that X crosses Y if and only if Y crosses X. We will often write $X^{(*)} \cap Y^{(*)}$ instead of listing the four sets $X \cap Y$, $X^* \cap Y$, $X \cap Y^*$ and $X^* \cap Y^*$.

Definition 2.19. — Let U be a nontrivial H–almost invariant subset of G and let V be a nontrivial K–almost invariant subset of G. We will say that $U \cap V$ is *small* if it is H–finite.

Remark 2.20. — This terminology will be extremely convenient, particularly when we want to discuss translates U and V of X and Y, as we do not need to mention the stabilisers of U or of V. However, the terminology is symmetric in U and V and makes no reference to H or K, whereas the definition is not symmetric and does refer to H, so some justification is required. If U is also H'–almost invariant for a subgroup H' of G, then H' must be commensurable with H. Thus $U \cap V$ is H–finite if and only if it is H'–finite. In addition, Remark 2.18 tells us that $U \cap V$ is H–finite if and only if it is K–finite. This provides the needed justification of our terminology.

The term crossing has often been used in the literature for a somewhat different concept which we call strong crossing. As the name suggests, strong crossing implies crossing but the converse need not be true. We will now define this notion. Let G be a finitely generated group and let H and K be subgroups of G. Let X be a nontrivial H–almost invariant subset of G and let Y be a nontrivial K–almost invariant subset of G. It will be convenient to think of δX as a set of edges in Γ or as a set of points in G, where the set of points will simply be the collection of endpoints of all the edges of δX.

Definition 2.21. — We say that Y *crosses* X *strongly* if both $\delta Y \cap X$ and $\delta Y \cap X^*$ project to infinite sets in $H \backslash G$. If Y crosses X but not strongly, we say that Y *crosses* X *weakly*.

Remark 2.22. — These definitions are independent of the choice of generators for G which is used to define Γ. Clearly, if Y crosses X strongly, then Y crosses X. Note that Y does not cross X strongly if and only if δY is contained in a bounded neighbourhood of X or X^*.

An interesting point about strong crossing of X and Y is that it depends only on the subgroups H and K. More precisely, we have the following result.

Lemma 2.23. — *Let G be a finitely generated group and let H and K be subgroups of G. Let X be a nontrivial H–almost invariant subset of G and let Y and Y' be nontrivial K–almost invariant subsets of G. Then Y crosses X strongly if and only if Y' crosses X strongly.*

Proof. — By Remark 2.22, Y does not cross X strongly if and only if δY is contained in a bounded neighbourhood of X or X^*. As $H\backslash\delta Y$ and $H\backslash\delta Y'$ are both finite, it follows that δY is contained in a bounded neighbourhood of $\delta Y'$ and vice versa. Thus Y crosses X strongly if and only if Y' crosses X strongly. □

Definition 2.24. — Let σ_1 and σ_2 be splittings of G over C_1 and C_2, and let X_i be one of the standard C_i–almost invariant subsets of G associated to the splitting σ_i, for $i = 1, 2$. Then σ_1 *crosses* σ_2 *(strongly)* if there is g in G such that X_1 crosses gX_2 (strongly).

Remark 2.25. — As the standard C_i–almost invariant subsets of G associated to the splitting σ_i are all C_i–almost equal or C_i–almost complementary, this definition does not depend on the choice of the X_i's.

Next we give a simple example to show that strong crossing is not symmetric, in general.

Example 2.26. — Consider an essential two-sided simple closed curve S on a compact surface F which intersects a simple arc L transversely in a single point. Let G denote $\pi_1(F)$, and let H and K respectively denote the subgroups of G carried by S and L, so that H is infinite cyclic and K is trivial. Then S and L each define a splitting of G over H and K respectively. Let X and Y denote associated standard H–almost invariant and K–almost invariant subsets of G. These correspond to submanifolds of the universal cover of F bounded respectively by a line $\widetilde{S}$ lying above S and by a compact interval $\widetilde{L}$ lying above L, such that $\widetilde{S}$ meets $\widetilde{L}$ transversely in a single point. Clearly, X crosses Y strongly but Y does not cross X strongly.

If σ_1 and σ_2 are splittings of G over finitely generated subgroups C_1 and C_2 respectively, Sela introduced in [**49**] the following notion of crossing of σ_1 and σ_2. He says that σ_1 is *hyperbolic with respect to* σ_2 if C_1 is not conjugate into a vertex group of the splitting σ_2. It is easy to show that this idea is the same as strong crossing, and we give the proof below.

Lemma 2.27. — *If σ_1 and σ_2 are splittings of a finitely generated group G over finitely generated subgroups, then σ_1 is hyperbolic with respect to σ_2 if and only if σ_1 crosses σ_2 strongly.*

Proof. — Consider the G–tree T_2 corresponding to the splitting σ_2 and let the amalgamating group of the splitting σ_i be H_i. Consider the action of H_1 on T_2. It is immediate from the definition that σ_1 is hyperbolic with respect to σ_2 if and only if H_1 does not fix a vertex of T_2. As we are assuming that H_1 is finitely generated, it fixes some vertex of T_2 if and only if every element of H_1 fixes some vertex of T_2. Thus σ_1 is hyperbolic with respect to σ_2 if and only if there is an element h of H_1 which fixes no vertex of T_2, and thus has an axis l. We claim that this implies that σ_1 crosses σ_2 strongly. Let X_i denote a H_i–almost invariant subset of G associated to the splitting σ_i as discussed just after Lemma 2.10. As the quotient $H_i \backslash \delta X_i$ is finite, it follows that δX_i must lie within a bounded distance of H_i. Now let s denote an oriented edge of T_2 with stabiliser H_2, and choose the terminal vertex w of s to be the basepoint of T, so that $\varphi(g) = gw$. Then $\varphi(\delta X_2) = \partial s$. Now some translate gs of s lies on the axis l of the element h of H_1. It follows that $\varphi(H_1)$ contains points arbitrarily far from gs and on each side of gs. Hence X_1 crosses gX_2 strongly, so that σ_1 crosses σ_2 strongly as claimed. On the other hand, if σ_1 is not hyperbolic with respect to σ_2, then H_1 fixes a vertex v of T_2. It follows that $\varphi(\delta X_1)$ must lie within a bounded distance of v, and hence must lie within a bounded distance of gs, for every g in G. Thus δX_1 lies within a bounded distance of gX_2 and gX_2^*, for every g in G, so that σ_1 cannot cross σ_2 strongly. □

Remark 2.28. — A key point in the above argument is that if every element of H_1 fixes some vertex of T_2, then H_1 itself fixes some vertex. If H_1 is not finitely generated, this can fail, but a result of Tits, Lemma 3.4 of [**55**], can be used instead to show that the above lemma holds for splittings over infinitely generated subgroups.

Sela showed in [**49**] that if C_1 and C_2 are two-ended and if G does not split over a finite group, then his crossing, and hence our strong crossing, is symmetric.

The following two technical results play an important role in the theory of almost invariant sets. They are well known to experts. We will need the following terminology.

Definition 2.29. — Two subsets U and V of a set G are called *nested* if one of the four sets $U^{(*)} \cap V^{(*)}$ is empty.

Lemma 2.30. — *Let G be a finitely generated group with finitely generated subgroups H and K, a nontrivial H–almost invariant subset X and a nontrivial K–almost invariant subset Y. Then $\{g \in G : gX \text{ and } Y \text{ are not nested}\}$ consists of a finite number of double cosets KgH.*

Proof. — Let Γ denote the Cayley graph of G with respect to some finite generating set for G. Let P denote the almost invariant subset $H \backslash X$ of $H \backslash G$ and let Q denote the almost invariant subset $K \backslash Y$ of $K \backslash G$. Recall from the start of this chapter, that if we identify P with the 0–cochain on $H \backslash \Gamma$ whose support is P, then P is an almost invariant subset of $H \backslash G$ if and only if δP is finite. Thus δP is a finite collection of edges in $H \backslash \Gamma$ and similarly δQ is a finite collection of edges in $K \backslash \Gamma$. Now let C denote a finite connected subgraph of $H \backslash \Gamma$ such that C contains δP and the natural map

$\pi_1(C) \to H$ is onto, and let E denote a finite connected subgraph of $K\backslash\Gamma$ such that E contains δQ and the natural map $\pi_1(E) \to K$ is onto. Thus the pre-image D of C in Γ is connected and contains δX, and the pre-image F of E in Γ is connected and contains δY. Let Δ denote a finite subgraph of D which projects onto C, and let Φ denote a finite subgraph of F which projects onto E. If gD meets F, there must be elements h and k in H and K such that $gh\Delta$ meets $k\Phi$. Now $\{\gamma \in G : \gamma\Delta \text{ meets } \Phi\}$ is finite, as G acts freely on Γ. It follows that $\{g \in G : gD \text{ meets } F\}$ consists of a finite number of double cosets KgH.

The result would now be trivial if X and Y were each the vertex set of a connected subgraph of Γ. As this need not be the case, we need to make a careful argument as in the proof of Lemma 5.10 of [**47**]. Consider g in G such that gD and F are disjoint. We will show that gX and Y are nested. As D is connected, the vertex set of gD must lie entirely in Y or entirely in Y^*. Suppose that the vertex set of gD lies in Y. For a set S of vertices of Γ, let $\overline{S}$ denote the maximal subgraph of Γ with vertex set equal to S. Each component W of $\overline{X}$ and $\overline{X^*}$ contains a vertex of D. Hence gW contains a vertex of gD and so must meet Y. If gW also meets Y^*, then it must meet F. But as F is connected and disjoint from gD, it lies in a single component gW. It follows that there is exactly one component gW of $\overline{gX}$ and $\overline{gX^*}$ which meets Y^*, so that we must have $gX \subset Y$ or $gX^* \subset Y$. Similarly, if gD lies in Y^*, we will find that $gX \subset Y^*$ or $gX^* \subset Y^*$. It follows that in either case gX and Y are nested as required. □

A special case of the following result is proved in [**44**], but is not formulated as a separate statement, so we give here a brief description of the proof.

Lemma 2.31. — *Let G be a finitely generated group with finitely generated subgroups H and K, a nontrivial H–almost invariant subset A and a nontrivial K–almost invariant subset U. Then $\{g \in G : gU^{(*)} \leqslant A\}$ is contained in a bounded neighbourhood of A in the Cayley graph of G.*

Proof. — Lemma 2.30 tells us that $\{g \in G : gU \text{ and } A \text{ are not nested}\}$ consists of a finite number of double cosets HgK. Now suppose that $gU \leqslant A$ but gU is not contained in A. As $gU \cap A^*$ is small, i.e. it projects to a finite subset in $H\backslash G$, it follows that gU is contained in a bounded neighbourhood of A. The fact that g must lie in a finite number of double cosets HgK implies that there is a uniform bound on the size of this neighbourhood of A. This means that there is a number d such that for all g such that $gU^{(*)} \leqslant A$, we have $gU^{(*)}$ is contained in a d–neighbourhood of A. Now let W denote $\{g \in G : gU^{(*)} \leqslant A\}$. The preceding discussion shows that if $g \in W$, then $g\delta U$ lies in the $(d+1)$–neighbourhood of A. If we let c denote the distance of the identity of G from δU, then g must lie in the $(c+d+1)$–neighbourhood of A, so that W lies in the $(c+d+1)$-neighbourhood of A, as required. □

Now we come to the definition of the intersection number of two almost invariant sets.

Definition 2.32. — Let H and K be subgroups of a finitely generated group G. Let P denote a nontrivial almost invariant subset of $H\backslash G$, let Q denote a nontrivial almost invariant subset of $K\backslash G$ and let X and Y denote the pre-images of P and Q respectively in G. Then the *intersection number* $i(P,Q)$ of P and Q equals the number of double cosets KgH such that gX crosses Y.

Remark 2.33. — The following facts about intersection numbers are proved in Lemmas A.2.7 and A.2.8.

(1) Intersection numbers are symmetric, i.e. $i(P,Q) = i(Q,P)$.

(2) $i(P,Q)$ is finite when G, H and K are all finitely generated. This follows immediately from Lemma 2.30.

(3) If P' is an almost invariant subset of $H\backslash G$ which is almost equal to P or to P^* and if Q' is an almost invariant subset of $K\backslash G$ which is almost equal to Q or to Q^*, then $i(P',Q') = i(P,Q)$.

When [**42**] was written, we knew of no examples of almost invariant sets with infinite intersection number, so we are grateful to Vincent Guirardel for showing us the following example. This demonstrates that $i(P,Q)$ may be infinite if one of H or K is not finitely generated.

Example 2.34. — Let G denote the free group of rank 2 and let C denote any subgroup of G which is not finitely generated. It is easy to show that, as C has infinite index in G, the number of ends of the pair (G,C) is infinite. Let σ denote any splitting of G over the trivial subgroup K. There are only two such splittings up to isomorphism of G. One is HNN, the other is not, and in both cases the vertex groups are infinite cyclic. This last fact is the key to the simple argument which follows.

Let X denote any nontrivial C–almost invariant subset of G, and let Y denote an almost invariant subset of G associated to the splitting σ. Then $i(X,Y)$ is infinite.

To see this, let T denote the G–tree associated to the splitting σ of G, pick a basepoint for T and let $\varphi : G \to T$ be a G–equivariant map. Also let s denote an edge of T such that Z_s is equivalent to Y. Now $C\backslash T$ naturally yields a graph of groups structure for C in which every edge group is trivial, and every vertex group is cyclic. As C is not finitely generated, it follows that there are infinitely many edges of $C\backslash T$ which determine a nontrivial splitting of C. If t denotes an edge of T which projects to such an edge of $C\backslash T$, it immediately follows that $\varphi(C)$ contains points arbitrarily far from t and on each side of t. If $t = gs$, this implies that X crosses gY strongly. If t and t' are two such edges, they have the same image in $C\backslash T$ if and only if $ct = t'$, for some c in C. It follows that the number of double cosets CgK such that X crosses gY is infinite.

One can also define the intersection number of two splittings to be the intersection number of the almost invariant sets associated to the splittings. Now if two curves on a surface have intersection number zero, they can be isotoped to be disjoint. There is a natural algebraic analogue of this fact. We define a collection of n splittings of a group G to be *compatible* if G can be expressed as the fundamental group of a graph of groups with n edges, such that the edge splittings of the graph are conjugate to the given splittings. The following result is a slight rewording of Theorem B.2.5.

Theorem 2.35. — *Let G be a finitely generated group with n splittings over finitely generated subgroups. Then the splittings are compatible if and only if each pair of splittings has intersection number zero. Further, in this situation, the graph of groups structure on G obtained from these splittings is unique up to isomorphism.*

In [**2**], Bass gives a discussion of isomorphisms of graphs of groups. In particular, if two graphs of groups structures on a group G are isomorphic, they have isomorphic underlying graph, and edges and vertices which correspond under the isomorphism carry conjugate subgroups of G.

This discussion of intersection numbers leads naturally to the concept which we call a canonical splitting. Recall from chapter 1 that an embedded essential annulus or torus F in an orientable 3–manifold M is called canonical if any essential map of the annulus or torus into M can be properly homotoped to be disjoint from F. An equivalent formulation of this definition is that an embedded essential annulus or torus F in an orientable 3–manifold M is canonical if any essential map of the annulus or torus into M has intersection number zero with F (where the intersection number is defined as in [**19**]). In [**45**], we showed that the canonical annuli and tori in an orientable Haken 3–manifold M are the same (up to isotopy) as the frontier of the characteristic submanifold, thus yielding the classical JSJ-decomposition. There is another natural approach to this idea. As discussed at the end of chapter 1, Neumann and Swarup [**34**] defined annuli and tori embedded in a 3–manifold M to be canonical if they have intersection number zero with every *embedded* essential annulus or torus in M.

Each of these approaches has natural algebraic generalisations. Generalising our idea of canonical would involve considering splittings of a group over subgroups isomorphic to $\mathbb{Z}$ or $\mathbb{Z}\times\mathbb{Z}$ which have intersection number zero with many almost invariant sets. Generalising Neumann and Swarup's idea of canonical would involve considering splittings of a group over subgroups isomorphic to $\mathbb{Z}$ or $\mathbb{Z}\times\mathbb{Z}$ which have intersection number zero with many splittings. For our purposes, the first idea turns out to be most useful, but it seems to be important to consider a much larger class of subgroups. The following is the algebraic definition which we derive from the above discussion.

Definition 2.36. — Let G be a one-ended finitely generated group and let X be a nontrivial almost invariant subset over a subgroup H of G.

For $n \geqslant 1$, we will say that X is *n–canonical* if X has zero intersection number with any nontrivial almost invariant subset of any $K\backslash G$, for which K is VPC of length at most n.

For $n \geqslant 1$, we will say that X is *n–canonical with respect to abelian groups* if X has zero intersection number with any nontrivial almost invariant subset of any $K\backslash G$, for which K is virtually free abelian of rank at most n.

If X is associated to a splitting σ of G and X is n–canonical (with respect to abelian groups), we will say that σ is *n–canonical (with respect to abelian groups)*.

If H is virtually infinite cyclic, and X is 1–canonical, we will often say simply that X is *canonical*.

Remark 2.37. — If H is not finitely generated, we will only use these ideas when X is associated to a splitting over H.

If $n \leqslant 2$, then X is n–canonical if and only if it is n–canonical with respect to abelian groups.

The definitions above make perfectly good sense when G has more than one end, and when $n = 0$. However we show in Lemma 6.9 that if a group G has infinitely many ends, then nontrivial almost invariant subsets of G are never 0–canonical. Of course, an almost invariant subset of G which is not 0–canonical is certainly not n–canonical for any n.

Many other related ideas can be defined by changing the class of groups in which K lies. For example, one could insist that K be free abelian, or that K be VPC of length equal to some fixed number k. In particular, in [**45**], we gave a similar definition, but restricted H to be infinite cyclic, K to be free abelian, X to be associated to a splitting, and n to equal 1 or 2. In this situation, almost invariant subsets of $K\backslash G$ correspond to (possibly singular) annuli or tori in a 3–manifold. One could also restrict attention to those almost invariant subsets of $K\backslash G$ which are associated to splittings. This is analogous, in the 3–manifold situation, to considering the Enclosing Property only for embedded annuli or tori, which is effectively what Neumann and Swarup were doing in [**34**].

In [**45**], we also considered the connection between the canonical annuli and tori in M and the 2–canonical splittings of $G = \pi_1(M)$ over subgroups isomorphic to $\mathbb{Z}$ or $\mathbb{Z} \times \mathbb{Z}$. We showed that every such 2–canonical splitting of G arises from a canonical annulus or torus in M, and that every canonical annulus in M determines a 2–canonical splitting of G. Further every canonical torus in M determines a 1–canonical splitting of G. However, we also showed that often M will have canonical tori which determine splittings of G which are not 2–canonical. See Example 11.7.

In order to start developing our algebraic theory of regular neighbourhoods in the next chapter, we will need to consider collections of almost invariant subsets of a given group. The following generalisation of Definition 2.29 will be useful.

Definition 2.38. — A collection E of subsets of G which are closed under complementation is called *nested* if any pair of sets in E is nested, i.e. if U and V are sets in E, then one of the four sets $U^{(*)} \cap V^{(*)}$ is empty.

If each element U of E is a H_U–almost invariant subset of G for some subgroup H_U of G, we will say that E is *almost nested* if for any pair U and V of sets in the collection, one of the four sets $U^{(*)} \cap V^{(*)}$ is small.

If one is given a H–almost invariant subset X of G, it is natural to ask whether there is a K–almost invariant set Y which is equivalent to X and is associated to a splitting of G over K. This is analogous to asking whether a given codimension–1 immersion in a manifold can be homotoped to cover an embedding. We state below Theorem B.1.12, which we will need later. To prove this result, we used the almost nested assumption to construct a tree with G–action. There are more general results of this type in section 2 of Appendix B, which will be extended and used later on.

Theorem 2.39

(1) *Let H be a finitely generated subgroup of a finitely generated group G. Let X be a nontrivial H–almost invariant subset of G such that $E = \{gX, gX^* : g \in G\}$ is almost nested and if two of the four sets $X^{(*)} \cap gX^{(*)}$ are small, then at least one of them is empty. Then G splits over the stabiliser H' of X and H' contains H as a subgroup of finite index. Further, one of the H'–almost invariant sets Y determined by the splitting is equivalent to X.*

(2) *Let $H_1, \ldots, H_k$ be finitely generated subgroups of a finitely generated group G. Let X_i be a nontrivial H_i–almost invariant subset of G such that $E = \{gX_i, gX_i^* : 1 \leqslant i \leqslant k, g \in G\}$ is almost nested. Suppose further that, for any pair of elements U and V of E, if two of the four sets $U^{(*)} \cap V^{(*)}$ are small, then at least one of them is empty. Then G can be expressed as the fundamental group of a graph of groups whose i–the edge corresponds to a conjugate of a splitting of G over the stabiliser H_i' of X_i, and H_i' contains H_i as a subgroup of finite index.*

We end this chapter by discussing some closely related ideas which appear in the literature. Recall that, in this paper, G always denotes a finitely generated group.

In [**30**], Kropholler and Roller gave a different definition of a H–almost invariant subset X of a group G, which they used in several later papers. They omitted the condition that X be invariant under the left action of H. Thus they defined X to be H–almost invariant if X and Xg are H–almost equal for all g in G, and to be nontrivial if X and X^* are not H–finite. Thus X is H–almost invariant in their sense if and only if its coboundary δX is H–finite. They also introduced a way to count such subsets of G, which they called the number of relative ends $\widetilde{e}(G, H)$ of the pair (G, H). In particular, $\widetilde{e}(G, H) \geqslant 2$ if and only if G possesses a nontrivial H–almost invariant subset in their sense. As a subset of G which is H–almost invariant in our sense is automatically H–almost invariant in their sense, we will always have the inequality $\widetilde{e}(G,H) \geqslant e(G,H)$. But this inequality is often strict. We will now discuss the connection between these two invariants in more detail.

Suppose that H is a finitely generated subgroup of G and that X is a H–almost invariant subset of G, in the sense of [**30**]. Since δX is H–finite, it has finite image in $H\backslash\Gamma$. Thus we can find a finite connected subcomplex of $H\backslash\Gamma$ which contains the image of δX and carries H. Let L denote the pre-image of this subcomplex in Γ, so that L is a connected, H–finite subcomplex of Γ which contains δX. Further $HL = L$. Now let Y denote the complement $\Gamma - L$, so that $HY = Y$ and each component of Y is contained in X or in X^*. As the quotient $H\backslash L$ is a finite subcomplex of $H\backslash\Gamma$, the complement of $H\backslash L$ in $H\backslash\Gamma$ has finitely many components. Thus the components of Y fall into finitely many orbits under the action of H. Let X' denote the union of all those components of Y which are contained in X and are not H–finite. Then X' is H–almost equal to X. It follows that G has a nontrivial H–almost invariant subset, in the sense of [**30**], if and only if there is a connected H–finite subcomplex L of Γ such that $HL = L$ and $\Gamma - L$ has at least two H–infinite components. As mentioned above, this condition is equivalent to having $\widetilde{e}(G,H) \geqslant 2$, but it does not imply that $e(G,H) \geqslant 2$. An easy but interesting example occurs when G is the fundamental group of a closed non-orientable surface F, not P^2, and H is the subgroup of G carried by a simple but one-sided circle C on F. For then $e(G,H) = 1$, as the cover of F determined by H is an open Moebius band. But the pre-image of C in the universal cover $\widetilde{F}$ of F is a line L such that $HL = L$ and L divides $\widetilde{F}$ into two H–infinite components, so that $\widetilde{e}(G,H) = 2$. However this example is misleadingly simple. First note that if X is K–almost invariant in the sense of [**30**], then X is automatically H–almost invariant in this sense for any subgroup H of G which contains K. Surprisingly, if X is nontrivial over K, it must also be nontrivial over H, so long as H has infinite index in G. This follows immediately from Lemma 2.13.

This enables us to give the following connection between ends and relative ends.

Lemma 2.40. — *Let H be a finitely generated subgroup of infinite index in a group G. Then $\widetilde{e}(G,H) \geqslant 2$ if and only if there is a subgroup K of H such that $e(G,K) \geqslant 2$.*

Remark 2.41. — K may have infinite index in H.

Proof. — Suppose that $\widetilde{e}(G,H) \geqslant 2$, so that G has a nontrivial H–almost invariant subset X in the sense of [**30**]. As H is finitely generated, the previous discussion shows that there is a connected H–finite subcomplex L of Γ such that $HL = L$, and $Y = \Gamma - L$ has at least two H–infinite components. Let Z denote one of these H–infinite components of Y, and let K denote the subgroup of H which stabilises Z. Then K must equal the subgroup of H which stabilises δZ or have index 2 in this subgroup. As δZ is contained in L, it must be H–finite and hence must also be K–finite. Further, as the quotients $K\backslash Z$ and $K\backslash Z^*$ are both infinite and $K\backslash \delta Z$ is finite, it follows that $e(G,K) \geqslant 2$.

Conversely, suppose that H has a subgroup K such that $e(G,K) \geqslant 2$. Thus there is a nontrivial K–almost invariant (in our sense) subset Z of G. Trivially, Z is also K–almost invariant, and hence H–almost invariant, in the sense of [**30**]. As H has infinite index in G, Lemma 2.13 shows that Z must be nontrivial over H, so that $\widetilde{e}(G,H) \geqslant 2$, as required. □

In [**8**], Bowditch considered a notion which he called 'coends'. In [**21**], Geoghegan independently considered the same notion which he called 'filtered coends'. In [**8**], Bowditch considers a pair (G,H) of finitely generated groups where G is one-ended and H is two-ended. Let G act properly discontinuously and cocompactly on a locally finite one-ended graph Γ, for example the Cayley graph of G. Let $S(H)$ denote the set of H–invariant connected subgraphs of Γ with finite quotient. If $A \in S(H)$, let $C(A)$ denote the collection of all complementary components of A and let $C_\infty(A)$ denote those components which are not contained in any element of $S(H)$. If A and B are elements of $S(H)$ such that $A \subset B$, then there is a surjective map from $C_\infty(B)$ to $C_\infty(A)$. Let $E(H)$ denote the inverse limit in the category of topological spaces, where A and B range over $S(H)$. Then $E(H)$ is a compact and totally disconnected space and an element of $E(H)$ is called a coend of H. It is clear from this discussion that the concept of coends extends to any pair of finitely generated groups (G,H). It is also clear that there is a close connection between coends of H in G and H–almost invariant subsets of G in the sense of [**30**]. Geoghegan showed that the number of coends of a subgroup H of a group G is the same as the number $\widetilde{e}(G,H)$ of relative ends of H in G as defined by Kropholler and Roller in [**30**]. (See also the introduction of [**8**].)

Bowditch calls an element of $S(H)$ an axis of H if it satisfies some further technical conditions. An axis A is called proper if $C_\infty(A)$ has at least two elements. If a proper axis corresponding to a subgroup K crosses A, then Bowditch shows that H and K both have two coends (see section 10 of [**5**]). If the number of coends of H is 2 and H does not interchange the coends, there is an essentially unique (up to equivalence and complementation) H–almost invariant subset. Thus in this case, crossing of axes is the same as strong crossing of the corresponding almost invariant sets. The arguments used by Bowditch in [**8**] are geared to the above case where G is one-ended and H is two-ended, but his ideas work just as well in the case where H is VPC and G does not split over any subgroup of length less than that of H.

CHAPTER 3

ALGEBRAIC REGULAR NEIGHBOURHOODS: CONSTRUCTION

In this and the following three chapters, we will develop our algebraic analogue of the topological idea of a regular neighbourhood. It would be possible to start by giving an abstract definition of an algebraic regular neighbourhood, and then to prove existence and uniqueness results. But the precise definition is somewhat technical, so we will start with our construction of an algebraic regular neighbourhood. We will consider a finitely generated group G with finitely generated subgroups $H_1, \ldots, H_n$. For $i = 1, \ldots, n$, let X_i be a nontrivial H_i–almost invariant subset of G. In this chapter, we will construct a finite graph of groups structure $\Gamma(X_1, \ldots, X_n : G)$ for G. In chapter 6, we will define abstractly what constitutes an algebraic regular neighbourhood of the X_i's in G, and then will prove that $\Gamma(X_1, \ldots, X_n : G)$ satisfies the requirements of our definition. We will also prove that any algebraic regular neighbourhood of the X_i's in G is naturally isomorphic to $\Gamma(X_1, \ldots, X_n : G)$. These results are the algebraic analogue of the existence and uniqueness results for regular neighbourhoods in topology.

While the restriction to finite families of almost invariant sets is very natural, regular neighbourhoods of infinite families will play an important role in this paper. At the end of this chapter, we will briefly discuss how to modify our construction when the number of X_i's is infinite. At the end of chapter 5, we also discuss what happens to our construction if some of the H_i's are not finitely generated. It is a surprising fact that, even in this situation, our theory of algebraic regular neighbourhoods goes through in certain cases.

In order to introduce our ideas, consider a connected manifold M and let $\mathcal{T}$ be a compact (possibly disconnected) codimension–1 two-sided manifold properly embedded in M. (The reader will not miss anything by thinking of M as a surface, and $\mathcal{T}$ as a collection of circles and arcs.) If each component of $\mathcal{T}$ is π_1–injective in M, then $\mathcal{T}$ determines a graph of groups structure Γ on $G = \pi_1(M)$ whose underlying graph is dual to $\mathcal{T}$. The edge groups of Γ are the fundamental groups of the components of $\mathcal{T}$

and the vertex groups are the fundamental groups of the components of the complement of $\mathcal{T}$. If the components of $\mathcal{T}$ are not all π_1–injective, then $\mathcal{T}$ still determines a graph of groups structure for G, with the same underlying graph, but the edge and vertex groups are obtained from the above by replacing each group by its image in G.

Now suppose that we consider a compact (possibly disconnected) codimension–0 submanifold N of M. We will associate to this the graph of groups decomposition Γ of G determined as above by the frontier of N in M. (If M is closed, this frontier is the same as the boundary of N.) The vertices of Γ correspond to components of N and of $M - N$ and each edge of Γ joins a vertex of one type to a vertex of the other type. Thus Γ is naturally a bipartite graph. Throughout this paper, we will denote the collection of vertices of Γ which correspond to components of N by $V_0(\Gamma)$, or simply V_0 if the context is clear. The remaining vertices will be denoted by $V_1(\Gamma)$ or simply V_1. If we consider the pre-image $\widetilde{N}$ of N in the universal cover $\widetilde{M}$ of M, the dual graph to $\partial\widetilde{N}$ is a tree T on which G acts with quotient Γ, and the vertex and edge groups of Γ are simply the vertex and edge stabilisers for the action of G on T. Again T is naturally bipartite with some vertices corresponding to components of $\widetilde{N}$ and some vertices corresponding to components of $\widetilde{M} - \widetilde{N}$.

In the previous paragraph, we discussed how any codimension–0 submanifold N of M corresponds to a bipartite graph of groups structure for $\pi_1(M)$. In what follows, we will be interested in the situation where N is a regular neighbourhood of the union of a finite collection of codimension–1 manifolds C_λ properly immersed in a manifold M, and in general position. For simplicity, suppose that each C_λ lifts to an embedding in M_λ, the cover of M whose fundamental group equals that of C_λ, and let S_λ denote the pre-image in $\widetilde{M}$ of C_λ in M_λ. Let Σ denote the collection of all the translates of all the S_λ's, and let $|\Sigma|$ denote the union of all the elements of Σ. Thus $|\Sigma|$ is the complete pre-image in $\widetilde{M}$ of the union of the images of all the C_λ's and $\widetilde{N}$ is a regular neighbourhood of $|\Sigma|$. The fact that N is a regular neighbourhood of the C_λ's implies that the inclusion of the union of the C_λ's into N induces a bijection between components, and an isomorphism between the fundamental groups of corresponding components. It follows that the inclusion of $|\Sigma|$ into $\widetilde{N}$ also induces a bijection between components. Thus the V_0–vertices of T correspond to the components of $|\Sigma|$, and the V_1–vertices of T correspond to the components of $\widetilde{M} - |\Sigma|$. If two elements S and S' of Σ belong to the same component of $|\Sigma|$, there must be a finite chain $S = S_0, S_1, \ldots, S_n = S'$ of elements of Σ such that S_i intersects S_{i+1}, for each i. Thus the elements of Σ which form a component of $|\Sigma|$ are an equivalence class of the equivalence relation on Σ generated by saying that two elements of Σ are related if they intersect. This is what we want to encode in the algebraic setting, except that as we will be dealing with almost invariant sets, we will want to ignore all "inessential" intersections.

Now we return to our finitely generated group G with finitely generated subgroups $H_1, \ldots, H_n$ and nontrivial H_i–almost invariant subsets X_i of G, and will describe our construction of the finite graph of groups structure $\Gamma(X_1, \ldots, X_n : G)$ of G. Let $E = \{gX_i, gX_i^* : g \in G, 1 \leqslant i \leqslant n\}$. Previously we said that the X_i's are the algebraic analogue of the immersed C_λ's. However, it is neater to consider the pair $\{X, X^*\}$ as a single object, and regard this as the algebraic analogue of an immersion. We will denote the unordered pair $\{X, X^*\}$ by $\overline{X}$, and will say that $\overline{X}$ crosses $\overline{Y}$ if X crosses Y. Then our algebraic analogue of the set Σ is the set $\overline{E} = \{g\overline{X_i} : g \in G, 1 \leqslant i \leqslant n\}$. Our analogue of the equivalence relation on Σ which described the components of $|\Sigma|$ is the equivalence relation on $\overline{E}$ which is generated by saying that two elements A and B of $\overline{E}$ are related if they cross. We call an equivalence class a *cross-connected component* (CCC) of $\overline{E}$, and denote the equivalence class of A by $[A]$. Note that this is a purely combinatorial definition. The use of the word component is simply to emphasise the analogies with the topological situation of the preceding paragraph. We will denote the collection of all CCC's in $\overline{E}$ by P.

We want to construct in a natural way a bipartite G–tree T with P as its set of V_0–vertices. Then we will let $\Gamma(X_1, \ldots, X_n : G)$ be the quotient $G \backslash T$. Note that if one has a tree, there is a natural idea of betweenness for vertices. We will reverse this process and construct the required tree T starting from an idea of betweenness on the set P. In order to define this idea of betweenness, we will first introduce a partial order on E using ideas from the proof of Theorem 2.39. There we defined a partial order on the set E, and constructed a G–tree whose edges were the elements of E by showing that the partial order satisfied certain conditions. Unlike the situation of Theorem 2.39, we cannot expect to construct a tree with the elements of E as its edges. For this would imply that each X_i determined a splitting of G. However, the partial order will still play a crucial role in our situation.

If U and V are two elements of E such that $U \subset V$, then our partial order will have $U \leqslant V$. But we also want to define $U \leqslant V$ when U is "nearly" contained in V. Precisely, we want $U \leqslant V$ if $U \cap V^*$ is small. However, an obvious difficulty arises when two of $U^{(*)} \cap V^{(*)}$ are small, as we have no way of deciding between two possible inequalities. It turns out that we can avoid this difficulty if we know that whenever two of $U^{(*)} \cap V^{(*)}$ are small, then one of them is empty. Thus we consider the following condition on E:

Condition (*): If U and V are in E, and two of $U^{(*)} \cap V^{(*)}$ are small, then one of $U^{(*)} \cap V^{(*)}$ is empty.

If E satisfies Condition (*), we will say that the family $X_1, \ldots, X_n$ is in *good position.*

Assuming that this condition holds, we can define a relation $\leqslant$ on E by saying that $U \leqslant V$ if and only if $U \cap V^*$ is empty or is the only small set among the four sets $U^{(*)} \cap V^{(*)}$. Despite the seemingly artificial nature of this definition, one can show that $\leqslant$ is a partial order on E. This is not entirely trivial, but the proof is in Lemma 2.4 of [**41**] and repeated more group theoretically in Lemma B.1.14. Condition (*) plays a key role in the proof. If $U \leqslant V$ and $V \leqslant U$, it is easy to see that we must have $U = V$, using the fact that E satisfies Condition (*). Most of the proof of Lemma B.1.14 is devoted to showing that $\leqslant$ is transitive. We note here that the argument that $\leqslant$ is transitive does not require that the H_i's be finitely generated.

In general, the family $X_1, \ldots, X_n$ need not be in good position, but it turns out that this does not cause any problems. We will discuss this after Theorem 3.8.

Before we go any further, we need to discuss the idea of a pretree. As already mentioned, the vertices of a tree possess a natural idea of betweenness. The idea of a pretree formalises this. A *pretree* consists of a set P together with a ternary relation on P denoted xyz which one should think of as meaning that y is strictly between x and z. The relation should satisfy the following four axioms:

(T0) If xyz, then $x \neq z$.
(T1) xyz implies zyx.
(T2) xyz implies not xzy.
(T3) If xyz and $w \neq y$, then xyw or wyz.

A pretree is said to be *discrete*, if, for any pair x and z of elements of P, the set $\{y \in P : xyz\}$ is finite. Clearly, the vertex set of any simplicial tree forms a discrete pretree with the induced idea of betweenness. It is a standard result that a discrete pretree P can be embedded in a natural way into the vertex set of a simplicial tree T so that the notion of betweenness is preserved. We briefly describe the construction of T following Bowditch's papers [**5**] and [**8**]. For the proofs, see section 2 of [**5**] and section 2 of [**8**]. These are discussed in more detail in [**7**] and [**1**]. For any pretree P, we say that two distinct elements of P are *adjacent* if there are no elements of P between them. We define a *star* in P to be a maximal subset of P which consists of mutually adjacent elements. (This means that any pair of elements of a star are adjacent. It also means that any star has at least 2 elements.) We now enlarge the set P by adding in all the stars of P to obtain a new set V. One can define a pretree structure on V which induces the original pretree structure on P. A star is adjacent in V to each element of P that it contains. Next we give V the discrete topology and add edges to V to obtain a graph T with V as its vertex set. For each pair of adjacent elements of V, we simply add an edge which joins them. If P is discrete, then it can be shown that T is a tree with vertex set V. It follows easily from this construction that if a group G acts on the original pretree P, this action extends naturally to an action of G on the simplicial tree T. Moreover, G will act without inversions on T. This will then give a graph of groups decomposition for G, though this decomposition would be trivial if G fixed a vertex of T. The tree T is clearly bipartite with vertex set $V(T)$ expressed as the disjoint union of $V_0(T)$ and $V_1(T)$, where $V_0(T)$ equals P and

$V_1(T)$ equals the collection of stars in P. Note that each V_1–vertex of T has valence at least 2, because a star in P always has at least 2 elements.

Note also that if we start with a tree T', let P denote its vertex set with the induced idea of betweenness and then construct the tree T as above, then T is not the same as T'. In fact, T is obtained from T' by subdividing every edge into two edges.

Now we return to our discussion of the set E which we still assume satisfies Condition (*). We have the partial order $\leqslant$ on E and we write $U < V$ if $U \leqslant V$ but U is not equal to V. If $U < Z < V$, we will say that Z is *between* U and V. We summarise below some elementary facts about E.

Lemma 3.1. — *If E satisfies Condition (*), then E together with $\leqslant$ satisfies the following conditions.*

(1) *If U, $V \in E$ and $U \leqslant V$, then $V^* \leqslant U^*$,*

(2) *If U, $V \in E$, there are only finitely many $Z \in E$ such that Z is between U and V,*

(3) *If U, $V \in E$, one cannot have $U \leqslant V$ and $U \leqslant V^*$.*

Proof. — As U^* denotes the complement of U, the first part of this lemma is clear.

Since E consists of translates of a finite number of almost invariant sets over finitely generated subgroups, it is a standard fact that there are only finitely many elements of E between U and V (See Lemma 2.6 in [**41**] or Lemma B.1.15).

Finally, if U and V lie in E, one cannot have $U \leqslant V$ and $U \leqslant V^*$. For this would imply that $U \cap V^*$ and $U \cap V$ are each small, so that U itself would be small, contradicting the assumption that each X_i is a nontrivial H_i–almost invariant subset of G. □

Condition 2) of the above lemma will play an important role in our later discussions, so we will give this condition a name.

Definition 3.2. — Let E be a partially ordered set. We will say that E is *discrete* if for any elements U and V of E, there are only finitely many elements of E between U and V.

Next we use these properties of our partial order on E to define a notion of betweenness on the set $\overline{E}$. Recall that $\overline{X}$ denotes the unordered pair $\{X, X^*\}$ and that $\overline{E}$ denotes the set of all translates of all the $\overline{X_i}$'s.

Definition 3.3. — Let L, M, $N \in \overline{E}$. We say that M is *between* L and N if there exist $U \in L$, $V \in M$, and $W \in N$ such that $U < V < W$, and we write LMN or $\overline{U}\,\overline{V}\,\overline{W}$ with U, V, W chosen as above.

Note that it is clear that if $\overline{U}\,\overline{V}\,\overline{W}$ holds, then $\overline{W}\,\overline{V}\,\overline{U}$ also holds.

Recall that we say that $\overline{X}$ crosses $\overline{Y}$ if X crosses Y. This generates an equivalence relation on $\overline{E}$, whose equivalence classes we call *cross-connected components* (CCC).

We denote the equivalence class of $\overline{X}$ by $[\overline{X}]$, and denote the collection of all CCC's in $\overline{E}$ by P.

We extend the above idea of betweenness in $\overline{E}$ to one in P, as follows.

Definition 3.4. — Let A, B and C be distinct cross-connected components of $\overline{E}$. We say that B is *between* A and C and write ABC if there exist $\overline{U} \in A$, $\overline{V} \in B$ and $\overline{W} \in C$ such that $\overline{V}$ is between $\overline{U}$ and $\overline{W}$, i.e. $\overline{U}\,\overline{V}\,\overline{W}$.

In order for this definition to be useful, we need to know that it is independent of the choices of $\overline{U}$ and $\overline{W}$. This is what we prove in Corollary 3.7 below. We need two small results first.

Lemma 3.5. — *If U, V and Z are elements of E such that $U \leqslant V$, and Z crosses U but Z does not cross V, then either $Z \leqslant V$ or $Z^* \leqslant V$.*

Proof. — As Z crosses U, none of $Z^{(*)} \cap U^{(*)}$ is small. Since $U \leqslant V$, it follows that $Z \cap V$ and $Z^* \cap V$ are not small. As Z does not cross V, either $Z \cap V^*$ or $Z^* \cap V^*$ is small. Hence either $Z \leqslant V$ or $Z^* \leqslant V$ as claimed. □

Lemma 3.6. — *If U, V and Z are elements of E such that $U < V$, and Z crosses U but Z does not cross V, then either $Z < V$ or $Z^* < V$.*

Proof. — This follows from the preceding lemma, since if either Z or Z^* were equal to V, we would have one of the inequalities $U < Z$ or $U < Z^*$, which would contradict the assumption that Z crosses U. □

Corollary 3.7. — *Let A, B and C be distinct cross-connected components of $\overline{E}$, and suppose that $\overline{U}$ and $\overline{U'}$ lie in A, $\overline{V}$ lies in B, and $\overline{W}$ and $\overline{W'}$ lie in C. If $\overline{U}\,\overline{V}\,\overline{W}$, then $\overline{U'}\,\overline{V}\,\overline{W'}$.*

Proof. — It is easy to reduce this to the case when $\overline{U'}$ crosses $\overline{U}$ and $\overline{W} = \overline{W'}$. We can also assume that $U < V < W$. Now we write $U' = Z$ and apply the preceding lemma. This tells us that either $Z < V$ or $Z^* < V$. By the definition of betweenness, this implies that $\overline{V}$ is between $\overline{U'}$ and $\overline{W}$, as required. □

Now we are ready to show that if P denotes the collection of all CCC's of $\overline{E}$ equipped with the relation of betweenness defined above, then P is a discrete pretree.

Theorem 3.8. — *Let G denote a finitely generated group, and let $H_1, \ldots, H_n$ be finitely generated subgroups of G. For each $1 \leqslant i \leqslant n$, let X_i be a nontrivial H_i–almost invariant subset of G, and suppose that the X_i's are in good position, so that the set E of all translates of all the X_i's satisfies Condition (*). Form the set $\overline{E}$ as above, and consider the collection P of all cross-connected components of $\overline{E}$ equipped with the relation of betweenness introduced above. Then the following statements hold:*

(1) *P is a pretree, and G acts on P in a natural way.*

(2) *The pretree P is discrete, and the quotient $G\backslash P$ is finite. Further, the stabilisers of elements of P under this G–action are finitely generated.*

(3) *As P is discrete, it can be embedded in a natural way into the vertex set of a G–tree T so that the quotient $G\backslash T$ is a bipartite graph of groups $\Theta(X_1, \ldots, X_n : G)$. This graph is finite, and the V_0–vertex groups are finitely generated.*

Remark 3.9. — In most cases, $\Theta(X_1, \ldots, X_n : G)$ will be our algebraic regular neighbourhood of the X_i's in G. The V_0–vertex groups of $\Theta(X_1, \ldots, X_n : G)$ are called the *enclosing groups* of the X_i's. We will formally define a regular neighbourhood in chapter 6, and will define enclosing in chapter 4. Recall that $\Theta(X_1, \ldots, X_n : G)$ is the algebraic analogue of a regular neighbourhood of a finite family of immersed codimension–1 submanifolds of a manifold with fundamental group G, and the enclosing groups correspond to the fundamental groups of the components of the regular neighbourhood.

Note that even though the enclosing groups of $\Theta(X_1, \ldots, X_n : G)$ are finitely generated, the edge groups and the V_1–vertex groups need not be finitely generated, even when G is finitely presented. We give examples of this phenomenon in Example 6.11.

The result that $\Theta(X_1, \ldots, X_n : G)$ is finite can be strengthened greatly. We will see in Proposition 5.2 that T is a minimal G–tree.

Proof. — 1) The action of G on itself by left multiplication induces an action of G on E and hence on $\overline{E}$. As this action preserves crossing, it is immediate that it induces an action on P.

Now we verify that P satisfies the four axioms (T0)-(T3) for a pretree. For the convenience of the reader, we give these axioms again.

(T0) If xyz, then $x \neq z$.
(T1) xyz implies zyx.
(T2) xyz implies not xzy.
(T3) If xyz and $w \neq y$, then xyw or wyz.

Axioms (T0) and (T1) are immediate from our definition of betweenness.

To prove (T2), suppose that A, B and C are elements of P such that we have both ABC and ACB. As ABC holds, Corollary 3.7 tells us that there is $\overline{V} \in B$ such that we have $\overline{U}\,\overline{V}\,\overline{W}$ for any $\overline{U} \in A$ and $\overline{W} \in C$. As ACB also holds, there is $\overline{W} \in C$ such that $\overline{U}\,\overline{W}\,\overline{V}$ holds for any $\overline{U} \in A$. We choose some $\overline{U} \in A$. For these particular choices of U, V and W, we have both $\overline{U}\,\overline{V}\,\overline{W}$ and $\overline{U}\,\overline{W}\,\overline{V}$. As $\overline{U}\,\overline{V}\,\overline{W}$ holds, we can arrange that $U < V < W$, by replacing sets by their complement if needed. As $\overline{U}\,\overline{W}\,\overline{V}$ holds, there exist $X \in \overline{U}$, $Y \in \overline{V}$, $Z \in \overline{W}$ such that $X < Z < Y$. Now consider the inequalities $U < V$ and $X < Y$, and recall that X equals U or U^*, and Y equals V or V^*. It is easy to see that the only possibility is that $X = U$ and $Y = V$. For example, if we had $X = U^*$ and $Y = V$, the inequalities $U < V$ and $X < Y$ would imply that $U < V$ and $U^* < V$, which is impossible. Similarly, the inequalities $U < W$ and $X < Z$ imply that $X = U$ and $Z = W$. But now the inequality $V < W$ implies that $Y < Z$ which contradicts the inequality $Z < Y$. This completes the proof of (T2).

We next verify (T3). Suppose that we have A, B, C, $D \in P$ with ABC and $D \neq B$. We must show that ABD or DBC. Choose $\overline{U} \in A$, $\overline{V} \in B$, $\overline{W} \in C$ so that $U < V < W$, and choose $\overline{X} \in D$. The result is trivial if D equals A or C, so we will assume that D is not one of A, B or C. Thus X does not cross any of U, V or W. Since X does not cross U, one of $X^{(*)} \cap U^{(*)}$ is small. Thus, we can, by interchanging X and X^* if necessary, arrange that either $U < X$ or $X < U$. If $X < U$, then $X < V < W$ so that DBC holds, and we are done. If $U < X$, then we compare X and V. Since they do not cross, we should have one of the four inequalities $X < V$, $V < X$, $X < V^*$ or $V^* < X$. If $X < V$, then as above we have DBC, and if $V < X$, we have $U < V < X$ so that ABD holds. The inequality $X < V^*$ is impossible, as we already have $U < X$ and $U < V$. The inequality $V^* < X$ implies that $X^* < V$, which again implies DBC. This completes the proof that P is a pretree.

2) As the given family of X_i's is finite, part 2) of Lemma 3.1 tells us that P is discrete.

There is a natural surjective map from the given family of X_i's to the quotient $G\backslash\overline{E}$. As the quotient $G\backslash P$ is a quotient of $G\backslash\overline{E}$, it follows that $G\backslash P$ is finite, as required.

Note that the surjection from the given family of X_i's to $G\backslash\overline{E}$ may not be a bijection because it is possible that there are distinct i and j such that X_j has a translate which is equal to X_i or X_i^*. This is an annoyance which we can easily remove as follows. It is clear that the construction of $\Theta(X_1, \ldots, X_n : G)$ depends only on the G–orbits of the $\overline{X_i}$'s in $\overline{E}$ rather than the X_i's themselves. Thus, by choosing one X_i for each such G–orbit, we can always replace the given family of X_i's by a subfamily which yields the same G–orbits, and such that distinct $\overline{X_i}$'s lie in distinct G–orbits in $\overline{E}$. Such a family will yield the same graph of groups decomposition of G, so we will do this whenever it is convenient. In particular, we will assume this condition holds for the rest of this proof.

It remains to show that the stabilisers of elements of P are finitely generated. We start by noting that the stabiliser K_i of $\overline{X_i}$ contains H_i as a subgroup of finite index and so must be finitely generated. Now let v denote an element of P, and consider those $\overline{X_i}$'s which have a translate in v. By renumbering, we can assume that $\overline{X_i}$ has a translate Y_i in v for $1 \leqslant i \leqslant k$. We let S_i denote the stabiliser of Y_i. As S_i is conjugate to K_i, it is also finitely generated.

There is a natural distance function on the elements of v described as follows. If U and V are elements of v, then $d(U,U) = 0$, and $d(U,V) \leqslant r$ if and only if there is a sequence $U = U_0, U_1, \ldots, U_r = V$ of elements of v such that U_i and U_{i+1} cross for $0 \leqslant i \leqslant r-1$. We call such a sequence a *chain*. Thus $d(U,V)$ is the length of the shortest chain from U to V. This function extends to a distance function on finite subsets of v in the usual way. Now let A denote the set $\{Y_1, \ldots, Y_k\}$, and let $G(A)$ denote the subgroup of G generated by $X = \{g \in G : d(gA, A) \leqslant 1\}$.

We claim that $G(A)$ is finitely generated, and that $\mathrm{Stab}(v)$, the stabiliser of v, equals $G(A)$. This will show that $\mathrm{Stab}(v)$ is finitely generated, as required.

First we show that $G(A)$ is finitely generated. Recall that if U is an element of A, it has finitely generated stabiliser. We will denote this stabiliser by $S(U)$. Now let g denote an element of X. If $d(gA, A) = 0$, there are elements U and V of A such that $gV = U$. As we have arranged that distinct elements of A belong to distinct G–orbits, it follows that $U = V$. Thus the collection of all such elements g is simply $S(U)$. If $d(gA, A) = 1$, then g stabilises no element of A, but there are elements U and V of A, possibly equal, such that gV crosses U. Lemma 2.30 shows that the collection of all such elements g consists of the union of a finite number of double cosets $S(U)gS(V)$. It follows that $G(A)$ is generated by the S_i's, $1 \leqslant i \leqslant k$, and by a finite collection of double cosets of these groups, and so is finitely generated.

Next we show that $\mathrm{Stab}(v)$ equals $G(A)$. Clearly $G(A)$ is contained in $\mathrm{Stab}(v)$. We will prove that any element g of $\mathrm{Stab}(v)$ lies in $G(A)$, by induction on the integer $d(gA, A)$. If $d(gA, A) \leqslant 1$, this is immediate from the definition of $G(A)$. Now suppose that $d(gA, A) = D > 1$. Thus there are elements U and V of A such that $d(gA, A) = d(U, gV) = D$. Pick a chain from U to gV of length D, and let W denote the element of A with a translate hW which immediately precedes gV in this chain. Thus $d(hA, A) \leqslant d(U, hW) = D-1$, so that h lies in $G(A)$ by our induction hypothesis. As hW crosses gV, we see that W crosses $h^{-1}gV$, so that $d(h^{-1}gA, A) \leqslant 1$. Hence $h^{-1}g$ also lies in $G(A)$, so that g itself lies in $G(A)$, as required. This completes the proof of part 2).

3) Recall that a discrete pretree P can be embedded in a natural way into the vertex set of a tree T, and that an action of G on P which preserves betweenness will automatically extend to an action without inversions on T. Also T is a bipartite tree with vertex set $V(T) = V_0(T) \cup V_1(T)$, where $V_0(T)$ equals P, and $V_1(T)$ equals the collection of all stars in P. It follows that the quotient $G \backslash T$ is naturally a bipartite graph of groups Θ with V_0–vertex groups conjugate to the stabilisers of elements of P and V_1–vertex groups conjugate to the stabilisers of stars in P. Further, part 2) implies that Θ has only finitely many V_0–vertices, and that each such vertex carries a finitely generated group. Now we can show that Θ must be a finite graph. For as G is finitely generated, there is a finite subgraph Θ_1 of Θ with fundamental group G. Thus $\Theta - \Theta_1$ has finitely many components each of which must be a tree in which every edge determines a trivial splitting of G. Hence if Θ were infinite, one of these trees would be infinite. Part 2) tells us that Θ has only finitely many V_0–vertices. Thus any infinite subtree of Θ must have a V_1–vertex v which is terminal in Θ. As the edge incident to v determines a trivial splitting of G, it follows that T has a terminal

V_1–vertex which projects to v, which contradicts the fact that every V_1–vertex of T has valence at least 2. This contradiction completes the proof of the theorem. □

As we remarked immediately after the statement of Theorem 3.8, in most cases, the graph of groups $\Theta(X_1, \ldots, X_n : G)$ will be our algebraic regular neighbourhood of the X_i's in G. However, there are a few situations where our construction of Θ does not yield what we want. There is an important point here, and we will discuss some relevant examples. The most basic example occurs when one has a single almost invariant set X which is associated to a splitting. Surprisingly, a subtle problem can occur even in this case. An important property of our algebraic regular neighbourhood of a family of almost invariant subsets of a group is that it should be unchanged if we replace the sets in the family by equivalent sets. However our construction does not always have this property. Recall from Definition 2.12 that if X is an H–almost invariant subset of a group G associated to a splitting of G over H, then X is invertible if there is an element g in G such that $gX = X^*$. We will see that whether or not X is invertible has an important effect on our construction of $\Theta(X)$.

Example 3.10. — Let G be any group which splits over a subgroup H. This implies that G acts on a tree T with quotient a graph Γ with a single edge which yields the given splitting of G. We will let X denote a H–almost invariant subset of G associated to the given splitting and apply the preceding construction to the set E of all translates of X and X^*, to obtain a new graph of groups structure $\Theta = \Theta(X)$ for G.

If X is not invertible, then we will see that Θ can be obtained from Γ by subdividing its edge into two edges. The original vertices are V_1–vertices and the new vertex is the single V_0–vertex. This is the graph of groups structure on G which we want to be the algebraic regular neighbourhood of X.

To see that Θ has the claimed structure, let s be an edge of T with stabiliser H, oriented so that Z_s is equivalent to X, and consider the map η from the oriented edges of T to the set E, given by sending gs to gX and $g\overline{s}$ to gX^*. Note that this map is well defined, because if $gs = s$, then g lies in H, so that $gX = X$. Conversely, if $gX = X$, then some power g^k of g must lie in H. Thus $g^k Z_s = Z_s$ which implies that $gs = s$, so that g lies in H. As we are assuming that there is no element g in G such that $gX = X^*$, it follows that η is a bijection. Now the partial order $\leqslant$ on E corresponds to the natural ordering of the oriented edges of T. As each CCC of $\overline{E}$ consists of a single translate of $\overline{X}$, the collection P of all CCC's corresponds to the edges of T and has the same pretree structure as the collection of midpoints of edges of T. It follows that the G–tree constructed from the pretree P can be obtained from T by dividing each edge into two. Thus Θ has the claimed structure.

If X is invertible, there is an element k in G such that $kX = X^*$, and Lemma 2.15 implies that $G = A *_H B$, where H has index 2 in B and k lies in B. We will see that Θ is the same graph of groups as Γ, and it has one V_0–vertex and one V_1–vertex. The

V_0–vertex is the vertex which carries B. This means that if H has index 2 in both A and B, there are two distinct possibilities for Θ, as either of the two vertices could be the V_0–vertex, depending on the choice of X.

To see that Θ has the claimed structure, we again let s be an edge of T with stabiliser H, oriented so that Z_s is equivalent to X, and consider the map η from the oriented edges of T to the set E, given by sending gs to gX and $g\overline{s}$ to gX^*. This time, η is not a bijection as it identifies ks with $\overline{s}$. Of course, it follows that for any g in G, the map η identifies gks with $g\overline{s}$. However, as before η does not identify distinct translates of s. The proof of Lemma 2.15 shows that s has a vertex w of valence 2 and that ks meets s in w. It follows that the interiors of $ks \cup s$ and $g(ks \cup s)$ are equal or disjoint. Hence the only identifications induced by η are the identifications of gks with $g\overline{s}$ and of $gk\overline{s}$ with gs, for all g in G. Thus η induces a bijection between the set of all translates of w and the set $\overline{E}$. Further this bijection preserves the pretree structures on these two sets. It follows that the G–tree constructed from the pretree P is the same as T. Thus Θ has the claimed structure.

If one considers several compatible splittings of G, the situation is similar.

Example 3.11. — If G has a finite family of compatible splittings $\sigma_1, \ldots, \sigma_n$, then G acts on a tree T with quotient a graph Γ with n edges which correspond to the given splittings. If we let X_i denote a H_i–almost invariant subset of G associated to the splitting σ_i, chosen not to be invertible, and we apply the preceding construction to the set E of all translates of X_i and X_i^*, $1 \leqslant i \leqslant n$, we obtain a new graph of groups structure Θ for G which is obtained from Γ by subdividing each edge into two edges. The original vertices are V_1–vertices and the new vertices are the V_0–vertices. This is the graph of groups structure on G which we want to be the algebraic regular neighbourhood of the X_i's.

If some X_i's are invertible, then, as in the preceding example, the corresponding edges of Γ will not be subdivided when forming Θ.

Next we show that this fits neatly with simple topological examples.

Example 3.12. — Let M be a surface and let C denote a finite family of essential, two-sided, simple arcs or closed curves on M.

(1) If the arcs and curves in C are all disjoint, then C defines a graph of groups structure Γ for $G = \pi_1(M)$ such that the underlying graph of Γ is dual to C. As in the previous example, so long as we avoid invertible almost invariant sets, our algebraic regular neighbourhood construction yields a graph of groups structure Θ for G which is obtained from Γ by subdividing each edge into two edges, and this corresponds exactly to a topological regular neighbourhood of C. If a component S of C is a circle which bounds a Moebius band W, it is possible to choose the associated almost invariant set to be invertible. In this case Θ corresponds to a topological regular neighbourhood of the family C' obtained from C by replacing S by the core of W.

(2) If C denotes two essential simple arcs or closed curves on M which have minimal possible intersection, then again it is true that, so long as we avoid invertible almost invariant sets, our algebraic regular neighbourhood construction yields a graph of groups structure Θ for G which corresponds exactly to a topological regular neighbourhood of C.

(3) However, if C denotes three essential simple arcs or closed curves on M such that each pair has minimal possible intersection, the algebraic regular neighbourhood may be a little different from the topological one. For example, cutting M along C may yield a disc component. In this case, the dual graph for the regular neighbourhood would have a terminal vertex carrying the trivial group, but our algebraic regular neighbourhood construction can never yield such a vertex.

As we have seen, there are occasions on which the graph of groups $\Theta(X_1, \ldots, X_n : G)$ constructed in Theorem 3.8 is not what we want for our algebraic regular neighbourhood. We will see shortly that this problem of invertible almost invariant sets is essentially the only problem. Before discussing this, recall that in the preceding construction and the proof of Theorem 3.8, we used the assumption that the family $X_1, \ldots, X_n$ was in good position. This condition need not always be satisfied, so we need to discuss how to modify our construction to handle the general situation.

Suppose that we consider any finite family of nontrivial H_i–almost invariant subsets X_i of G. Recall that the basic idea in our construction of $\Theta(X_1, \ldots, X_n : G)$ was that of a cross-connected component (CCC) of $\overline{E}$. We can consider the equivalence relation generated by crossing of elements of $\overline{E}$ whether or not the family $X_1, \ldots, X_n$ is in good position. Thus we can always define the family P of all CCC's of $\overline{E}$ and there will always be a natural action of G on P. The importance of good position was that it enabled us to define the inequality $\leqslant$ on E and hence to define the relation of betweenness on P. Suppose that we have distinct elements U and V of E such that two of $U^{(*)} \cap V^{(*)}$ are small, but neither is empty. This means that, when we attempt to define the inequality $\leqslant$ on E, the elements U and V are not comparable. However, note that U and V must be equivalent (see Definition 2.8) up to complementation. Thus if there is an element W of E which crosses U, then W also crosses V, so that $\overline{U}$ and $\overline{V}$ will lie in the same CCC. We will say that the family $X_1, \ldots, X_n$ is *in good enough position* if whenever we find incomparable elements U and V of E which do not cross, there is some element W of E which crosses them. It is easy to see that all the preceding discussion in this chapter applies essentially unchanged if the family $X_1, \ldots, X_n$ is in good enough position. The point is that any pair of incomparable elements already lie in the same CCC, and so we never need to be able to compare them.

Next we consider the case when the family $X_1, \ldots, X_n$ is not in good enough position. We will say that an element of E which crosses no element of E is *isolated in E*. Note that this condition depends on the set E, but we will often omit the phrase "in E" when the context is clear. As any translate of an isolated element is also isolated, such elements can occur only if the original family $X_1, \ldots, X_n$ contains elements which are isolated in E. By re-labelling, we can arrange that $X_1, \ldots, X_k$ are the only isolated elements of the X_i's, for some k between 1 and n. The following result tells us that we can replace the isolated X_i's by equivalent sets such that the new family is in good enough position.

Lemma 3.13. — *Let G denote a finitely generated group, and let $H_1, \ldots, H_n$ be finitely generated subgroups of G. For each $i \geqslant 1$, let X_i be a nontrivial H_i–almost invariant subset of G, such that $X_1, \ldots, X_k$ are the only isolated elements of the X_i's. Then, for each i, $1 \leqslant i \leqslant k$, there is an almost invariant set Z_i equivalent to X_i, such that the translates of all the Z_i's and Z_i^*'s are nested and the family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$ is in good enough position.*

Proof. — If X_i and X_j are two isolated elements of E with stabilisers H_i and H_j respectively, then the almost invariant subsets $H_i \backslash X_i$ of $H_i \backslash G$ and $H_j \backslash X_j$ of $H_j \backslash G$ have intersection number zero with each other, and each has self-intersection number zero. In [**44**], we discussed almost invariant sets with intersection number zero, and the main results of that paper are exactly what we need to understand the present situation. As $H_i \backslash X_i$ has self-intersection number zero, Theorem B.2.8 tells us that X_i is equivalent to Y_i such that Y_i is associated to a splitting of G over a subgroup K_i commensurable with H_i. We will replace each of $X_1, \ldots, X_k$ by $Y_1, \ldots, Y_k$, chosen so that each Y_i is associated to a splitting σ_i of G over K_i. In topological terms, this corresponds to starting with some closed curves C_i on a surface such that each C_i is homotopic to some power of a simple closed curve S_i, and then replacing each C_i by S_i. Note that as Y_i is equivalent to X_i, the splittings $\sigma_1, \ldots, \sigma_k$ have intersection number zero with each other. Now Theorem 2.35 tells us that these splittings are compatible. This means that we can replace the Y_i's by equivalent almost invariant sets Z_i over K_i, whose translates are nested. It follows immediately that the new family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$ is in good enough position. In topological terms, replacing the Y_i's by the Z_i's corresponds to starting with some simple closed curves on a surface such that each pair has intersection number zero, and then replacing the curves by homotopic but disjoint simple closed curves. □

Now we want to define $\Theta(X_1, \ldots, X_n : G)$ to be $\Theta(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$. There is an obvious difficulty here, as we will need to show the resulting structure is independent of the choices of the Z_i's. Recall that we also need to handle the problem of invertible almost invariant sets. It turns out that it will suffice to ensure that the Z_i's are not invertible.

Thus we can now give our plan for constructing the graph of groups structure $\Gamma(X_1, \ldots, X_n : G)$ for G which will turn out to be the algebraic regular neighbourhood of the X_i's. In general terms, we will simply replace the isolated X_i's by equivalent

Z_i's so that each Z_i is not invertible, and the new family is in good position. It remains to show that this construction can be made so that the result is independent of the choices of the Z_i's. This is not automatically the case as the new family may not have the same number of G–orbits in $\overline{E}$ as the original family. If there are distinct i and j such that X_i and X_j are isolated, and $\overline{X_i}$ and $\overline{X_j}$ lie in the same G–orbit in $\overline{E}$, it is possible that $\overline{Z_i}$ and $\overline{Z_j}$ do not lie in the same G–orbit in $\overline{E}$, so that the new family has more G–orbits. On the other hand, if X_i and X_j are isolated and equivalent to each other, the method of the preceding lemma could well yield Z_i equal to Z_j, so that the new family has less G–orbits. In either case, $\Theta(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$ would clearly not give the "correct" structure. In topological terms, this would correspond to replacing a single curve by two homotopic curves or replacing two homotopic simple curves by a single curve. In order to deal with these difficulties, we first consider a special case. It will be convenient to introduce the following terminology. If X_i and X_j are isolated and some translate of X_j is equivalent to X_i or X_i^*, we will say that the G–orbits of $\overline{X_i}$ and $\overline{X_j}$ are *parallel*. We use this word because we are thinking of parallel G–orbits as corresponding to parallel simple closed curves on a surface. We obtain the following result.

Lemma 3.14. — *Let G denote a finitely generated group, and let $H_1, \ldots, H_n$ be finitely generated subgroups of G. For each $i \geqslant 1$, let X_i be a nontrivial H_i–almost invariant subset of G, such that $X_1, \ldots, X_k$ are the only isolated elements of the X_i's. Suppose that each parallelism class of the G–orbits of isolated $\overline{X_i}$'s contains a single element. Then, for $1 \leqslant i \leqslant k$, there is an almost invariant set Z_i equivalent to X_i, such that each Z_i is not invertible, and the family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$ is in good enough position. Further the graph of groups structure $\Theta(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$ is independent of the choices of the Z_i's.*

Remark 3.15. — The uniqueness part of this result can be deduced from Theorem 6.7 where we prove the uniqueness of regular neighbourhoods in general, but it seems worth giving this more direct argument, which does not depend on any of the theory of the next two chapters.

Proof. — We start by applying the statement of Lemma 3.13. The hypothesis that each parallelism class of the G–orbits of isolated $\overline{X_i}$'s contains a single element, implies that the same holds for the $\overline{Z_i}$'s. In particular, there are k distinct G–orbits of the $\overline{X_i}$'s and of the $\overline{Z_i}$'s. It also follows that if σ_i denotes the splitting of G determined by Z_i, then distinct σ_i's are non-conjugate. If some Z_i is invertible, it is trivial to replace it by an equivalent set which is not invertible, but the resulting collection of Z_i's need not be nested. However, as distinct Z_i's are not parallel, the new Z_i's will be almost nested and the family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$ will still be in good enough position. This proves the first part of the lemma.

For the uniqueness result, it will be convenient to consider first the graph of groups structure $\Theta(Z_1, \dots, Z_k : G)$. In the proof of Lemma 3.13, when we replaced X_i by Y_i, we obtained a splitting σ_i to which Y_i is associated. Lemma B.2.3 implies that if two splittings of G have equivalent associated almost invariant sets, then the splittings are conjugate. Thus the splitting σ_i is unique up to conjugacy. Now Theorem 2.35 combined with the assumption that no Z_i is invertible tells us that $\Theta(Z_1, \dots, Z_k : G)$ is determined by the conjugacy classes of the splittings σ_i. It follows that $\Theta(Z_1, \dots, Z_k : G)$ is independent of the choices of the Z_i's.

A more useful way of putting this is the following. Suppose that $W_1, \dots, W_k$ are chosen in the same way as $Z_1, \dots, Z_k$, and are also non-invertible. Then the natural bijection between the set $F(Z) = \{gZ_i, gZ_i^* : g \in G, 1 \leqslant i \leqslant k\}$ and the set $F(W) = \{gW_i, gW_i^* : g \in G, 1 \leqslant i \leqslant k\}$ is order preserving. Now we consider the sets $E(Z)$ and $E(W)$ obtained from E by replacing each X_i by Z_i or by W_i, $1 \leqslant i \leqslant k$. The natural bijection between these two sets extends the bijection between $F(Z)$ and $F(W)$. It is the identity on the translates of $X_{k+1}, \dots, X_n$ and so trivially preserves the partial order on the translates of these elements. Thus it remains to check that our bijection is order preserving when we compare a translate of one of these elements with a translate of some Z_j. Suppose, for example, that $gZ_j \leqslant hX_l$, where $l > k$, so that $gZ_j \cap hX_l^*$ is small. As Z_j and W_j are equivalent, the intersection $gW_j \cap hX_l^*$ must also be small. We know that gW_j and hX_l are comparable, as the family $W_1, \dots, W_k, X_{k+1}, \dots, X_n$ is in good enough position. Thus either $gW_j \cap hX_l^*$ is the only small set out of the four sets $gW_j^{(*)} \cap hX_l^{(*)}$, or there is a second small set and one of the two small sets is empty. As gW_j and hX_l, and their complements, are not small, the second small set could only be $gW_j^* \cap hX_l$. Thus we must have either $gW_j \leqslant hX_l$, or gW_j and hX_l are equivalent. If the second case occurred, the fact that W_j is isolated would imply that X_l was also isolated, which contradicts our definition of k and the fact that $l > k$. It follows that $gW_j \leqslant hX_l$. This shows that the natural bijection between the sets $E(Z)$ and $E(W)$ is order preserving, and hence that $\Theta(Z_1, \dots, Z_k, X_{k+1}, \dots, X_n : G)$ is naturally isomorphic to $\Theta(W_1, \dots, W_k, X_{k+1}, \dots, X_n : G)$ as required. This completes the proof of the lemma. □

This result gives us a well defined construction of our regular neighbourhood $\Gamma(X_1, \dots, X_n : G)$ on the assumption that each parallelism class of the G–orbits of isolated $\overline{X_i}$'s contains a single element. Now we are ready to complete the construction of $\Gamma(X_1, \dots, X_n : G)$ in the general case. As discussed earlier, by replacing the X_i's by a subfamily, we can assume that distinct $\overline{X_i}$'s lie in distinct G–orbits. Next we choose a subfamily of the X_i's by selecting all the non-isolated X_i's and one isolated X_i from each parallelism class. We will renumber so that this subfamily is $X_1, \dots, X_m$. By our choice, this subfamily satisfies the hypotheses of Lemma 3.14. This lemma gives us a well defined regular neighbourhood $\Gamma(X_1, \dots, X_m : G)$. Now we add back $X_{m+1}, \dots, X_n$ to the family. Let X_i denote one of these sets, so that $i > m$.

Recall that $\overline{X_i}$ is equivalent to a translate of an isolated $\overline{X_j}$, for some $j \leqslant m$. The construction above shows that X_j is equivalent to an almost invariant non-invertible set Z_j associated to a splitting, and that $\Gamma(X_1,\ldots,X_m : G)$ has a V_0–vertex v_j which comes from the CCC which consists only of $\overline{Z_j}$. Proposition 5.4 implies that v_j has valence 2, and the fact that Z_j is non-invertible implies that each of the incident edge groups includes by an isomorphism into the vertex group $G(v_j)$. (Note that Proposition 5.4 has not yet been proved. However, the half of the result which we need has a proof which could have been given now, so there is no logical problem in using it.) We now construct $\Gamma(X_1,\ldots,X_m,X_i : G)$ from $\Gamma(X_1,\ldots,X_m : G)$ by subdividing each of the two edges incident to v_j, to obtain four edges and three interior vertices, each with associated group $G(v)$. The two new vertices are V_0–vertices and the original V_0–vertex becomes a V_1–vertex. Effectively what this does is to replace X_i and X_j by the same almost invariant set Z_j, which we then use twice in the construction of the regular neighbourhood. The topological analogue is to replace a single simple closed curve by two parallel curves. If several of $X_{m+1},\ldots,X_n$ are equivalent to the same X_j, we split the two edges incident to v_j several times, in order to obtain the requisite number of V_0–vertices. By repeating this process for each of $X_{m+1},\ldots,X_n$ we will construct the regular neighbourhood $\Gamma(X_1,\ldots,X_n : G)$ from $\Gamma(X_1,\ldots,X_m : G)$.

We can summarise our construction of $\Gamma(X_1,\ldots,X_n : G)$ as follows.

Summary 3.16. — Let G denote a finitely generated group, and let $H_1,\ldots,H_n$ be finitely generated subgroups of G. For each $i \geqslant 1$, let X_i be a nontrivial H_i–almost invariant subset of G, and let $E = \{gX_i, gX_i^* : g \in G, 1 \leqslant i \leqslant n\}$. There are four cases to our construction of $\Gamma(X_1,\ldots,X_n : G)$.

(1) If the X_i's are in good position, we have a partial order $\leqslant$ on E and can then directly construct $\Theta(X_1,\ldots,X_n : G)$ using a pretree as in Theorem 3.8. If in addition, no isolated X_i is invertible, we define $\Gamma(X_1,\ldots,X_n : G)$ to be $\Theta(X_1,\ldots,X_n : G)$.

(2) If the X_i's are in good enough position, we can make essentially the same construction as in Case 1) using the same partial order where it is defined.

(3) If the X_i's are not in good enough position or if some isolated X_i is invertible, we pick one X_i for each G–orbit in $\overline{E}$, thus replacing the X_i's by a subfamily with the same G–orbits in $\overline{E}$ such that distinct $\overline{X_i}$'s lie in distinct G–orbits. By renumbering, we can assume that $X_1,\ldots,X_k$ are the only isolated X_i's.

(a) Suppose that each parallelism class of the G–orbits of isolated $\overline{X_i}$'s contains a single element. Then for each i such that $1 \leqslant i \leqslant k$, we replace the isolated set X_i by an equivalent almost invariant set Z_i, which is not invertible, such that the new family $Z_1,\ldots,Z_k,X_{k+1},\ldots,X_n$ is in good enough position, and we define $\Gamma(X_1,\ldots,X_n : G)$ to be $\Theta(Z_1,\ldots,Z_k,X_{k+1},\ldots,X_n : G)$.

(b) In general, we choose a subfamily of the X_i's by selecting all the non-isolated X_i's and one isolated X_i from each parallelism class. We will renumber so that this subfamily is $X_1,\ldots,X_m$. Then we apply Case 3a) to construct

$\Gamma(X_1,\ldots,X_m : G)$. Finally, we construct $\Gamma(X_1,\ldots,X_n : G)$ by subdividing those edges of $\Gamma(X_1,\ldots,X_m : G)$ which are incident to V_0–vertices arising from isolated X_i's. For each isolated X_i, the corresponding V_0–vertex is replaced by a number of V_0–vertices which equals the number of elements in the parallelism class of $\overline{X_i}$.

This allows us to summarise our construction rather more briefly as follows:

First we omit some X_i's and replace others by equivalent almost invariant sets such that the new family is in good enough position and no isolated X_i is invertible. Now apply the direct construction using a pretree to the new family to obtain a graph of groups structure Γ for G. Finally, $\Gamma(X_1,\ldots,X_n : G)$ is constructed from Γ by subdividing certain edges of Γ.

The above arguments raise the more general question of how our regular neighbourhood construction changes when one replaces X_i's, which need not be isolated, by equivalent sets. This is an important question because usually one is not very interested in a particular almost invariant set but rather in its equivalence class. Our next result shows that our regular neighbourhood construction essentially depends only on the equivalence classes of the X_i's.

Lemma 3.17. — *Let G denote a finitely generated group, and let $H_1,\ldots,H_n$ be finitely generated subgroups of G. For each $i \geqslant 1$, let X_i be a nontrivial H_i–almost invariant subset of G and let W_i be a nontrivial K_i–almost invariant subset of G which is equivalent to X_i. Suppose that the correspondence between the isolated X_i's and the isolated W_i's induces a bijection between the G–orbits of isolated $\overline{X_i}$'s and the G–orbits of isolated $\overline{W_i}$'s. Then $\Gamma(X_1,\ldots,X_n : G)$ and $\Gamma(W_1,\ldots,W_n : G)$ are isomorphic.*

Remark 3.18. — There need not be a bijection between the G–orbits of non-isolated $\overline{X_i}$'s and the G–orbits of non-isolated $\overline{W_i}$'s, but that does not affect the result.

Proof. — Our construction of $\Gamma(X_1,\ldots,X_n : G)$ and $\Gamma(W_1,\ldots,W_n : G)$ shows that it will suffice to handle the case when each parallelism class of G–orbits of isolated $\overline{X_i}$'s has a single element. This is the situation of Lemma 3.14. Let E denote the set of all translates of the X_i's and let F denote the set of all translates of the W_i's. As the stabilisers of X_i and W_i are commensurable but need not be equal, there is no natural map between E and F, but we will show that there is a natural bijection between the CCC's of $\overline{E}$ and $\overline{F}$, which preserves betweenness.

As in Lemma 3.14, we will renumber so that $X_1,\ldots,X_k$ are the only isolated elements of the X_i's. Thus $W_1,\ldots,W_k$ are the only isolated elements of the W_i's. As in our construction of an algebraic regular neighbourhood, we can replace $X_1,\ldots,X_k$ by equivalent sets $Z_1,\ldots,Z_k$ such that the translates of the Z_i's are nested, the Z_i's are not invertible, and the family $Z_1,\ldots,Z_k,X_{k+1},\ldots,X_n$ is in good enough position. We can also replace $W_1,\ldots,W_k$ by the same sets $Z_1,\ldots,Z_k$. This is because of the technical hypothesis on the isolated X_i's and W_i's. Lemma 3.14 allows us to define

$\Gamma(X_1, \ldots, X_n : G)$ as $\Theta(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$ and to define $\Gamma(W_1, \ldots, W_n : G)$ as $\Theta(Z_1, \ldots, Z_l, W_{k+1}, \ldots, W_n : G)$. Thus, by changing notation, we can suppose that in our original families $X_1, \ldots, X_n$ and $W_1, \ldots, W_n$, we have $X_i = W_i$, for $1 \leqslant i \leqslant k$.

Now if U is an isolated element of E, the CCC which contains $\overline{U}$ consists only of $\overline{U}$. Conversely, if a CCC consists of a single element $\overline{U}$ of $\overline{E}$, then U must be isolated. We call such a CCC *isolated*. There is an obvious bijection between the isolated CCC's of $\overline{E}$ and the isolated CCC's of $\overline{F}$. Further, this bijection preserves betweenness for isolated CCC's.

Now consider a non-isolated X_i. If gX_i is equivalent to X_i, then $g\overline{X_i}$ must lie in the CCC $[\overline{X_i}]$ of $\overline{E}$. As W_i is equivalent to X_i, it follows that gW_i is equivalent to W_i and hence that $g\overline{W_i}$ must lie in the CCC $[\overline{W_i}]$ of $\overline{F}$. If kX_j crosses X_i, then kW_j crosses W_i. It follows that if S_{ij} denotes the collection of elements s of G such that sX_j lies in $[\overline{X_i}]$, then S_{ij} must also equal the collection of elements s of G such that sW_j lies in $[\overline{W_i}]$. This yields a natural bijection between the non-isolated CCC's of $\overline{E}$ and the non-isolated CCC's of $\overline{F}$. Further, it is clear that this bijection preserves betweenness for non-isolated CCC's.

It follows that we have a natural bijection between the CCC's of $\overline{E}$ and the CCC's of $\overline{F}$, and that this preserves betweenness except possibly when we consider a mixture of isolated CCC's and non-isolated CCC's. As in the proof of the previous lemma, it is easy to see that betweenness must be preserved here also. It follows that $\Gamma(X_1, \ldots, X_n : G)$ and $\Gamma(W_1, \ldots, W_n : G)$ are naturally isomorphic, as required. $\square$

Our regular neighbourhood construction always expresses G as the fundamental group of a graph $\Gamma(X_1, \ldots, X_n : G)$ of groups, but this graph may consist of a single point. This occurs precisely when the set $\overline{E}$ has only one cross-connected component, so that P and hence T consists of a single point. We give two examples of situations where this will occur. Our first example is from the topology of 3–manifolds and is due to Rubinstein and Wang [**37**].

Example 3.19. — In this case $n = 1$, and we denote X_1 by X and the stabiliser of X by H. The group G is the fundamental group of a closed graph manifold M, and the subgroup H is isomorphic to the fundamental group of a closed surface F. One can choose a π_1–injective map $f : F \to M$ so that $f_*\pi_1(F) = H$, and F lifts to an embedding in the cover M_F of M with fundamental group equal to H. Considering one side of F in M_F determines the H–almost invariant subset X of G. Rubinstein and Wang show that, for many choices of the manifold M, the surface F cannot lift to an embedding in any finite cover of M. They do this by showing that the pre-image of F in the universal cover $\widetilde{M}$ of M consists of a family of embedded planes such that any two cross in the sense of [**19**]. This implies that any two distinct translates of X cross, so that the set $\overline{E}$ has only one cross-connected component.

Our second example is also rather special.

Example 3.20. — Let H denote any finitely generated group, and let G denote the group $H \times \mathbb{Z}$. Let X denote the H–almost invariant subset of G associated to the splitting of G as the HNN-extension $H*_H$. Thus the translates of X by G are all equivalent to X. In particular, none of these translates cross each other. Now suppose that there is a subgroup K of G and a K–almost invariant subset Y of G such that Y crosses X. Then Y crosses every translate of X, and hence also X crosses every translate of Y. It follows that if we let E denote the set of all translates of X and Y, then the set $\overline{E}$ has only one cross-connected component. A simple way to generate such examples of K and Y is to choose $H = A *_C B$, choose $K = C \times \mathbb{Z}$ and to choose Y to be associated to the splitting $G = (A \times \mathbb{Z}) *_K (B \times \mathbb{Z})$.

We have now completed discussing our regular neighbourhood construction when one is given a finite family of nontrivial almost invariant subsets of a group G. We will end this chapter by discussing what happens if one is given an infinite collection of such subsets, as this will play an important role in this paper. At first glance this may seem to be a very unreasonable thing to consider. In topology, one never discusses regular neighbourhoods of infinite collections of submanifolds. But if one considers a subsurface N of a surface M, then N contains curves representing each element of $\pi_1(N)$, and we want to regard N as a regular neighbourhood of this infinite family of curves. Of course, such an idea cannot make sense for arbitrary infinite families of curves. For example, an infinite collection of disjoint essential circles in M cannot reasonably be said to have a regular neighbourhood.

Now let G denote a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. We will proceed as we did earlier in this chapter, and just note the differences in the case when Λ is infinite.

As before, we let E denote the collection of all translates of the X_λ's and their complements. First we will assume that the X_λ's are in good position, i.e. that E satisfies Condition (*). This allows us to define the partial order $\leqslant$ on E exactly as before. The first and crucial difference between the infinite case and the finite case occurs when we consider Lemma 3.1. While parts 1) and 3) still hold, part 2) need not hold, i.e. E need not be discrete, so that there may be elements U and V of E with infinitely many elements of E between them. The analogous situation occurs in topology if one considers infinitely many disjoint simple closed curves on a surface. However, we can still define the set $\overline{E}$ of pairs $\{X, X^*\}$ for $X \in E$, and can define P to be the collection of all CCC's of $\overline{E}$. Further the arguments following Lemma 3.1 still apply, and show that the idea of betweenness can be defined on P as before. Now we consider the proof of Theorem 3.8. The argument for the first part which asserts that P is a pretree remains correct. But the second part which asserts that P is discrete depends on the discreteness of E, and so P may not be discrete. In fact, if there are infinitely many X_λ's which are all equivalent, then P

will not be discrete. However, if it happens that P is discrete, then as before P can be embedded in a G–tree T and $G\backslash T$ is a graph of groups structure for G which we will denote by $\Theta(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. Further, if no isolated X_λ is invertible, we define $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ to be $\Theta(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. Of course, the V_0–vertex groups of this graph need not be finitely generated, and it appears that it may have infinitely many V_0–vertices. However, as G is finitely generated, it will follow from Proposition 5.2 that $\Theta(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is always a finite graph.

Having dealt with the case when the X_λ's are in good position, we say that the family $\{X_\lambda\}_{\lambda\in\Lambda}$ is *in good enough position* if whenever we find incomparable elements U and V of E which do not cross, there is some element W of E which crosses them. As in the finite case, the above discussion applies equally well if the X_λ's are in good enough position. If this condition does not hold, or if some isolated X_λ is invertible, we want to apply the proof of Lemma 3.14 and the discussion following. The arguments work exactly as before, so long as E has only finitely many G–orbits of isolated elements. If E has infinitely many such, then it is not possible to construct a regular neighbourhood of the X_λ's. For then $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ would have to have infinitely many V_0–vertices, one for each G–orbit of isolated elements, whereas Proposition 5.2 implies that $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is finite. Note that it would be simpler to replace the condition that E has only finitely many G–orbits of isolated elements by the condition that only finitely many of the X_λ's are isolated. But this second condition is more restrictive than needed.

Finally the proof of Lemma 3.17 still applies to show that our construction essentially depends only on the equivalence classes of the X_λ's.

We summarise our conclusions as follows. Suppose that we are given a finitely generated group G, finitely generated subgroups H_λ, $\lambda \in \Lambda$, and a nontrivial H_λ–almost invariant subset X_λ of G. Suppose that E has only finitely many G–orbits of isolated elements, so that after replacing the X_λ's by equivalent sets we can assume that they are in good enough position. Then one can define the idea of betweenness on the set P of all CCC's of $\overline{E}$, and P is a pretree. If P is discrete, one can construct an algebraic regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ of the X_λ's, and this depends only on the equivalence classes of the X_λ's.

CHAPTER 4

ENCLOSING

In this chapter we will consider graphs of groups in general. We will discuss the idea of a vertex of a graph of groups enclosing an almost invariant set. In the following chapter, we will apply these ideas to our regular neighbourhood construction.

Let Γ be a graph of groups and write $\pi_1(\Gamma)$ for the fundamental group of Γ. We emphasise that although we will mostly consider situations where $\pi_1(\Gamma)$ is finitely generated, we will not assume that the edge and vertex groups of Γ are finitely generated unless this is specifically stated. To avoid some degeneracy phenomena, we will often assume that Γ is minimal, meaning that for any proper connected subgraph K, the natural inclusion of $\pi_1(K)$ into $\pi_1(\Gamma)$ is not an isomorphism. Note that if Γ is minimal and $\pi_1(\Gamma)$ is finitely generated, then Γ must be finite.

If Γ' is a (not necessarily connected) subgraph of Γ, we say that a graph of groups structure Γ_1 is *obtained from* Γ *by collapsing* Γ' if the underlying graph of Γ_1 is obtained from Γ by collapsing each component of Γ' to a point. In addition, if $p : \Gamma \to \Gamma_1$ denotes the natural projection map, we require that each vertex v of Γ_1 has associated group equal to $\pi_1(p^{-1}(v))$. These conditions imply that $\pi_1(\Gamma)$ and $\pi_1(\Gamma_1)$ are naturally isomorphic. Two important special cases of this construction occur when we consider an edge e of Γ. If the subgraph Γ' equals e, we say that Γ_1 is *obtained from* Γ *by collapsing* e. If the subgraph Γ' equals the complement of the interior of e, then Γ_1 has a single edge which determines a splitting σ of $\pi_1(\Gamma)$, and we call σ *the splitting of* $\pi_1(\Gamma)$ *associated to* e. Note that so long as we assume that Γ is minimal then σ really is a splitting, i.e. a nontrivial decomposition of $\pi_1(\Gamma)$. Such splittings of G will be referred to as the edge splittings of Γ.

It will be very convenient to introduce some terminology to describe the process which is the reverse of collapsing an edge. If a graph of groups structure Γ_1 for a group G is obtained from a graph of groups structure Γ by collapsing an edge e, and if e projects to the vertex v_1 of Γ_1, we will say that Γ is a *refinement* of Γ_1 obtained by splitting at the vertex v_1.

Definition 4.1. — Let Γ_1 be a graph of groups structure for a group G, and let σ denote a splitting of G. Then the vertex v_1 of Γ_1 *encloses* σ if there is a refinement Γ of Γ_1 obtained by splitting Γ_1 at v_1, such that the edge e of Γ which projects to the vertex v_1 has σ as its associated splitting.

To understand the reasons for our terminology, the reader should consider a subsurface N of a surface M, and let Γ_1 be the graph of groups structure for $G = \pi_1(M)$ determined by ∂N. Let C be some simple closed curve in N, and let Γ be the graph of groups structure for G determined by $\partial N \cup C$. If e denotes the edge of Γ which corresponds to C, then Γ_1 is obtained from Γ by collapsing e, and the vertex v_1 of Γ_1 corresponds to the component of N which contains C. Thus saying that v_1 encloses the splitting of G associated to C mirrors the fact that N contains C.

Next we introduce a little more terminology. We will say that a vertex v of Γ is *redundant* if it has valence at most two, it is not the vertex of a loop, and each edge group includes by an isomorphism into the vertex group at v. Now assume that Γ is minimal. Then a redundant vertex must have valence two. Clearly, these two edges determine conjugate edge splittings of G. Conversely, it is easy to see that if Γ has two edges with conjugate edge splittings, then these edges are the end segments of a path all of whose interior vertices are redundant. If Γ has a redundant vertex v, we can amalgamate the two edges incident to v into a single edge to obtain a new graph of groups structure for G. If Γ is finite, we can repeat this to obtain a graph of groups structure Γ' for G with no redundant vertices. Clearly Γ is obtained from Γ' by subdividing some edges.

In Definition 4.1, we defined what it means for a vertex v of a graph of groups Γ to enclose a splitting of $G = \pi_1(\Gamma)$. Next we want to extend this notion to define what it means for v to enclose an almost invariant subset of G. This is meant to be an analogue of the topological idea of a subsurface containing a possibly singular curve. In order to avoid problems with conjugates, it is better to consider a G–tree T rather than the quotient graph of groups Γ. The condition that $\Gamma = G\backslash T$ be minimal in the sense above is equivalent to the condition that T have no proper G–invariant subtree. Such a G–tree is also called minimal. Note that any G–tree possesses a minimal subtree T_0. If G fixes more than one vertex of T, then T_0 is not unique, but otherwise T_0 is unique and can be described simply as the intersection of all the G–invariant subtrees of T. Note that a minimal G–tree has no vertices of valence one. For if a G–tree has such vertices, one can obtain a proper G–invariant subtree by simply removing each such vertex together with the interior of the incident edge.

We recall some notation from chapter 2. An oriented edge s of a tree T determines a natural partition of $V(T)$ into two sets, namely the vertices of the two subtrees obtained by removing the interior of s from T. Let Y_s denote the collection of all the vertices of the subtree which contains the terminal vertex v of s, and let Y_s^* denote the complementary collection of vertices. If a group G acts without inversions on T, then choosing a basepoint w in T determines a G–equivariant map $\varphi : G \to V(T)$ given by the formula $\varphi(g) = gw$. We define the sets $Z_s = \varphi^{-1}(Y_s)$ and $Z_s^* = \varphi^{-1}(Y_s^*)$. Lemma 2.10 shows that, if S denotes the stabiliser of s, then Z_s is S–almost invariant, and its equivalence class is independent of the choice of basepoint w.

Note that all the results in this chapter also hold if G acts on T with inversions. However, in all our applications T will be the universal covering G–tree of a graph of groups, so that G will automatically act without inversions.

We now define enclosing of almost invariant sets.

Definition 4.2. — Let A be a nontrivial H–almost invariant subset of a group G and let v a vertex of a G–tree T. Choose a basepoint w for T. For each edge s of T, this determines the subsets Z_s and Z_s^* of G as above. We say that the vertex v *encloses* A, if for all edges s of T which are incident to v and directed towards v, we have $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small.

Remark 4.3. — This definition is independent of the choice of basepoint w, because changing w replaces each Z_s by an equivalent almost invariant set.

Note that it does not make much sense to consider enclosing of trivial almost invariant subsets of G, because any such subset of G would automatically be enclosed by every vertex of T.

It would seem more natural to say that v encloses A, if for all edges s of T which are incident to v and directed towards v, we have $A^* \geqslant Z_s^*$ or $A \geqslant Z_s^*$, but we want to ensure that any set equivalent to Z_s or Z_s^* is enclosed by v, and not all such sets are comparable with Z_s. See part 4) of Lemma 4.6 for precise statements.

It will also be convenient to define enclosing by a vertex of a graph of groups.

Definition 4.4. — Let A be a nontrivial H–almost invariant subset of a group G, let T be a G–tree and let Γ denote the associated graph of groups structure for G with underlying graph $G \backslash T$. We say that a vertex u of Γ *encloses* A if there is a vertex v of T which encloses A and projects to u.

Remark 4.5. — If u encloses A, then it also encloses any translate of A and any almost invariant set equivalent to A.

We now have two natural ideas of what it means for a vertex u of Γ to enclose a splitting σ of G over a subgroup H. One is given in Definition 4.1, and the other is that a H–almost invariant set associated to σ is enclosed by u. In Lemma 4.10, we will show that these ideas are equivalent when H is finitely generated. In Lemma 5.10 we will be able to show that this equivalence holds even when H is not finitely generated.

We will need the following basic properties of enclosing almost invariant sets, which we state in several lemmas. Again we emphasise that we are not assuming that the subgroup H of G is finitely generated. Note that several of the statements below require a choice of basepoint for T, but are independent of that choice.

Lemma 4.6. — *Let A be a nontrivial H–almost invariant subset of a group G and let v a vertex of a G–tree T. Then the following statements all hold:*

(1) *A is enclosed by v if and only if A^* is enclosed by v.*

(2) *If B is an almost invariant set equivalent to A, then v encloses A if and only if v encloses B.*

(3) *If s is an edge of T with stabiliser S, and if Z_s is a nontrivial S–almost invariant subset of G, then Z_s is enclosed by each of the two vertices to which s is incident.*

(4) *A is enclosed by v if and only if either*

(a) *A is equivalent to Z_s or Z_s^* for some edge s incident to v, or*

(b) *we have $A^* \geqslant Z_s^*$ or $A \geqslant Z_s^*$ for every edge s incident to v and oriented towards v.*

(5) *If A is enclosed by v, then for any edge t of T which is oriented towards v, we have $A \cap Z_t^*$ or $A^* \cap Z_t^*$ is small.*

Proof. — Parts 1), 2) and 3) are all trivial.

4) First suppose that A is enclosed by v, so that for all edges s of T which are incident to v and directed towards v, we have $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small. If two of the four sets $A^{(*)} \cap Z_s^{(*)}$ are small, then A is equivalent to Z_s or Z_s^*. Otherwise, we have $A^* \geqslant Z_s^*$ or $A \geqslant Z_s^*$ for every edge s incident to v and oriented towards v, as required.

Conversely, if $A^* \geqslant Z_s^*$ or $A \geqslant Z_s^*$ for every edge s incident to v and oriented towards v, then $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small for each such edge s, so that A is enclosed by v. And if A is equivalent to Z_s or Z_s^* for some edge s incident to v, then parts 2) and 3) imply that A is enclosed by v. The result follows.

5) Consider the oriented path in T which joins t to v. Let s be the edge of this path incident to v. As s is oriented towards v and A is enclosed by v, it follows that $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small. As $t \leqslant s$ in the natural ordering on oriented edges of T, it follows that $Z_s \subset Z_t$, so that $A \cap Z_t^*$ or $A^* \cap Z_t^*$ is small, as required. □

Lemma 4.7. — *Let A and B be nontrivial H–almost invariant subsets of a group G such that each is enclosed by a vertex v of a G–tree T. Then $A \cup B$, $A \cap B$ and $A + B$ are also H–almost invariant and each is also enclosed by v, if it is nontrivial.*

Proof. — It is clear that $A \cup B$, $A \cap B$ and $A + B$ are H–almost invariant. Let s be an edge of T incident to v and oriented towards v. Thus $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small, and $B \cap Z_s^*$ or $B^* \cap Z_s^*$ is small. We will examine the four possibilities in turn, and show that if C denotes any one of $A \cup B$, $A \cap B$ and $A + B$, then either $C \cap Z_s^*$ or $C^* \cap Z_s^*$ is small. As this holds for every such edge s, it will follow that if C is a nontrivial H–almost invariant subset of G it is enclosed by v, as required.

Case 1: $A \cap Z_s^*$ and $B \cap Z_s^*$ are both small.

Then the union of these sets, $(A \cup B) \cap Z_s^*$, is small. As $A \cap B$ and $A + B$ are each contained in $A \cup B$, it follows that $(A \cap B) \cap Z_s^*$ and $(A + B) \cap Z_s^*$ are also small, completing the proof in this case.

Case 2: $A \cap Z_s^*$ and $B^* \cap Z_s^*$ are both small.

As $(A \cup B)^* \subset B^*$, it follows that $(A \cup B)^* \cap Z_s^*$ is small. As $A \cap B \subset A$, it follows that $(A \cap B) \cap Z_s^*$ is small. As $(A + B)^* \subset A \cup B^*$, it follows that $(A + B)^* \cap Z_s^*$ is small, completing the proof in this case.

Case 3: $A^* \cap Z_s^*$ and $B \cap Z_s^*$ are both small.

We apply the preceding paragraph with the roles of A and B interchanged. Again $(A \cup B)^* \cap Z_s^*$, $(A \cap B) \cap Z_s^*$ and $(A + B)^* \cap Z_s^*$ are small, completing the proof in this case.

Case 4: $A^* \cap Z_s^*$ and $B^* \cap Z_s^*$ are both small.

As $(A \cap B)^* = A^* \cup B^*$, it follows that $(A \cap B)^* \cap Z_s^*$ is small. As $(A \cup B)^* \subset (A \cap B)^*$, it follows that $(A \cup B)^* \cap Z_s^*$ is also small. As $(A + B) \subset (A^* \cup B^*)$, it follows that $(A + B) \cap Z_s^*$ is small.

This completes the proof of the lemma. □

Lemma 4.8. — *Let G be a group, and let T be a G-tree.*

(1) *If s is an edge of T, and s has stabiliser S, then Z_s is a nontrivial S-almost invariant subset of G if and only if s lies in the minimal subtree of T.*

(2) *If a vertex v of T is not fixed by G, then v encloses some nontrivial H-almost invariant subset of G if and only if v lies in the minimal subtree of T.*

Proof. — 1) Let T_0 denote the minimal subtree of T. If G fixes more than one vertex of T so that T_0 is not unique, we let T_0 denote one of the vertices fixed by G.

First suppose that s does not lie in T_0. Note that in the special case when T_0 is not unique, this condition is automatic as then T_0 has no edges. We will show that Z_s must be trivial. We can assume that s is oriented towards T_0. As we are free to choose the basepoint w for T, we will choose it to lie in T_0. This implies that $\varphi(G)$, which equals the orbit of w, also lies in T_0. Hence $Z_s = G$, so that Z_s is a trivial S-almost invariant set as claimed.

Now suppose that s lies in T_0. Note that in this case, T_0 must be unique. Lemma A.3.3 tells us that s lies in T_0 if and only if there exists an element g of G such that s and gs are distinct and coherently oriented. (Two oriented edges in T are coherently oriented if there is an oriented path in T which begins with one and ends with the other.) By repeatedly applying g or g^{-1}, we see that on each side of s in T there are translates of s which are arbitrarily far from s. It follows that $\varphi(G)$ contains points in Y_s and in Y_s^* which are arbitrarily far from s. This immediately implies that Z_s must be nontrivial, as required.

2) We choose T_0 as in part 1), and recall that we are assuming that v is not fixed by G. It follows that if v lies in T_0, there is an edge s of T_0 incident to v. Part 1) implies that Z_s is a nontrivial S–almost invariant set, and part 3) of Lemma 4.6 implies that Z_s is enclosed by v.

If v does not lie in T_0 and v encloses some nontrivial H–almost invariant subset A of T, we will obtain a contradiction. Let s denote the edge of T which is incident to v and on the path joining v to T_0 and we choose s to be oriented towards v. If we choose the basepoint w of T to lie in T_0, we will have $Z_s^* = G$. But as A is enclosed by v, we have $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small, which implies that A or A^* is small. It follows that A is a trivial H–almost invariant set, which is the required contradiction. □

Lemma 4.9. — *Let A be a nontrivial H–almost invariant subset of a group G and let T be a G–tree.*

(1) *If A is enclosed by two distinct vertices u and v of T, then A is equivalent to Z_s or to Z_s^* for each edge s on the path λ joining u and v. Further λ is contained in the minimal subtree T_0 of T, and each interior vertex of λ has valence 2 in T_0.*

(2) *If A is enclosed by a vertex v of T, then $Hv = v$.*

Proof. — 1) Let λ denote the path in T which joins u and v. Let l denote the edge of λ incident to u and oriented towards u, and let m denote the edge of λ incident to v and oriented towards v. Our choice of orientations on l and m implies that $Z_l^* \supset Z_m$. As A is enclosed by u, we know that $A \cap Z_l^*$ or $A^* \cap Z_l^*$ is small. Without loss of generality, we can assume that $A \cap Z_l^*$ is small. As $Z_l^* \supset Z_m$, it follows that $A \cap Z_m$ is also small. As A is enclosed by v, we know that $A \cap Z_m^*$ or $A^* \cap Z_m^*$ is small. But if $A \cap Z_m^*$ were small, then A would itself be small, which contradicts our hypothesis that A is nontrivial. It follows that $A^* \cap Z_m^*$ must be small. As $A \cap Z_m$ and $A^* \cap Z_m^*$ are both small, it follows that A is equivalent to Z_m^*. Similarly, A must be equivalent to Z_l. If n denotes any edge of λ oriented towards u, we have the inclusions $Z_l \subset Z_n \subset Z_m^*$, and it immediately follows that A is also equivalent to Z_n. The fact that λ is contained in the minimal subtree T_0 of T follows at once from part 1) of Lemma 4.8.

Finally, let z denote an interior vertex of λ, and let r and s denote the edges of λ incident to z and oriented towards z. The above discussion implies that Z_r and Z_s^* are equivalent. In particular, their stabilisers are commensurable. Let K denote $\mathrm{Stab}(r) \cap \mathrm{Stab}(s)$, so that Z_r and Z_s^* are equivalent K–almost invariant subsets of G. Thus $Z_r \cap Z_s$ is K–finite. If there is another edge t of T incident to z, we orient t towards z also. As K fixes z, for each $k \in K$, the edge $k(t)$ is also incident to z and oriented towards z. Thus $Z_{k(t)}^*$ is disjoint from Z_r^* and Z_s^*, and so is contained in $Z_r \cap Z_s$. Hence $\cup_{k \in K} Z_{k(t)}^*$ is also K–finite. As the $Z_{k(t)}^*$ are disjoint from each other, it follows that Z_t^* is $\mathrm{Stab}(t)$-finite. Equivalently Z_t^* is a trivial almost invariant set over $\mathrm{Stab}(t)$. Now part 1) of Lemma 4.8 implies that t does not lie in the minimal subtree T_0 of T. Hence every interior vertex of λ has valence 2 in T_0 which completes the proof of part 1).

2) As A is enclosed by v, it follows that gA is enclosed by gv. If g lies in H, then $gA = A$, so that A is enclosed by both v and gv. If $gv = v$ for all $g \in H$, then $Hv = v$, as required.

Now suppose that there is g in H such that gv is not equal to v. Then part 1) implies that A is equivalent to Z_s or to Z_s^* for each edge s on the path λ joining v and gv. Further λ is contained in the minimal subtree T_0 of T, and each interior vertex of λ has valence 2 in T_0. Fix an edge s of λ oriented so that A is equivalent to Z_s, and let S denote the stabiliser of s. Thus H and S must be commensurable. In particular, as g lies in H, some power of g lies in S. As $gA = A$, and A is equivalent to Z_s, it follows that gZ_s is equivalent to Z_s. Let L denote the maximal subinterval of T_0 which contains λ and has all interior vertices of valence 2 in T_0. As gL meets L, we must have $gL = L$. If g reversed the orientation of L, we would have $gZ_s \subset Z_s^*$ or $gZ_s^* \subset Z_s$, but either of these is impossible as gZ_s is equivalent to Z_s. Hence g preserves the orientation of L. As g does not act on L by the identity, it follows that L is doubly infinite and that g acts on L by a nontrivial translation. But this contradicts the fact that some power of g lies in S and so fixes s. This contradiction completes the proof. □

Now we are ready to show that the two ideas of enclosing a splitting are equivalent.

Lemma 4.10. — *Let T be a G–tree, let Γ denote the graph of groups structure for G given by the quotient $G\backslash T$, and let u denote a vertex of Γ. Suppose that A is associated to a splitting σ of G over H. Then the following statements hold.*

(1) *If σ is enclosed by u, then A is enclosed by u.*

(2) *If H is finitely generated and A is enclosed by u, then σ is enclosed by u.*

Remark 4.11. — In Lemma 5.10 we will show that part 2) holds even when H is not finitely generated.

Proof. — 1) Suppose that σ is enclosed by u. Thus there is a graph of groups Γ_1 which is a refinement of Γ obtained by splitting at u so that the extra edge has σ as its associated edge splitting. Let $q : \Gamma_1 \to \Gamma$ denote the projection map, and let e denote the extra edge of Γ_1 so that $q(e) = u$. Let v denote a vertex in the pre-image of u in the G–tree T, and let T_1 denote the universal covering G–tree of Γ_1. Thus T_1 is obtained from T by splitting at each vertex in the orbit of v. Let e also denote the extra edge inserted at v. Then T_1 has an induced projection $p : T_1 \to T$ such that $p(ge) = gv$. Pick a basepoint w for T_1, and choose $p(w)$ to be the basepoint for T. Consider the set E_1 of all the sets Z_s, for each oriented edge s of T_1. If s is not equal to ge for any g, then $p(s)$ is an oriented edge of T and $Z_s = Z_{p(s)}$. There is a translate ge of e such that Z_{ge} is equivalent to A or A^*. We will choose e so that this translate is e itself. As E_1 is nested, we know that for any edge s of $T_1 - \{e\}$, oriented towards e, we have $Z_e \supset Z_s^*$ or $Z_e^* \supset Z_s^*$. Turning to the tree T, it follows that either A is equivalent to Z_s or Z_s^* for some edge s incident to v, or we have $A^* \geqslant Z_s^*$ or $A \geqslant Z_s^*$ for every edge s incident to v and oriented towards v. Now part 4) of Lemma 4.6 implies that A is enclosed by v and hence by u.

2) Suppose that A is enclosed by the vertex v of T. Without loss of generality we can suppose that A is not invertible. Part 2) of Lemma 4.8 implies that either v is fixed by G or it lies in the minimal subtree of T. Thus it suffices to consider the case when T itself is minimal. Now part 1) of Lemma 4.8 shows that for any edge s of T, the sets Z_s and Z_s^* are nontrivial S–almost invariant subsets of G. Recall that the collection E of all these sets is nested. We now enlarge this set to a set F by adding A and A^* and all their translates. As A is associated to a splitting, its translates are nested. Part 4) of Lemma 4.6 implies that either A is comparable with every Z_s or that A is equivalent to some Z_s. In the first case, we can apply Theorem 2.39 to obtain the required refinement of the graph of groups $G\backslash T$. (This is where we use the assumptions that H is finitely generated and that A is not invertible.) In the second case, the required refinement of $G\backslash T$ can be constructed by simply subdividing the image of the edge s of T such that A is equivalent to Z_s. $\square$

Recall from part 5) of Lemma 4.6 that if A is a nontrivial H–almost invariant subset of a group G which is enclosed by a vertex v of a G–tree T, then for each edge s of T which is directed towards v, we have $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small. If both sets are small, then Z_s^* itself must be small and hence is a trivial S–almost invariant subset of G. Part 1) of Lemma 4.8 shows that this cannot occur if T is minimal. Thus, if T is minimal, we have a naturally defined H–invariant partition of the edges of T, where one set consists of those s with $A \cap Z_s^*$ small and the other set consists of those s with $A^* \cap Z_s^*$ small. This induces a H–invariant partition of all the vertices of $T - \{v\}$ as in the following definition.

Definition 4.12. — Let A be a nontrivial H–almost invariant subset of a group G and let v a vertex of a minimal G–tree T. Suppose that A is enclosed by v, and that s is an edge of T which is directed towards v. We will say that s *lies on the A–side of v* if $A^* \cap Z_s^*$ is small.

We will say that a vertex u of $T - \{v\}$ *lies on the A–side of v* if the path from u to v ends in an edge s which lies on the A–side of v. The collection of all vertices of $T - \{v\}$ which lie on the A–side of v will be denoted by $\Sigma_v(A)$ or by $\Sigma(A)$ if the context is clear.

Remark 4.13. — As usual, this definition does not depend on the choice of basepoint for T.

It is easy to see that if an edge s of T lies on the A–side of v, then the same holds for every edge (and hence every vertex) in the path joining s to v. Also if a vertex u of $T - \{v\}$ lies on the A–side of v, then the path from u to v consists entirely of edges which lie on the A–side of v. Thus the ideas of an edge and a vertex being on the A–side of v are compatible. Clearly, every vertex of $T - \{v\}$ lies in $\Sigma_v(A)$ or $\Sigma_v(A^*)$, so that these two sets partition the vertices of $T - \{v\}$. Also $\Sigma_v(A)$ and $\Sigma_v(A^*)$ are each clearly H–invariant, under the left action of H.

To understand the reason for our terminology, think of A as determined by a closed curve on a surface M which lies inside a subsurface N of M, think of G as being $\pi_1(M)$, and think of T as being the G–tree determined by N, so that the picture in T corresponds to the picture in the universal cover of M.

It is natural to ask if we can replace A by an equivalent almost invariant subset B of G such that for all edges s of T which are directed towards v, we have $B \cap Z_s$ or $B^* \cap Z_s$ is empty. Equivalently, can we replace A by B which is nested with respect to every Z_s? We will show that this is indeed the case.

Suppose that we are given a nontrivial H–almost invariant subset A of G which is enclosed by a vertex v of T. The following result shows how to replace A by a subset $B(A)$ which is nested with respect to every Z_s.

As usual, given a basepoint w for T, we define $\varphi : G \to V(T)$ by the formula $\varphi(g) = gw$. Then we define

$$B(A) = \varphi^{-1}(\Sigma_v(A)) \cup (A \cap \varphi^{-1}(v)),$$
$$C(A) = \varphi^{-1}(\Sigma_v(A^*)) \cup (A^* \cap \varphi^{-1}(v)).$$

Note that these definitions are clearly equivariant, i.e. $B(kA) = kB(A)$, for all k in G.

Lemma 4.14. — *Let G be a finitely generated group with a finitely generated subgroup H, and a nontrivial H–almost invariant subset A. Suppose that T is a minimal G–tree with basepoint w. If A is enclosed by a vertex v of T, then $C(A) = B(A)^*$, and $B(A)$ is H–almost invariant and is equivalent to A. Further, $B(A)$ is nested with respect to Z_s, for every oriented edge s of T.*

Proof. — To simplify notation, we will write $\Sigma(A)$ in place of $\Sigma_v(A)$ throughout this proof.

It is clear from their definitions that B and C are disjoint and that $B \cup C = G$. Thus $C = B^*$. It is also clear that $HB = B$, because $H(\Sigma(A)) = \Sigma(A)$. Finally, it is also clear that if s is any edge of T which is directed towards v, and lies on the A–side of v, then $B(A) \supset \varphi^{-1}(\Sigma(A)) \supset Z_s^*$. If s lies on the A^*–side of v, then $C(A) \supset Z_s^*$. Hence $B(A)$ is nested with respect to Z_s, for every oriented edge s of T.

It remains to show that $B(A)$ is H–almost invariant and is equivalent to A.

As $\varphi(g) = gw$, we see that $\varphi^{-1}(\Sigma(A)) = \{g \in G : gw \in \Sigma(A)\}$. As T is minimal, part 2) of Lemma 4.8 shows that the basepoint w encloses some nontrivial K–almost invariant subset U of G. This implies that $\varphi^{-1}(\Sigma(A)) \subset \{g \in G : gU^{(*)} < A\}$. Now Lemma 2.31 implies that $\varphi^{-1}(\Sigma(A))$ lies in a bounded neighbourhood of A. It follows that B itself is contained in a bounded neighbourhood of A. Similarly B^* is contained in a bounded neighbourhood of A^*. It follows that δB lies in a bounded neighbourhood of δA. As A is H–almost invariant, we know that δA projects to a finite subset of the quotient graph $H\backslash\Gamma$. It follows that δB also projects to a finite subset of $H\backslash\Gamma$, so that B projects to an almost invariant subset of $H\backslash G$ and hence is H–almost invariant. As B is contained in a bounded neighbourhood of A, and B^* is contained in a bounded neighbourhood of A^*, it now follows that B is equivalent to A, which completes the proof of the lemma. $\square$

Before our final result of this chapter, we will need the following simple proposition.

Proposition 4.15. — *Let G be a finitely generated group, and let X be a H–almost invariant subset of G which is contained in a proper subgroup K of G. Then X is H–finite.*

Proof. — As K is a proper subgroup of G, there is an element $g \in G - K$. As X is H–almost invariant, we must have X and Xg being H–almost equal. The assumption that X is contained in K implies that X and Xg are disjoint. It follows that X is H–finite. $\square$

Now we can prove the following useful result.

Corollary 4.16. — *Let G be a finitely generated group with finitely generated subgroups H and K, and let T be a minimal G–tree which is not a single point. Let U be a nontrivial H–almost invariant subset of G and let V be a nontrivial K–almost invariant subset of G. Then the following statements hold:*

(1) *If U is enclosed by a vertex v of T, then both $\Sigma_v(U)$, the U–side of v, and $\Sigma_v(U^*)$, the U^*–side of v, are nonempty, so that U determines a nontrivial partition of the vertices of $T - \{v\}$.*

(2) *If U is enclosed by a vertex v of T, then H is contained in $\mathrm{Stab}(\Sigma_v(U))$ with finite index.*

(3) *If U and V are enclosed by a vertex v of T, and if they determine the same partition of the vertices of $T - \{v\}$, then U and V are equivalent.*

Proof. — We start by fixing the vertex v and choosing it to be the basepoint of T. As usual, this determines the G–equivariant map $\varphi : G \to V(T)$ given by $\varphi(g) = gv$.

1) By applying Lemma 4.14, we can assume that

$$U = B(U) = \varphi^{-1}(\Sigma_v(U)) \cup (U \cap \varphi^{-1}(v)).$$

If $\Sigma_v(U)$ is empty, this implies that $U \subset \varphi^{-1}(v)$. Now $\varphi^{-1}(v) = \mathrm{Stab}(v)$, and the assumption that T is minimal and not a single point implies that $\mathrm{Stab}(v)$ must be a proper subgroup of G. Thus Lemma 4.15 implies that U must be trivial. We conclude that $\Sigma_v(U)$ cannot be empty, and similarly that $\Sigma_v(U^*)$ cannot be empty. Thus U determines a nontrivial partition of the vertices of $T - \{v\}$, as claimed.

2) Clearly $H \subset \mathrm{Stab}(\Sigma_v(U))$. Further, it is clear that $\Sigma_{gv}(gU) = g\Sigma_v(U)$, for any element g of G. Hence if $g \in \mathrm{Stab}(\Sigma_v(U))$, then $\Sigma_{gv}(gU) = \Sigma_v(U)$. Let U_1 be a translate of U enclosed by a vertex of T on the U–side of v, and let U_2 be a translate of U enclosed by a vertex of T on the U^*–side of v. Then, for any $g \in \mathrm{Stab}(\Sigma_v(U))$, the set $g\overline{U}$ lies between $\overline{U_1}$ and $\overline{U_2}$. As the number of translates of $\overline{U}$ between $\overline{U_1}$ and $\overline{U_2}$ is finite, it follows that the orbit of U under the action of $\mathrm{Stab}(\Sigma_v(U))$ is finite, so that H has finite index in $\mathrm{Stab}(\Sigma_v(U))$.

3) Apply Lemma 4.14 so that we can assume that

$$U = B(U) = \varphi^{-1}(\Sigma_v(U)) \cup (U \cap \varphi^{-1}(v)),$$

and

$$V = B(V) = \varphi^{-1}(\Sigma_v(V)) \cup (V \cap \varphi^{-1}(v)).$$

Suppose that $\Sigma_v(U) = \Sigma_v(V)$. It follows that the symmetric difference $U + V$ of U and V is contained in $\varphi^{-1}(v) = \mathrm{Stab}(v)$. As part 2) implies that the stabilisers H_U and H_V of U and V are commensurable, both U and V are H–almost invariant, where $H = H_U \cap H_V$. It follows that $U + V$ is also H–almost invariant. As in part 1), $\mathrm{Stab}(v)$ must be a proper subgroup of G, and now Proposition 4.15 shows that $U + V$ must be trivial, so that U and V are equivalent as claimed. □

CHAPTER 5

ALGEBRAIC REGULAR NEIGHBOURHOODS: ENCLOSING

In this chapter, we will apply the results of the previous chapter to graphs of groups which are obtained by the regular neighbourhood construction. Let G denote a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. In chapter 3, we discussed how to construct the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. If the X_λ's are in good position and isolated X_λ's are not invertible, the construction produces a bipartite G–tree T whose V_0–vertices are the CCC's of $\overline{E}$. Recall that the construction always works when Λ is finite. In order to show that this construction and our ideas about enclosing all fit together, our next result shows that the V_0–vertices of $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ enclose the given X_λ's. Recall from Summary 3.16 the following brief description of our construction. "First we omit some X_i's and replace others by equivalent almost invariant sets such that the new family is in good enough position. Now apply the direct construction using a pretree to the new family to obtain a graph of groups structure Γ for G. Finally, $\Gamma(X_1, \dots, X_n : G)$ is constructed from Γ by subdividing certain edges of Γ." Recall from part 2) of Lemma 4.6 that if a vertex encloses one almost invariant set it also encloses all equivalent almost invariant sets. This fact combined with the above description shows that it suffices to prove that the V_0–vertices of $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ enclose the given X_λ's in the case when the X_λ's are in good position and isolated X_λ's are not invertible. This we now do.

Lemma 5.1. — *Let G denote a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G, and suppose that the X_λ's are in good position, that isolated X_λ's are not invertible, and that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Let T denote the bipartite G–tree produced in that chapter in order to define this regular neighbourhood.*

If v is a V_0–vertex of T, and the corresponding CCC of $\overline{E}$ contains an element $\overline{U}$ of $\overline{E}$, then v encloses U.

Proof. — We choose v as our basepoint for T, so that we have the equivariant map $\varphi : G \to V(T)$ given by $\varphi(g) = gv$. Thus $\varphi(G)$ is contained in $V_0(T)$, which we can identify with the pretree P consisting of all the CCC's of $\overline{E}$. Let s be an edge of T which is oriented towards v. We will show that $U \cap Z_s^*$ or $U^* \cap Z_s^*$ is small. Recall that $Z_s = \varphi^{-1}(Y_s)$, where Y_s denotes the collection of all the vertices of T which lie on the terminal vertex side of s. Thus Y_s includes v. Also $Z_s^* = \varphi^{-1}(Y_s^*) = \{g \in G : gv \in Y_s^*\}$. Hence if g and h lie in Z_s^*, then v does not lie between gv and hv. Recall that the idea of betweenness which we defined on the set P of all CCC's of $\overline{E}$ is the same as the idea of betweenness for the V_0–vertices of T. Thus the CCC $[\overline{U}]$ does not lie between $g[\overline{U}]$ and $h[\overline{U}]$. Our definition of betweenness for P implies that $\overline{U}$ does not lie between $g\overline{U}$ and $h\overline{U}$. As Y_s^* does not include v, we know that g cannot fix v, so that gU and U are in distinct CCC's. Thus they are comparable using our partial order $\leqslant$ on the elements of E. Similarly hU and U are comparable. By replacing U by U^* if necessary, we can arrange that $gU^{(*)} \leqslant U$, and the fact that $\overline{U}$ does not lie between $g\overline{U}$ and $h\overline{U}$ then implies that we must also have $hU^{(*)} \leqslant U$. We conclude that, by replacing U by U^* if necessary, we can arrange that $gU^{(*)} \leqslant U$, for every $g \in Z_s^*$. Now Lemma 2.31 implies that Z_s^* lies in a bounded neighbourhood of U, so that $U^* \cap Z_s^*$ is small. It follows that for any edge s of T which is oriented towards v, we have $U \cap Z_s^*$ or $U^* \cap Z_s^*$ is small, so that U is enclosed by v as required. This completes the proof of the lemma. □

Our first application of this result gives us some new information about our regular neighbourhood construction. The examples at the end of chapter 3 show that this construction can yield a graph of groups Γ consisting of a single point, but they also suggest that this is quite unusual. A more delicate question is whether the graph of groups decomposition of G which we obtain can decompose G trivially when Γ is not a point. The answer is that this cannot happen. In fact, we can show the far stronger result that $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is always minimal.

Proposition 5.2. — *Let G be a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G, and suppose that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Let T denote the universal covering G–tree of this regular neighbourhood. Then T is a minimal G–tree, so that $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is also minimal.*

Proof. — Let T_0 denote the minimal subtree of T. If G fixes more than one vertex of T so that T_0 is not unique, we let T_0 denote one of the vertices fixed by G. We will show that $T_0 = T$. Now Lemma 5.1 tells us that every V_0–vertex of T encloses some translate of some X_λ, and part 2) of Lemma 4.8 shows that any vertex of T which encloses a nontrivial almost invariant subset of G must lie in T_0. Thus T_0 contains every V_0–vertex of T. If the X_λ's are in good position, the construction of T

from a pretree shows that each V_1–vertex of T is joined by edges of T to at least two V_0–vertices. (It is possible that T has a single V_0–vertex, but in this case T consists only of this vertex.) The brief description of our construction given above implies that even if the X_λ's are not in good position, it is still true that each V_1–vertex is joined by edges of T to at least two V_0–vertices. As T is a tree, it follows that $T_0 = T$ as required. $\square$

Next we will apply Lemma 4.14 to the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. Let T denote the universal covering G–tree of $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. As each X_λ is enclosed by a vertex of T, this lemma tells us how to replace each X_λ by an equivalent set $B(X_\lambda)$, which is nested with respect to Z_s, for every edge s of T. Note that before we can define Z_s, we should first choose a basepoint w for T. Recall from Lemma 3.17 that replacing each X_λ by an equivalent set does not alter the regular neighbourhood. Thus we obtain the following interesting result which can be thought of as asserting that we can replace the X_λ's by equivalent sets which are in "very good position".

Lemma 5.3. — *Let G be a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G, and suppose that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Then we can replace each X_λ by an equivalent almost invariant set $B(X_\lambda)$ so that if U and V are any two elements of the set $E = \{gX_\lambda, gX_\lambda^* : g \in G\}$, then either $B(U)$ and $B(V)$ are nested or U and V lie in the same CCC of $\overline{E}$.*

Proof. — If U and V lie in distinct CCC's of $\overline{E}$, they are enclosed by distinct vertices of T. Let s denote an edge of T on the path joining these two vertices. As $B(U)$ and $B(V)$ are nested with respect to every Z_s, it follows that $\overline{Z_s}$ lies between $\overline{B(U)}$ and $\overline{B(V)}$, so that the two sets $B(U)$ and $B(V)$ are nested, as claimed. $\square$

We next examine the special features of isolated elements of E in the case when the elements of E are in good enough position, and isolated X_λ's are not invertible. Recall that if A is an isolated element of E, then the translates of A are nested and since gA is also isolated in E, each gA determines exactly one cross-connected component of $\overline{E}$, which we call isolated. The following result characterises isolated CCC's of $\overline{E}$ in terms of the G–tree T.

Proposition 5.4. — *A V_0–vertex v of T has valence 2 if and only if v corresponds to an isolated CCC of $\overline{E}$.*

For such a vertex v of T, its stabiliser equals the stabilisers of the two incident edges.

Proof. — Suppose that v is a V_0–vertex of valence 2 and let s and t denote the edges which are incident to v. Let U and V denote elements of E enclosed by v. Part 1) of Corollary 4.16 implies that s and t must lie one on the U–side of v and the other on the U^*–side of v, and the same holds with U replaced by V. Thus U and V determine the same partition of the vertices of $T - \{v\}$. Now part 3) of Corollary 4.16 shows that U and V must be equivalent, up to complementation. As required, it now follows that the CCC which corresponds to v contains exactly one element of $\overline{E}$.

To prove the converse, suppose that A is an isolated element of $\overline{E}$. Consider the V_0–vertex v which is the CCC of A and, in the pretree P, let S be a star which contains v. Thus S can also be identified with a V_1–vertex v_1 in T which is adjacent to v. Let $v_2 \neq v$ be a V_0–vertex corresponding to another CCC in S. This means that both v_2 and v are adjacent to v_1 in the tree T. If U is an element of E which is enclosed by v_2, then we must have $U^{(*)} > A$ or $U^{(*)} < A$. If $U^{(*)} < A$, then we have the same inequality for any other CCC in S other than v. For otherwise, v would lie between two of these vertices. Conversely, if $[\overline{U}] = u$ is any vertex of a star which contains v and if $U^{(*)} < A$, then u must lie in S since v cannot lie between u and v_2. Thus, there are at most two stars which contain v. Equivalently, there are at most two edges of T incident to v. Now the fact that T is minimal implies that it has no vertices of valence 1. It follows that v has valence 2 as required.

For the last part of the proposition, we note that the stabiliser of v equals the stabiliser of $\overline{U}$, where $\overline{U}$ is the unique element of $\overline{E}$ which corresponds to v. As we are assuming that U is not invertible, this stabiliser equals the stabiliser of U. Now the stabiliser of U must stabilise each of s and t, so it follows that the stabiliser of v equals the stabilisers of each of s and t as required. □

Before proceeding, we will need the following simple fact about splittings of groups.

Lemma 5.5. — *Let G be a group with a splitting σ over a subgroup H. Let X be one of the standard H–almost invariant subsets of G associated to σ. Then one of the following holds:*

(1) *there is an element g of G such that $gX \subset X^*$, and an element k of G such that $kX^* \subset X$.*

(2) *σ is a HNN extension $G = H*_H$ in which at least one of the inclusions of the edge group in the vertex group is an isomorphism.*

Remark 5.6. — In case 2), such a HNN extension is often called ascending.

Proof. — The given splitting σ determines a G–tree T such that the quotient $G\backslash T$ has a single edge. We pick an orientation of this edge, thereby fixing a G–invariant orientation on all the edges of T. In terms of our previous notation, there is an edge s of T such that X or X^* is equivalent to Z_s. We will assume that X is equivalent to Z_s. Let v denote the vertex of T at the initial end of s, and let gs denote another edge incident to v. Then $gX \subset X^*$ or $gX^* \subset X^*$, depending on whether gs points away from v or towards v. Suppose that there is no element g of G such that $gX \subset X^*$. Then $gX^* \subset X^*$, and gs must point towards v. It follows that s is the only edge of T

which is incident to v and points away from v. Hence the stabilisers of s and of v are equal, which implies the conclusion of case 2) of the lemma.

A similar argument considering the terminal vertex w of s shows that if there is no element k of G such that $kX^* \subset X$, then the stabilisers of s and of w are equal, which again implies the conclusion of case 2) of the lemma. □

Our next result is crucial for understanding algebraic regular neighbourhoods. It is the algebraic analogue of the topological fact that if N is a regular neighbourhood of a finite collection C_λ of closed curves on a surface M, and if C is a closed curve which has zero intersection number with each C_λ, then we can homotop C into $M - N$. In our result, the curve C is replaced by a H–almost invariant subset X of G, and the conclusion is that if X crosses no element of E, then X is enclosed by a V_1–vertex of $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. As usual, we argue with the G–tree T rather than with the graph of groups Γ itself. As for Lemma 5.1, it suffices to prove this for the case when the X_λ's are in good position.

Proposition 5.7. — *Let G be a finitely generated group with a family of finitely generated subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G, and suppose that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Further suppose that the X_λ's are in good position, and that there is more than one CCC, so that the pretree P is not a single point. Let T denote the G–tree with $V_0(T) = P$. Let X be a nontrivial H–almost invariant subset of G which does not cross any element of E. Then the following statements hold.*

(1) *If H is finitely generated, then X is enclosed by a V_1–vertex of T.*

(2) *If X is a standard H–almost invariant set associated to a splitting of G over H, which need not be finitely generated, then X is enclosed by a V_1–vertex of T.*

Remark 5.8. — For this proof, it does not matter whether any isolated X_λ's are invertible.

Proof. — We start by showing that we can assume that X is in good position, i.e. that the set of all translates of X satisfies Condition (*). If H is not finitely generated, we are in case 2), and this is true by hypothesis. If H is finitely generated, we first observe that Condition (*) is automatic for X and gX, unless gX is equivalent to X or X^*. Now we apply Proposition B.2.14, which asserts that X is equivalent to an almost invariant set Y whose translates are nested with respect to the subgroup $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$. This immediately implies that Y satisfies Condition (*), and so it suffices to replace X by Y.

If X is equivalent to Z_s for any edge s of T, the result follows from part 3) of Lemma 4.6. If X is equivalent to an element U of E, our hypothesis on X implies that U must be isolated in E. This implies that X is equivalent to Z_s for each of the two edges of T incident to the V_0–vertex which corresponds to the CCC $[\overline{U}]$. In particular, the result follows in this case also. So we will assume that X is not equivalent to any element of E. As the X_λ's are in good position and X is in good position, this implies that the relation $\leqslant$ on E can be extended to the set obtained from E by adding in all translates of X and X^*. Note that $\leqslant$ is a partial order on this larger set even if the stabiliser H of X is not finitely generated, because the proof of Lemma B.1.14 still applies.

Our first step is to show that X is sandwiched between two elements of E, i.e. that there are elements U_1 and U_2 of E such that $U_1 < X < U_2$. Let Y denote an element of E. As X crosses no element of E, we know that, for each element g of G, one of the four inequalities $gY \leqslant X$, $gY^* \leqslant X$, $gY \leqslant X^*$, $gY^* \leqslant X^*$ must hold. If the stabiliser H of X is finitely generated, then Lemma 2.31 tells us that $\{g \in G : gY^{(*)} \leqslant X\}$ is contained in a bounded neighbourhood of X, and that $\{g \in G : gY^{(*)} \leqslant X^*\}$ is contained in a bounded neighbourhood of X^*. As G is the union of these two sets it follows that neither is empty so that there are elements U_1 and U_2 of E such that $U_1 < X < U_2$, as required.

If H is not finitely generated, we use the hypothesis that X is associated to a splitting σ of G. Suppose that X is not sandwiched between two elements of E. Then, by replacing X by X^*, if necessary, we have that for every element U of E, either $U < X$ or $U^* < X$. As E is G–invariant, it follows that for every element U of E, and for every element g of G, either $U < gX$ or $U^* < gX$. Now suppose that there is an element g of G such that $gX \subset X^*$. This implies that for every element U of E, either $U < X^*$ or $U^* < X^*$, which contradicts the fact that either $U < X$ or $U^* < X$. It remains to handle the situation where there is no element g of G such that $gX \subset X^*$. Then Lemma 5.5 shows that σ is an ascending HNN extension $G = H*_H$. Let $f : G \to \mathbb{Z}$ denote the natural homomorphism associated to this HNN extension, and let U be an element of E such that $U < X$. As U does not cross X, it follows, in particular, that the image in $\mathbb{Z}$ of the coboundary δU of U is bounded above or below. Hence the stabiliser of U must be contained in $\ker(f)$. It follows that the image of δU in $\mathbb{Z}$ must be finite. As the image in $\mathbb{Z}$ of δgX is finite for all g in G, there must be an element g of G such that $gX \subset U$. But we know that $U < gX$ or $U^* < gX$. The first implies that U is equivalent to gX, which contradicts our assumption that X is not equivalent to any element of E, and the second implies that $U^* < U$ which is also a contradiction. This contradiction completes the proof that, in all cases, X must be sandwiched between two elements U_1 and U_2 of E.

By considering the path in T which joins the V_0–vertices $[\overline{U_1}]$ and $[\overline{U_2}]$, it is easy to see that there is a V_1–vertex v with two V_0–vertices v_1 and v_2 adjacent to v, and an element V_i of E enclosed by v_i, such that $V_1 < X < V_2$. We will show that X is enclosed by v.

If U is any element of E, then either $U^{(*)} < X$ or $U^{(*)} < X^*$ but not both. Further if V lies in the same CCC as U, then the same inequality must hold as for U. Thus we obtain a partition of the V_0–vertices of $T - \{v\}$ into two subsets Φ and Φ^*, where the vertices of Φ enclose those elements U of E such that $U^{(*)} < X$ and the vertices of Φ^* enclose those elements U of E such that $U^{(*)} < X^*$. If w lies in Φ, then every V_0–vertex on the path from w to v also lies in Φ. This enables us to define a partition of all the vertices of $T - \{v\}$ into two subsets Ψ and Ψ^*. We will say that a vertex w of $T - \{v\}$ lies in Ψ if the last V_0–vertex on the path from w to v lies in Φ. Thus Φ is contained in Ψ. Note that if we already knew that X was enclosed by v, then Ψ and Ψ^* would constitute the X–side of v and the X^*–side of v. If $U < X$ and $h \in H$, then $hU < X$. Thus Ψ and Ψ^* are H–invariant. In particular, H must be a subgroup of $\mathrm{Stab}(v)$.

Now, as in the proof of Lemma 4.14, we choose v as our basepoint for T, define $\varphi : G \to V(T)$ by the formula $\varphi(g) = gv$, and define $B(X) = \varphi^{-1}(\Psi) \cup (X \cap \varphi^{-1}(v))$, and $C(X) = \varphi^{-1}(\Psi^*) \cup (X^* \cap \varphi^{-1}(v))$. Clearly $B(X)$ and $C(X)$ partition G, so that $C(X) = B(X)^*$. If s is an edge of T which is directed towards v, then $Z_s^* \subset B(X)$ or $Z_s^* \subset C(X)$. It is also clear that $HB = B$, because Ψ, X and v are all H–invariant.

We will show that $B(X)$ is H–almost invariant and is equivalent to X. It will then follow that $Z_s^* \leqslant X$ or $Z_s^* \leqslant X^*$, for every edge s of T which is directed towards v, so that X is enclosed by v.

If H is finitely generated, we will proceed very much as in the proof of Lemma 4.14. As $\varphi(g) = gv$, we have $\varphi^{-1}(\Psi) = \{g \in G : gv \in \Psi\}$. As $V_1 < X < V_2$, we have $gV_1 < gX < gV_2$, for all $g \in G$. If gv lies in Ψ, then gv_1 and gv_2 must also lie in Ψ. Thus we have the inequalities $gV_i^{(*)} < X$, for $i = 1, 2$. It follows that one of the inequalities $gX^{(*)} < X$ also holds. Thus $\varphi^{-1}(\Psi) = \{g \in G : gv \in \Psi\}$ is contained in $\{g \in G : gX^{(*)} < X\}$, which is contained in a bounded neighbourhood of X, by Lemma 2.31. It follows that $B(X)$ itself is contained in a bounded neighbourhood of X. Similarly $C(X) = B(X)^*$ is contained in a bounded neighbourhood of X^*. Exactly as in the proof of Lemma 4.14 it follows that $B(X)$ is H–almost invariant and is equivalent to X.

If H is not finitely generated, we will choose X carefully and then show that $X = B(X)$. As discussed above, this implies that X is enclosed by v as required. Recall our assumption that X is associated to a splitting of G over H. As discussed after the proof of Lemma 2.10, we can choose X to satisfy $X = \{g \in G : gX^{(*)} \subset X\}$. As the translates of X are nested, we have

$$\{g \in G : gX^{(*)} \subset X\} = \{g \in G : gX^{(*)} \leqslant X\}$$

which allows us to express X as the disjoint union

$$\{g \in G : g \notin \mathrm{Stab}(v),\ gX^{(*)} < X\} \cup (X \cap \mathrm{Stab}(v)).$$

Recall that $B(X) = \varphi^{-1}(\Psi) \cup (X \cap \varphi^{-1}(v))$. As $\varphi(e) = v$, we have $\varphi^{-1}(v) = \mathrm{Stab}(v)$, so that the second terms in our expressions for X and $B(X)$ are equal.

Now we consider the first terms. As above, we have

$$\varphi^{-1}(\Psi) = \{g \in G : gv \in \Psi\} \subset \{g \in G : g \notin \mathrm{Stab}(v), gX^{(*)} < X\}.$$

Suppose that $g \notin \mathrm{Stab}(v)$ and $gX < X$. As $g \notin \mathrm{Stab}(v)$, the vertices gv, gv_1 and gv_2 of T must all lie in Ψ or all in Ψ^*. As $gV_1 < gX < X$, we see that gv_1 must lie in Ψ. Thus $gv \in \Psi$, so that $g \in \varphi^{-1}(\Psi)$. A similar argument applies if $g \notin \mathrm{Stab}(v)$ and $gX^* < X$. Thus $\{g \in G : g \notin \mathrm{Stab}(v), gX^{(*)} < X\} \subset \varphi^{-1}(\Psi)$. We conclude that $\varphi^{-1}(\Psi) = \{g \in G : g \notin \mathrm{Stab}(v), gX^{(*)} < X\}$, so that $X = B(X)$ as claimed. □

Recall from the definition of betweenness on the pretree of CCC's of $\overline{E}$ that if three V_0–vertices v_1, v_2 and v_3 lie on a path in T with v_2 between v_1 and v_3, then there is an element X of E which is enclosed by v_2 and such that for any elements Y and Z of E with Y enclosed by v_1 and Z enclosed by v_3, we have $\overline{Y}\,\overline{X}\,\overline{Z}$. Hence if a V_0–vertex v_2 lies between edges s and t which point towards v_2, there is an element X of E enclosed by v_2 such that $Z_s^* \leqslant X \leqslant Z_t$. The following result is an immediate consequence and is the analogous result for V_1–vertices.

Proposition 5.9. — *Let Y be a nontrivial H–almost invariant subset of G, and let Z be a nontrivial K–almost invariant subset of G. If H is not finitely generated, suppose in addition that Y is a standard H–almost invariant set associated to a splitting of G over H, and similarly for Z. If Y and Z are enclosed by distinct V_1–vertices v_1 and v_3, then there is a V_0–vertex v_2 and an element X of E which is enclosed by v_2 such that $Y^{(*)} \leqslant X \leqslant Z^{(*)}$.*

Proof. — Let v_2 be any V_0–vertex on the path joining v_1 and v_3. Let s and t be the edges of this path which are incident to v_2 and point towards v_2, labelled so that s is nearer to v_1 than is t. As we pointed out above, there is an element X of E enclosed by v_2 such that $Z_s^* \leqslant X \leqslant Z_t$. As s points away from v_1, we have $Y^{(*)} \leqslant Z_s^*$, and as t points away from v_3, we have $Z^{(*)} \geqslant Z_t$. It follows that $Y^{(*)} \leqslant Z_s^* \leqslant X \leqslant Z_t \leqslant Z^{(*)}$, so that $Y^{(*)} \leqslant X \leqslant Z^{(*)}$, as required. □

Next we apply the preceding results on regular neighbourhoods to obtain new results about more general graphs of groups.

First we can use the argument of Proposition 5.7 to show that the two ideas of enclosing a splitting of G over a subgroup H are equivalent even when H is not finitely generated.

Lemma 5.10. — *Suppose that A is associated to a splitting σ of G over H, and T is a G–tree. Let Γ denote the graph of groups structure for G given by the quotient $G \backslash T$, and let u denote the image of v in Γ. Then A is enclosed by u if and only if σ is enclosed by u.*

Proof. — In Lemma 4.10, we showed that if σ is enclosed by u then A is enclosed by u. We also proved the converse in the case when H is finitely generated. It remains to prove the converse in the case when H is not finitely generated.

Suppose that A is enclosed by the vertex u of Γ. This means that A is enclosed by some translate in T of v, and we will assume that A is enclosed by v itself. As always there is a special case if v is fixed by G, but the result is trivial in this case as v encloses all splittings and all nontrivial almost invariant subsets of G. So we will now assume that v is not fixed by G. Part 2) of Lemma 4.8 shows that v lies in the minimal subtree of T. By replacing T by its minimal subtree, we can assume that T itself is minimal.

Now the argument at the end of the proof of Proposition 5.7 applies to show that, by replacing A by an equivalent set, we can arrange that $A = B(A)$. As $B(A)$ is clearly nested with respect to every Z_s, the set E of all Z_s, for all oriented edges of T, together with all the translates of A and A^*, forms a nested set F of subsets of G. Thus F is a G–set partially ordered by inclusion. We want to apply Dunwoody's construction in [**10**] to F to obtain a G–tree T' whose oriented edges naturally correspond to the elements of F. This will yield a graph of groups structure $\Gamma' = G\backslash T'$ for G which is a refinement of Γ, with one extra edge whose associated edge splitting is σ. This will show that σ is enclosed by u as required. In order for his construction to be applicable, we need to know that F is discrete. Recall that the elements of E correspond to the edges of T and have the corresponding partial order. It follows at once that E is discrete. Also the fact that A is associated to a splitting implies that the set of all translates of A is discrete. (In fact, the translates of A correspond to the edges of the G–tree determined by the splitting.) As T is minimal, each vertex has valence at least two, so that A, and hence any translate of A, is sandwiched between two elements of E. Also any element of E lies between two translates of A or A^*. Combining these facts shows that F is discrete, as required. We should point out that there is a possible problem here, because A might be invertible. But we can simply replace it by an equivalent non-invertible almost invariant set. The new set F might not be nested, but it will still be in good position, which is enough to be able to apply Dunwoody's construction. □

Next we introduce a natural generalisation of the idea of a splitting of a group being enclosed by a vertex of a graph of groups. Recall that the precise definition of this idea was motivated by considering a surface with a subsurface which contains a simple closed curve. We now generalise this to the case of a subsurface which contains another subsurface. Let Γ and Γ' denote graphs of groups structures for a given group G. We will define what it means for a vertex of Γ' to enclose a vertex of Γ. If v is a vertex of Γ, we let Γ_v denote the graph of groups structure for G obtained from Γ by collapsing the closure of $\Gamma - \mathrm{star}(v)$. Recall that this means that each component of the closure of $\Gamma - \mathrm{star}(v)$ is collapsed separately to a point.

Definition 5.11. — Let Γ and Γ' denote minimal graphs of groups structures for a given group G, and let v and v' denote vertices of Γ and Γ' respectively. Then v is *equivalent* to v' if there is a graphs of groups isomorphism from Γ_v to $\Gamma'_{v'}$ which sends v to v'.

Definition 5.12. — Let Γ and Γ' denote minimal graphs of groups structures for a given group G, and let v and v' denote vertices of Γ and Γ' respectively. Then v is *enclosed* by v' if there is a graph of groups structure Γ'' for G and a projection $p : \Gamma'' \to \Gamma'$ and a vertex v'' of Γ'' such that $p(\text{star}(v'')) = v'$, and v'' is equivalent to v.

Remark 5.13. — If v is enclosed by v' and s is an edge of Γ incident to v, there is a corresponding edge s'' of Γ'' incident to v'' such that the edge splitting associated to s'' is the same as that associated to s. In particular, this splitting is enclosed by v'.

In the previous chapter, we also defined enclosing using almost invariant sets. In Lemma 5.10, we completed the proof that these two ideas of enclosing a splitting of G over a subgroup H are equivalent, even if H is not finitely generated. We will carry out the analogous arguments here.

Definition 5.14. — Let T and T' be minimal G–trees on which G acts without inversions. Then a vertex w of T is *enclosed* by a vertex w' of T' if $\text{Stab}(w) \subset \text{Stab}(w')$ and, for each edge s of $\text{star}(w)$ with stabiliser S, the S–almost invariant subset Z_s is enclosed by w'.

Note that in this definition, the group S may not be finitely generated. The following result brings these two ideas together as in Lemma 5.10.

Lemma 5.15. — *Let T and T' be minimal G–trees on which G acts without inversions. Let Γ and Γ' denote the quotients $G\backslash T$ and $G\backslash T'$, and let v and v' denote vertices of Γ and Γ' respectively. Then v is enclosed by v' if and only if there are vertices w and w' of T and T' which project to v and v' respectively such that w is enclosed by w'.*

Proof. — Suppose that v is enclosed by v'. Thus there is a graph of groups structure Γ'' for G and a projection $p : \Gamma'' \to \Gamma'$ and a vertex v'' of Γ'' such that $p(\text{star}(v'')) = v'$ and v'' is equivalent to v. Let T, T' and T'' denote the universal covering G–trees of Γ, Γ' and Γ'' respectively, let w denote a vertex of T which lies above v, let w'' denote the corresponding vertex of T'' and let w' denote the image of w'' in T'. As $p(\text{star}(v'')) = v'$, it follows that the image of $\text{star}(w'')$ equals w'. In particular, it follows immediately that w is enclosed by w'.

Conversely suppose that there are vertices w and w' of T and T' which project to v and v' respectively such that w is enclosed by w'. Thus for each edge s of $\text{star}(w)$ with stabiliser S, the S–almost invariant subset Z_s is enclosed by w'. We will choose w as the basepoint of T, choose s to be oriented towards w, and let X denote Z_s. As discussed after the proof of Lemma 2.10, it follows that $X = \{g \in G : gX^{(*)} \subset X\}$, where we use $X^{(*)}$ to denote a set which might be X or X^*. Now we consider the tree T', and the construction of the set $B(X)$ using this tree. The argument at the end of the proof of Proposition 5.7 shows that $X = B(X)$. Let F denote the family of all the translates of Z_s and Z_s^*, for each edge s of star(w). As the Z_s's are associated to distinct edges of T, it follows that F is nested. Now let E'' denote the union of F with the set E' of all the Z_t and Z_t^*, for all oriented edges t of T'. As $B(Z_s)$ is nested with respect to every Z_t, for every edge t of T', it follows that the G–set E'' is nested.

In the rest of our argument we will assume that, if s_1 and s_2 are distinct edges of star(w), then the associated almost invariant sets Z_{s_1} and $Z_{s_2}^*$ are not equivalent. We will also assume that, if s is an edge of star(w) and t is an edge of T', then Z_s is not equivalent to Z_t. If either of these conditions fails, it is not difficult to modify the argument.

We want to apply Dunwoody's construction in [**10**] to E'' to obtain a G–tree T'' whose oriented edges naturally correspond to the elements of E''. This will yield a graph of groups structure $\Gamma'' = G\backslash T''$ for G. In order for his construction to be applicable, we need to know that E'' is discrete. Recall that the elements of E' correspond to the edges of T' and have the corresponding partial order. It follows at once that E' is discrete. Also the fact that the Z_s's are associated to edge splittings of T implies that F is discrete. As T' is minimal, each vertex has valence at least two, so that each Z_s, and hence any translate of Z_s, lies between two elements of E'. Also any element of E' lies between two translates of Z_s or Z_s^*. Combining these facts shows that E'' is discrete, as required.

In order to show that v is enclosed by v', we need to show that there is a vertex v'' of Γ'' which projects to v', such that v is equivalent to v''. Let a and b denote distinct oriented edges of star(w) oriented towards w, and let α and β denote the edges of T'' which correspond to Z_a and Z_b respectively, oriented so that Z_α is equivalent to Z_a and Z_β is equivalent to Z_b. As $Z_a^* \subset Z_b$, it follows that $Z_\alpha^* \leqslant Z_\beta$. As Z_a and Z_b are each enclosed by the vertex w' of T', no element of E' can lie between them. It follows that α and β are adjacent and are oriented towards their common vertex. As this argument applies to all such edges, we see that the edges of T'' which correspond to the Z_s's have a common vertex, which we denote by w'', and they are all oriented towards w''. Next we claim that no edge of T'' which corresponds to an element of E' can meet w''. For suppose there is an edge e of T'' which meets w'', is oriented towards w'' and corresponds to an element of E'. Then $Z_e^* \subset Z_\alpha$ for each edge α of T'' corresponding to an edge a of star(w). Hence $Z_e^* \leqslant Z_a$, for each edge a of star(w). Thus Z_e is enclosed by the vertex w of T. Now recall that an edge s of star(w) lies on the Z_e–side of w if $Z_e^* \cap Z_s^*$ is small. (See Definition 4.12.) Thus every edge of star(w) lies on the Z_e–side of w. But part 1) of Lemma 4.16 tells us that this is impossible. This contradiction completes the proof of the claim. We have shown that the edges of T'' which correspond to the edges of star(w) form star(w''). Let v'' denote the image of w'' in Γ''.

Now we will show that v and v'' are equivalent. Recall that T'' is constructed from the partially ordered set E'' using Dunwoody's construction. As the graph of groups

$\Gamma''_{v''}$ is constructed from Γ'' by collapsing every edge not incident to v'', it follows that the universal covering G–tree $T''_{v''}$ of $\Gamma''_{v''}$ is constructed from T'' by collapsing every edge not incident to some translate of w''. Hence $T''_{v''}$ can be constructed from the partially ordered set F using Dunwoody's construction. One can also think of T as constructed from the partially ordered set E of all Z_s and Z_s^* for all edges s of T. Now the universal covering G–tree T_v of Γ_v is constructed from T by collapsing every edge not incident to some translate of w, and it follows that T_v can also be constructed from the partially ordered set F using Dunwoody's construction. Hence T_v and $T''_{v''}$ are G–equivariantly isomorphic, so that the graphs of groups Γ_v and $\Gamma''_{v''}$ are also isomorphic, as required. □

The preceding results allow us to give a surprising generalisation of Theorem 2.35, which we proved in [**44**]. In that theorem, we showed that if G is a finitely generated group with n splittings over finitely generated subgroups, then the splittings are compatible if and only if each pair of splittings has intersection number zero. Now we will show that the hypothesis that the splittings be over finitely generated subgroups of G can be removed. The precise result we obtain is the following.

Theorem 5.16. — *Let G be a finitely generated group with n splittings over possibly infinitely generated subgroups. Then the splittings are compatible if and only if each pair of splittings has intersection number zero. Further, in this situation, the graph of groups structure on G obtained from these splittings is unique up to isomorphism.*

Proof. — For $1 \leqslant i \leqslant n$, let σ_i be a splitting of G over a subgroup H_i, and let X_i be an associated H_i–almost invariant subset of G, chosen so that X_i is not invertible.

We start by discussing the existence proof. In order to use the previous arguments directly, we will proceed by induction on n. The result is trivial when $n = 1$. Thus we consider the situation where we have k compatible splittings $\sigma_1, \ldots, \sigma_k$ and then another splitting σ_{k+1} which has intersection number zero with each of $\sigma_1, \ldots, \sigma_k$. Thus G has a graph of groups decomposition with k edges and the edge splittings are conjugate to the σ_i's, $1 \leqslant i \leqslant k$. By subdividing each edge into two, we obtain the regular neighbourhood $\Gamma(X_1, \ldots, X_k : G)$. Now part 2) of Proposition 5.7 shows that X_{k+1} is enclosed by a V_1–vertex of $\Gamma(X_1, \ldots, X_k : G)$, and then Lemma 5.10 implies that $\sigma_1, \ldots, \sigma_{k+1}$ are compatible as required.

In order to prove the uniqueness, we want to use the proof of the second part of Theorem B.1.12. This argument never directly uses the hypothesis that the H_i's are finitely generated. It does quote one result, Lemma B.2.3, and this in turn uses Lemma B.2.2. Both lemmas are about splittings of a finitely generated group G, and they each contain the hypothesis that the splittings be over finitely generated subgroups. However, this hypothesis is not needed and the proofs work perfectly well even when the splittings are over infinitely generated subgroups of G. Thus the proof of the second part of Theorem B.1.12 yields the required result. □

The following result is an immediate consequence of the uniqueness part of Theorem 5.16.

Theorem 5.17. — *Let Γ_1 and Γ_2 be minimal graphs of groups structures for a finitely generated group G. If each Γ_i has no redundant vertices, and if they have the same conjugacy classes of edge splittings, then Γ_1 and Γ_2 are isomorphic.*

Remark 5.18. — As G is finitely generated, and Γ_1 and Γ_2 are minimal, it follows that Γ_1 and Γ_2 are finite.

Note that this result needs no assumptions on the edge groups involved.

The following result will be useful when we consider taking regular neighbourhoods of increasing finite families of almost invariant subsets of a fixed group.

Lemma 5.19. — *Let G be a finitely generated group with a finite family of finitely generated subgroups $H_1, \ldots, H_n$. For $1 \leqslant i \leqslant n$, let X_i denote a nontrivial H_i–almost invariant subset of G, such that the X_i's are in good position. Let E_n denote the set of all translates of the X_i's and their complements, let P_n denote the pretree of all CCC's of $\overline{E_n}$, let T_n denote the associated G–tree and let Γ_n denote the corresponding graph of groups structure for G, so that $\Gamma_n = G\backslash T_n$.*

Let $m < n$, and let f denote the natural map from P_m to P_n. Let A be a nontrivial H–almost invariant subset of G enclosed by a V_0–vertex v of T_m. Then A is enclosed by the V_0–vertex $f(v)$ of T_n.

Remark 5.20. — The point of this result is that A need not be an element of E_n.

Proof. — Lemma 4.14 implies that each edge splitting of Γ_n has zero intersection number with each X_i, $1 \leqslant i \leqslant n$, and it is trivial that the edge splittings of Γ_n have zero intersection number with each other. Now Theorem 5.16 implies that there is a common refinement $\Gamma_{m,n}$ of Γ_n and Γ_m obtained by splitting the V_1–vertices of Γ_m using the edge splittings of Γ_n. (Thus the number of edges of $\Gamma_{m,n}$ is the sum of the number of edges of Γ_n and of Γ_m.) Associated to the construction of $\Gamma_{m,n}$ there is a natural quotient map $p_m : T_{m,n} \to T_m$. There is also a natural quotient map $p_n : T_{m,n} \to T_n$ obtained by collapsing those edges of $T_{m,n}$ which correspond to edges of T_m. We choose a basepoint w for $T_{m,n}$, and let $p_m(w)$ and $p_n(w)$ be basepoints for T_m and T_n respectively. Now we consider the V_0–vertex v of T_m. The construction of $T_{m,n}$ means that the pre-image of v is a single vertex u. Further the pre-image of a small neighbourhood N of v maps homeomorphically to N.

Let A be a nontrivial H–almost invariant subset of G enclosed by the vertex v of T_m. Thus for any edge s of $T_{m,n}$ incident to u and oriented towards u, if t denotes $p_m(s)$, then we have $A \cap Z_t^*$ or $A^* \cap Z_t^*$ is small. As the pre-image in $T_{m,n}$ of a small neighbourhood N of v maps homeomorphically to N, our choice of basepoints implies that $Z_s = Z_t$, so that we have $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small, i.e. A is enclosed by the vertex u of $T_{m,n}$. Hence part 5) of Lemma 4.6 shows that $A \cap Z_s^*$ or $A^* \cap Z_s^*$ is small for every edge s of $T_{m,n}$ which is oriented towards u, whether or not s is incident to u. Let r denote an edge of T_n which is incident to $p_n(u)$ and oriented towards $p_n(u)$. The construction of p_n implies that there is a unique edge s of $T_{m,n}$, which need not be incident to u, such that $p_n(s) = r$. The edge s is automatically oriented towards u. Our choice of basepoints implies that $Z_r = Z_s$, so that we have $A \cap Z_r^*$ or $A^* \cap Z_r^*$ is small, as required. Hence A is enclosed by the vertex $p_n(u)$ of T_n. If $p_n(u) = f(v)$, we have completed the proof of the lemma.

Now suppose that $p_n(u) \neq f(v)$, and let U denote an element of E_m which belongs to the CCC of $\overline{E_m}$ corresponding to the V_0–vertex v of T_m. Lemma 4.10 tells us that U is enclosed by v, so that the above argument shows that U must also be enclosed by the vertex $p_n(u)$ of T_n. But we know that U is enclosed by the V_0–vertex $f(v)$ of T_n, as it belongs to the CCC of $\overline{E_n}$ associated to $f(v)$. As these vertices of T_n are distinct, part 1) of Lemma 4.9 tells us that U is equivalent to Z_s or to Z_s^* for each edge s on the path joining them. As each element of E_n is enclosed by some vertex of T_n, it follows that no element of E_n crosses any Z_s, so that U must be an isolated element of E_n. Hence U is an isolated element of E_m, and so Proposition 5.4 tells us that v has valence 2 in T_m. It follows from Corollary 4.16 that A is equivalent to U or U^*, so that again A is enclosed by the vertex $f(v)$ of T_n, as required. The lemma follows. □

We end this chapter by discussing how to generalise our construction of regular neighbourhoods of almost invariant subsets to the case of almost invariant subsets over infinitely generated subgroups. We will use the preceding results to show that this can be done provided that such sets are associated to splittings. While this may seem a little exotic, it is a natural extension once one realises that the edge groups in a regular neighbourhood may be infinitely generated even if the given almost invariant sets are all over finitely presented groups. (See Example 6.11.) Also very little more work is needed. Here is an outline of the theory.

As always we start with a finitely generated group G, a family $\{H_\lambda\}_{\lambda \in \Lambda}$ of subgroups of G, and for each $\lambda \in \Lambda$, a nontrivial H_λ–almost invariant subset X_λ of G. We no longer assume that every H_λ is finitely generated. Instead we assume that if H_λ is not finitely generated, then X_λ is associated to a splitting of G over H_λ. Now consider the construction in chapter 3. We will proceed exactly as we did there, and note the differences.

First we let E denote the collection of all translates of the X_λ's and their complements, and assume that the X_λ's are in good position, i.e. that E satisfies Condition (*), and that isolated X_λ's are not invertible. This allows us to define the partial order $\leqslant$ on E exactly as before. For the proof that $\leqslant$ is a partial order given in Lemma B.1.14 does not use the finite generation of the groups involved. As in our discussion at the end of chapter 3, E may not be discrete. However, we can still define the set $\overline{E}$ of pairs $\{X, X^*\}$ for $X \in E$, and can define P to be the collection of all CCC's of $\overline{E}$.

Further the arguments of chapter 3 still apply, and show that the idea of betweenness can be defined on P as before and this makes P into a pretree. Of course, P need not be discrete, but if it is, then as before P can be embedded in a G–tree T and $G\backslash T$ is a graph of groups structure for G which we denote by $\Theta(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. If no isolated X_λ is invertible, then $\Theta(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is a regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ of the X_λ's. Of course, the V_0–vertex groups of this graph need not be finitely generated, but Proposition 5.22 below shows that $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is minimal, and so must be a finite graph. In order to prove this we first need the following generalisation of Lemma 5.1.

Lemma 5.21. — *Let G denote a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. If H_λ is not finitely generated, then we assume that X_λ is associated to a splitting of G over H_λ. Suppose that the X_λ's are in good position and that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Let T denote the bipartite G–tree produced in that chapter in order to define this regular neighbourhood.*

If v is a V_0–vertex of T, and the corresponding CCC of $\overline{E}$ contains an element $\overline{U}$ of $\overline{E}$, then v encloses U.

Proof. — Lemma 5.1 shows that this result holds if $\mathrm{Stab}(U)$ is finitely generated, so it suffices to consider the case when $\mathrm{Stab}(U)$ is not finitely generated. Let s be an edge of T which is incident to v and oriented towards v, and consider the proof of Lemma 5.1. This proof used the finite generation of the H_i's only at the end when it used Lemma 2.31. Thus the proof of Lemma 5.1 shows that, by replacing U by U^* if needed, we can arrange that $gU^{(*)} \leqslant U$, for every $g \in Z_s^*$. Our assumption that U is associated to a splitting of G now implies that $gU^{(*)} \subset U$, for every $g \in Z_s^*$. Further our discussion just before Definition 2.17 shows that we can choose U so that $U = \{g \in G : gU^{(*)} \subset U\}$. It follows that $Z_s^* \subset U$. As the same argument applies to every such edge s, it follows that U is enclosed by v as required. □

Now we can prove the following generalisation of Proposition 5.2.

Proposition 5.22. — *Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. If H_λ is not finitely generated, then we assume that X_λ is associated to a splitting of G over H_λ. Suppose also that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as in chapter 3. Let T denote the universal covering G–tree of this regular neighbourhood. Then T is a minimal G–tree, so that $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ is also minimal.*

Proof. — The proof of Proposition 5.2 uses only Lemmas 5.1 and 4.8. The second lemma applies even if the groups involved are not finitely generated, and Lemma 5.21

shows that Lemma 5.1 remains true in this more general context. This completes the proof of the proposition. □

Having dealt with the case when the X_λ's are in good position, we say that the family $\{X_\lambda\}_{\lambda\in\Lambda}$ is *in good enough position* if whenever we find incomparable elements U and V of E which do not cross, there is some element W of E which crosses them. As before, the above discussion applies equally well if the X_λ's are in good enough position. If this condition does not hold, we need to insist that E has only finitely many G–orbits which consist of isolated elements. Now we recall the proof of Lemma 3.13. This consisted of first replacing each isolated X_λ by an equivalent almost invariant set Y_λ associated to a splitting, and then replacing the Y_λ's by equivalent almost invariant sets Z_λ such the collection of translates of the Z_λ's and their complements is nested. In the present situation some isolated X_λ may be over an infinitely generated group H_λ, but in this case our hypothesis is that X_λ is already associated to a splitting, so we can simply take Y_λ to equal X_λ. For the second stage, Theorem 5.16 tells us that we can still find the required sets Z_λ. Finally the proof of Lemma 3.17 still applies. We conclude that if E has only finitely many G–orbits which consist of isolated elements, we can arrange that the X_λ's are in good enough position and then can define the pretree P. Further, if P is discrete, one can construct an algebraic regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ of the X_λ's which depends only on the equivalence classes of the X_λ's, so long as we are careful about isolated X_λ's.

Finally, if the given family is finite, say $X_1, \ldots, X_n$, we claim that P is discrete, so that $\Gamma(X_1, \ldots, X_n : G)$ always exists in this case. By renumbering we can suppose that $H_1, \ldots, H_k$ are the only non-finitely generated H_i's. Let E_k denote the set of all translates of $X_1, \ldots, X_k$ and their complements, and let F denote the set of all translates of $X_{k+1}, \ldots, X_n$ and their complements. By replacing $X_1, \ldots, X_k$ by equivalent almost invariant sets, as discussed above, we can arrange that E_k is nested. As at the end of the proof of Lemma 5.10, any element of E_k lies between two elements of F and vice versa. But we already know that E_k and F are each discrete. It follows that E itself is discrete, so that P is discrete as required.

When we consider this more general situation, we will need the following generalisation of Proposition 5.7.

Proposition 5.23. — *Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. If H_λ is not finitely generated, then we assume that X_λ is associated to a splitting of G over H_λ. Suppose that the regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ can be constructed as discussed above. Further suppose that the X_λ's are in good position, and that there is more than one CCC, so that the pretree P is not a single point. Let T denote the G–tree with $V_0(T) = P$. Let X be a nontrivial H–almost invariant subset of G which does not cross any element of E. Then the following statements hold.*

(1) *If H is finitely generated, then X is enclosed by a V_1–vertex of T.*

(2) *If X is a standard H–almost invariant set associated to a splitting of G over H, which need not be finitely generated, then X is enclosed by a V_1–vertex of T.*

Proof. — Most of the proof of Proposition 5.7 still goes through as it does not depend on the H_λ's being finitely generated. The only place where this is used is in the third paragraph of the proof where we show that X must be sandwiched between two elements of E in the case when H is finitely generated. (If H is not finitely generated, the proof of Proposition 5.7 still applies.)

In order to complete the proof of Proposition 5.23, we will consider the situation where H is finitely generated and Y is an element of E whose stabiliser is not finitely generated. We will use much the same argument as in the proof of Proposition 5.7 to show that X must be sandwiched between two translates of Y or Y^*.

Suppose that X is not sandwiched between two translates of Y or Y^*. Then, by replacing X by X^*, if necessary, we have that for every translate U of Y or Y^*, either $U < X$ or $U^* < X$. In particular, by replacing Y by Y^*, if necessary, we can assume that $Y < X$. Thus $X^* < Y^*$, so that $gX^* < gY^*$, for all g in G. Now we use the hypothesis that Y is associated to a splitting of G, so that for every g in G, one of the inclusions $gY \subset Y$, $gY \subset Y^*$, $gY^* \subset Y$, $gY^* \subset Y^*$ holds. Suppose that there is g in G such that $gY^* \subset Y$. As $gX^* < gY^*$, it follows that $gX^* < Y$. If we write k for g^{-1}, this implies that $X^* < kY$. But we know that $kY < X$ or $kY^* < X$. The first implies that $X^* < X$ which is impossible, and the second implies that X is equivalent to kY^*, which is impossible as the stabiliser of X is finitely generated, but the stabiliser of Y is not.

If there is no g in G such that $gY^* \subset Y$, then Lemma 5.5 shows that Y is associated to an ascending HNN extension. Let $f : G \to \mathbb{Z}$ denote the natural homomorphism associated to this HNN extension. As X does not cross Y, it follows, in particular, that the image in $\mathbb{Z}$ of the coboundary δX of X is bounded above or below. Hence the stabiliser of X must be contained in $\ker(f)$. It follows that the image of δX in $\mathbb{Z}$ must be finite. As the image in $\mathbb{Z}$ of δgY is finite for all g in G, there must be an element g of G such that $X \subset gY$. But we know that $gY < X$ or $gY^* < X$. The first implies that X is equivalent to gY, and the second implies that $gY^* < gY$, which are both impossible as in the preceding paragraph. This contradiction completes the proof that, in all cases, X must be sandwiched between two elements of E. □

CHAPTER 6

ALGEBRAIC REGULAR NEIGHBOURHOODS: EXISTENCE AND UNIQUENESS

In this chapter, we finally give our precise definition of a regular neighbourhood, and prove existence and uniqueness results which correspond closely to the situation in topology. Then we discuss further generalisations and some applications. The rest of this paper uses heavily the existence and uniqueness of regular neighbourhoods.

In order to motivate our definition, we recall the discussion in the introduction of the characteristic submanifold $V(M)$ of a 3–manifold M and the graph of groups decomposition Γ of $G = \pi_1(M)$, whose underlying graph is dual to the frontier $\mathrm{fr}(V(M))$ of $V(M)$. The graph Γ is naturally bipartite, with V_0–vertices corresponding to components of $V(M)$ and V_1–vertices corresponding to components of $M - V(M)$. The following two properties of $V(M)$ have algebraic analogues. The first is the Enclosing Property, which says that any essential annulus or torus in M is homotopic into $V(M)$. The second is that if F is any embedded essential closed surface in M, not necessarily a torus, and if F has intersection number zero with every essential annulus and torus in M, then F is homotopic into $M - V(M)$. These conditions are not sufficient to characterise $V(M)$ up to isotopy, but they do contain much of the information needed for such a characterisation. The algebraic analogue of the Enclosing Property is that the almost invariant subsets of G which correspond to essential annuli or tori are enclosed by the V_0–vertices of Γ. The algebraic analogue of the second property is that the splitting associated to F is enclosed by a V_1–vertex of Γ.

Now let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. Then we want our algebraic regular neighbourhood of the X_λ's in G to be a bipartite graph of groups structure Γ for G such that the V_0–vertices of Γ enclose the X_λ's, and splittings of G which have intersection number zero with each X_λ are enclosed by the V_1–vertices of Γ. In the preceding chapters, we discussed how to construct a bipartite graph of groups structure $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. In chapter 3, we showed that this construction always works if Λ is finite and each H_λ is finitely generated. However Example 6.11

will show that the edge groups of $\Gamma(X_1, \ldots, X_n : G)$ need not be finitely generated even if G and each H_i is finitely presented. Thus we need to consider splittings over non-finitely generated subgroups of G. This is why we formulate our definitions without assuming that the H_λ's are finitely generated, although in most of this paper, we will restrict to the case when the H_λ's are finitely generated.

Example 3.11 shows that if the X_λ's are a finite family of compatible splittings of G, then our construction of a regular neighbourhood of X yields a bipartite graph of groups Γ whose V_0–vertices correspond to the splittings. Further each V_0–vertex of Γ is of valence two, and the incident edge groups include by an isomorphism into the vertex group. Thus each V_0–vertex is redundant. This motivates part of our definition below.

It will be convenient to say that a vertex v of a graph of groups Γ is *isolated* if it is redundant and has valence two. Thus v has exactly two incident edges and each edge group includes by an isomorphism into the vertex group at v.

Here is our definition of an algebraic regular neighbourhood of the X_λ's in G.

Definition 6.1. — Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. Let E denote the set of all translates of the X_λ's and their complements. Then an *algebraic regular neighbourhood of the X_λ 's in G* is a bipartite graph of groups structure Γ for G such that the following conditions hold:

(1) Each X_λ is enclosed by some V_0–vertex of Γ, and each V_0–vertex of Γ encloses some X_λ.

(2) If σ is a splitting of G over a subgroup H (which need not be finitely generated) such that σ does not cross any $H_\lambda\backslash X_\lambda$, then σ is enclosed by some V_1–vertex of Γ.

(3) Γ is minimal.

(4) There is a bijection f from the G–orbits of isolated elements of $\overline{E}$ to the isolated V_0–vertices of Γ, such that $f(\overline{X})$ encloses X.

(5) Any non-isolated V_0–vertex of Γ encloses some non-isolated element of E.

Remark 6.2. — As G is finitely generated and Γ is minimal, Γ must be finite. Now the existence of the bijection of Condition 4) implies that E contains only finitely many G–orbits of isolated elements.

It has been implicit in all our discussion so far that we are considering a non-empty family of X_λ's. However, this definition makes perfect sense if the family is empty. In this case, the fact that each V_0–vertex of Γ encloses some X_λ implies that Γ has no V_0–vertices, so that Γ must consist of a single V_1–vertex labelled G. Clearly such a graph of groups structure for G does satisfy all the above conditions, so that we have the existence and uniqueness of algebraic regular neighbourhoods in this case.

Before we prove our existence and uniqueness results, we discuss the motivation behind this definition.

Conditions 1) and 2) are expected from our previous discussions. Condition 3) is clearly necessary in order to be able to prove any uniqueness result. Conditions 4) and 5) only come into play when there are isolated X_λ's.

Consider a finite family of immersed circles C_λ in a surface M, and consider the problem of characterising their regular neighbourhood N. Condition 1) above is analogous to asserting that each C_λ is homotopic into N, and that for each component of N, some C_λ is homotopic into it. Condition 2) is analogous to asserting that any simple curve on M which has intersection number zero with each C_λ is homotopic into $M - N$. If we assume that the C_λ's are π_1–injective in M, and also assume that each component of ∂N is π_1–injective in M, then these conditions characterise N, unless some C_λ is homotopic to an embedding disjoint from all the other C_λ's. In this case, N clearly has some annulus components, and the two conditions above do not completely determine N, because they do not control the number of such components. First, one can always add annulus components to N parallel to other such components without affecting the above two conditions. Second if there are two C_λ's which are simple and disjoint from each other and from all other C_λ's, so that N has two parallel annulus components, the subsurface of M obtained from N by simply deleting one of these annulus components will still satisfy the above two conditions. Condition 4) is to deal with this problem of non-uniqueness.

If we consider the case when the C_λ's are arcs, there is a more subtle problem. Suppose that we have three disjoint simple arcs C_1, C_2 and C_3 in M, such that a component of M cut along the C_i's is a disc D with copies of C_1, C_2 and C_3 in its boundary. Let N denote a regular neighbourhood of the C_i's. If we enlarge N by adding a disjoint copy of D, the new submanifold still satisfies the two conditions above which correspond to conditions 1) and 2) of our definition of an algebraic regular neighbourhood. Condition 5) deals with this problem of non-uniqueness. This example has nothing to do with the triviality of the groups involved. Taking the product of this example with the circle S^1 yields three annuli $C_i \times S^1$ in the 3–manifold $M \times S^1$, and we can enlarge a regular neighbourhood of these three annuli by adding the solid torus $D \times S^1$. This example is related to some subtleties in the topological JSJ-decomposition of a Haken 3–manifold.

An immediate consequence of condition 4) is that if one considers a single H–almost invariant set X which is associated to a splitting, then its regular neighbourhood has one V_0–vertex which must be isolated. This is what we wanted to happen, but the point we are making here is that this fact is really built into the definition and does not follow from any theory.

Next we note the following result which is an immediate consequence of our definition of a regular neighbourhood and the fact that if a vertex of a G–tree encloses a nontrivial almost invariant subset X of G, it also encloses any almost invariant subset Y of G which is equivalent to X.

Lemma 6.3. — *Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ-almost invariant subset of G, and let W_λ be a nontrivial K_λ-almost invariant subset of G which is equivalent to X_λ. Assume further, that the correspondence between the isolated X_λ's and the isolated W_λ's induces a bijection between the G-orbits of isolated $\overline{X_\lambda}$'s and the G-orbits of isolated $\overline{W_\lambda}$'s. Then a bipartite graph of groups structure Γ for G is a regular neighbourhood of the X_λ's if and only if it is a regular neighbourhood of the W_λ's.*

Proof. — Checking the first three conditions of Definition 6.1 is trivial, as is checking condition 5). The technical hypothesis of the lemma about isolated elements is required in order to check condition 4). □

One more definition will be very useful later in this paper. The above lemma shows that a regular neighbourhood is essentially determined by a collection of equivalence classes of almost invariant subsets of G. Thus it will be convenient to define a regular neighbourhood of a family of such equivalence classes.

Definition 6.4. — Let G be a finitely generated group with a family $\mathcal{F}$ of equivalence classes of nontrivial almost invariant subsets. Then *an algebraic regular neighbourhood of $\mathcal{F}$ in G* is an algebraic regular neighbourhood of a family of almost invariant subsets of G obtained by picking a representative of each equivalence class in $\mathcal{F}$, subject to the condition that if A and B are elements of $\mathcal{F}$ such that $B = gA$, for some g in G, then the representatives X and Y chosen for A and B must satisfy $Y = gX$.

Remark 6.5. — The reason for requiring equivariance in the choice of representatives is to ensure that each equivalence class in $\mathcal{F}$ has a unique representative. This condition is not needed unless $\mathcal{F}$ has isolated elements.

Let F denote the collection of all the almost invariant subsets of G which represent elements of $\mathcal{F}$, and let Γ denote an algebraic regular neighbourhood of $\mathcal{F}$ in G. If no element of F is isolated, then Γ is also an algebraic regular neighbourhood of the collection F. However, if some element X of F is isolated, then the collection F does not have a regular neighbourhood. This is because it contains all the almost invariant subsets of G which are equivalent to X, and there are clearly infinitely many distinct such sets.

Now we are ready to prove existence and uniqueness results for algebraic regular neighbourhoods. For our existence result, we need to restrict to finite families of almost invariant subsets of G, but our uniqueness result does not need this restriction.

***Theorem 6.6* (Existence of algebraic regular neighbourhoods)**

Let G be a finitely generated group, and for each $1 \leqslant i \leqslant n$, let H_i be a subgroup of G, and let X_i be a nontrivial H_i-almost invariant subset of G. If H_i is not finitely generated, then we assume that X_i is associated to a splitting of G over H_i.

Then there exists an algebraic regular neighbourhood of the X_i's in G.

Proof. — Recall that we constructed $\Gamma(X_1, \ldots, X_n : G)$ in chapter 3 in the case when every H_i is finitely generated, and constructed it at the end of the previous chapter in the general case.

We start by considering the case when the X_i's are in good position and no isolated X_i is invertible.

We will write Γ for $\Gamma(X_1, \ldots, X_n : G)$. Proposition 5.4 shows that Γ satisfies conditions 4) and 5) of Definition 6.1, as the V_0–vertices of T are precisely the CCC's of $\overline{E}$. Note that Proposition 5.4 applies even when some H_i's are not finitely generated.

Suppose that each H_i is finitely generated. We showed in Lemma 5.1 that Γ satisfies condition 1). We showed in Proposition 5.7 that Γ satisfies condition 2), and we showed in Lemma 5.2 that T is a minimal G–tree so that Γ satisfies condition 3). Thus $\Gamma(X_1, \ldots, X_n : G)$ is an algebraic regular neighbourhood of the X_i's in G. Note that in Proposition 5.7 we proved a result which is stronger than condition 2) in the case when H is finitely generated, as it applies to H–almost invariant subsets of G which need not be associated to splittings.

Now suppose that some H_i is not finitely generated. Lemma 5.21 shows that Γ satisfies condition 1), Proposition 5.23 shows that Γ satisfies condition 2), and Lemma 5.22 tells us that T is a minimal G–tree, so that Γ satisfies condition 3). It follows that Γ is an algebraic regular neighbourhood of the X_i's in G.

If the X_i's are in good enough position, the same arguments apply. If the X_i's are not in good enough position or some isolated X_i is invertible, we proceed as in case 3) of Summary 3.16. First we can replace the X_i's by a subfamily such that distinct $\overline{X_i}$'s lie in distinct G–orbits. A bipartite graph of groups structure for G is a regular neighbourhood of the original family if and only if it is a regular neighbourhood of the subfamily. Renumber the X_i's so that only $X_1, \ldots, X_k$ are isolated. Recall from Lemma 3.13 that each isolated X_i is equivalent to an almost invariant set Z_i which is not invertible such that the family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$ is in good position. Suppose that we are in case 3a), so that if i and j are distinct and X_i and X_j are isolated, then no translate of X_j is equivalent to X_i or X_i^*. Then as discussed in Lemma 3.14, we define $\Gamma(X_1, \ldots, X_n : G)$ to be $\Gamma(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$. As we have just shown that $\Gamma(Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n : G)$ is a regular neighbourhood of the family $Z_1, \ldots, Z_k, X_{k+1}, \ldots, X_n$, Lemma 6.3 shows that $\Gamma(X_1, \ldots, X_n : G)$ is a regular neighbourhood of the family $X_1, \ldots, X_n$, as required. If we are in case 3b), we replace the X_i's by a subfamily $X_1, \ldots, X_m$ which satisfies the condition of case 3a). Thus $\Gamma(X_1, \ldots, X_m : G)$ is a regular neighbourhood of $X_1, \ldots, X_m$. We construct $\Gamma(X_1, \ldots, X_n : G)$ by subdividing certain edges of $\Gamma(X_1, \ldots, X_m : G)$ which are incident to V_0–vertices arising from isolated X_i's. It now follows easily that $\Gamma(X_1, \ldots, X_n : G)$ is a regular neighbourhood of the family $X_1, \ldots, X_n$, as required. □

Theorem 6.7 **(Uniqueness of algebraic regular neighbourhoods)**
Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. If Γ_1 and Γ_2 are algebraic regular neighbourhoods of the X_λ 's in G, then they are naturally isomorphic, preserving their bipartite structures.

Remark 6.8. — If Λ is finite, then the construction of chapter 3 yields a regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. It follows that all regular neighbourhoods of the X_λ 's in G are isomorphic to $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$. However it seems conceivable that, when Λ is infinite, the X_λ's could be in good position and possess a regular neighbourhood, but that the construction of chapter 3 does not yield a regular neighbourhood because the pretree P is not discrete.

Proof. — Let Γ denote any algebraic regular neighbourhood of the X_λ's in G, let T denote the universal covering G–tree of Γ, and let E denote the set of all translates of the X_λ's and their complements. Recall from Remark 6.2 that the existence of Γ implies that E contains only finitely many G–orbits of isolated elements. Now Lemma 6.3 implies that we will lose nothing by assuming that the X_λ's are in good enough position. We need to prove some general facts.

Suppose that an element U of E is enclosed by distinct vertices v_1 and v_2 of T. Part 1) of Lemma 4.9 tells us that U is equivalent to Z_s or to Z_s^* for each edge s on the path joining v_1 and v_2. As each element of E is enclosed by some vertex of T, no element of E crosses any Z_s. Hence U must be an isolated element of E. It follows that if U is a non-isolated element of E, then it is enclosed by a unique vertex of T. Next suppose that U_1 and U_2 are elements of E which are enclosed by distinct vertices v_1 and v_2 of T. Then U_1 and U_2 do not cross, as $\overline{Z_s}$ lies between $\overline{U_1}$ and $\overline{U_2}$, for any edge s of T between v_1 and v_2. It follows that if U is a non-isolated element of E, then U is enclosed by a unique vertex v of T, and any element of $\overline{E}$ which lies in the CCC containing $\overline{U}$ must also be enclosed by v.

Now let Γ_1 and Γ_2 denote two algebraic regular neighbourhoods of the X_λ's in G. We will first suppose that Γ_1 and Γ_2 do not have any redundant vertices. As Γ_1 and Γ_2 are minimal, this is equivalent to assuming that they have no isolated vertices. Hence condition 4) implies that no X_λ is isolated. Now it also follows that each X_λ is enclosed by a unique vertex in each graph of groups. We will show that Γ_1 and Γ_2 have the same conjugacy classes of edge splittings, which will then imply that they are isomorphic by Theorem 5.17 as required. Note that V_0–vertices of Γ_1 must correspond to V_0–vertices of Γ_2 under this isomorphism because the V_0–vertices enclose the X_λ's and the V_1–vertices do not. Thus the isomorphism of Γ_1 and Γ_2 automatically preserves their bipartite structure.

Let σ be an edge splitting of Γ_2 over a subgroup H of G. (Recall that σ is defined by collapsing Γ_2 with the interior of an edge removed.) If X is a H–almost invariant subset of G associated to σ, then X does not cross any translate of any X_λ. Thus condition 2) for Γ_1 implies that X is enclosed by some V_1–vertex of Γ_1. Now Lemma 5.10 shows that we can refine the graph of groups Γ_1 by splitting at this V_1–vertex using the edge splitting σ. If the new graph of groups has a redundant vertex, this can only be because Γ_1 already had an edge splitting conjugate to σ. Now let Γ_{12} denote the graph of groups structure for G obtained from Γ_1 by splitting at V_1–vertices using those edge splittings of Γ_2 which are not conjugate to any edge splitting of Γ_1. (Note that Theorem 5.16 shows that we can form Γ_{12}.) Thus Γ_{12} has no redundant vertices. Similarly let Γ_{21} be obtained from Γ_2 by splitting at V_1–vertices using those edge splittings of Γ_1 which are not conjugate to any edge splitting of Γ_2, so that Γ_{21} also has no redundant vertices. As Γ_{12} and Γ_{21} have exactly the same conjugacy classes of edge splittings, Theorem 5.16 tells us that they are isomorphic. Now consider the universal covering G–trees T_{12} and T_{21}. Although T_{12} and T_{21} are not bipartite, they still have the property that each element of E is enclosed by some vertex. Further, as each element of E is non-isolated, there is a unique vertex of T_{12} which encloses every element U of E which lies in a given CCC of $\overline{E}$, and there is a unique vertex of T_{21} with the same property. Thus the G-isomorphism between T_{12} and T_{21} preserves these vertices. For each such vertex, consider the edge splittings associated to the incident edges. In one case they are all edge splittings of Γ_1, and in the other case they are all edge splittings of Γ_2. The G–isomorphism between the trees now implies that each edge splitting of Γ_1 is conjugate to some edge splitting of Γ_2, and conversely. Thus Γ_1 and Γ_2 have the same conjugacy classes of edge splittings, as required.

Now consider the case when Γ_1 and Γ_2 may have isolated vertices. Note that even if there are no isolated X_λ's, it is possible for a V_1–vertex to be isolated. Now condition 5) tells us that a non-isolated V_0–vertex must enclose a non-isolated element of E, and hence is the unique vertex which encloses this element of E. We want to apply the construction of the previous paragraph but first we need to remove all the isolated vertices of Γ_1 and Γ_2, by amalgamating suitable segments to a single edge. The resulting graphs of groups Γ_1' and Γ_2' are no longer bipartite. But the non-isolated V_0–vertices do not get altered. The argument of the previous paragraph applies to show that Γ_1' and Γ_2' are isomorphic. Now Γ_1 and Γ_2 are obtained from this common graph of groups structure by subdividing some edges, and condition 4) implies that the same edges get subdivided the same number of times, so that Γ_1 and Γ_2 must be isomorphic. As before, this isomorphism must preserve the non-isolated V_0–vertices of Γ_1 and Γ_2, so it follows that the isomorphism must preserve the bipartite structure, except possibly when every V_0–vertex of Γ_1 and of Γ_2 is isolated. Now an isomorphism between Γ_1 and Γ_2 which does not preserve the bipartite structure must reverse it, i.e. the V_0–vertices of one must correspond to the V_1–vertices of the other. We conclude that the isomorphism must preserve the bipartite structure, except possibly when every vertex is isolated. In this case, Γ_1 and Γ_2 will each be a circle, and it is trivial to change the isomorphism to one which does preserve the bipartite structure. □

Now we can use our results about regular neighbourhoods to give the result promised in chapter 2, that nontrivial almost invariant subsets of a group with infinitely many ends are never 0–canonical. Note that regular neighbourhoods are not essential for this argument. They are simply convenient.

Lemma 6.9. — *Let G be a finitely generated group with infinitely many ends. If X is any nontrivial H–almost invariant subset of G, where H is finitely generated, or if X is associated to a splitting of G over H, then X is not 0–canonical.*

Proof. — As G has infinitely many ends, it admits a splitting σ over a finite subgroup K. Let Y denote a standard K–almost invariant set associated to σ. Suppose that X is 0–canonical. In particular, X has zero intersection number with Y. Now we consider the regular neighbourhood $\Gamma(X, Y : G)$, and its universal covering G–tree T. As σ is a splitting and has intersection number zero with X, it follows that Y is isolated and so there is a corresponding isolated V_0–vertex of T. In particular, the edges of T adjacent to this vertex also have finite stabiliser. Let v be a V_0–vertex of T which encloses X, and pick edges s and t of T with finite stabiliser such that one lies in $\Sigma_v(X)$ and the other lies in $\Sigma_v(X^*)$. Choose s and t to be oriented away from v. Then X crosses $Z_s \cup Z_t$, showing that X is not 0–canonical, as claimed. □

We can also give our example, promised in Remark 3.9, of a finitely presented group G which splits over finitely presented subgroups H and K, such that the algebraic regular neighbourhood Γ of these splittings has some edge and vertex groups which are not finitely generated. We start with a general construction.

Example 6.10. — This construction will give many examples of a group G which has two splittings with intersection number 1, and also yields the regular neighbourhood of these two splittings. Our construction is based on the following topological picture. Consider two arcs l and m embedded properly in a surface M so that each arc separates M and the two arcs meet transversely in a single point w. Thus a regular neighbourhood N of the union of the two arcs has four frontier arcs and $M - N$ has four components. Let Γ denote the graph of groups structure for G determined by the frontier $\mathrm{fr}(N)$ of N. Thus Γ is a tree which has a single vertex v_0 corresponding to N, has four edges corresponding to the components of $\mathrm{fr}(N)$ and four other vertices corresponding to the components of $M - N$. The vertex v_0 and the four edges all carry the trivial group. We will use this simple picture to guide us in constructing a group G which corresponds to $\pi_1(M)$ and possesses splittings over subgroups H and K which correspond to the arcs l and m. The group $H \cap K$ will correspond to the point w. Finally the regular neighbourhood of the two splittings will be a graph of the same combinatorial type as Γ, with a single V_0–vertex corresponding to v_0.

To understand our idea, consider constructing M by starting with a point w, adding the four halves of l and m to obtain $l \cup m$, then constructing N, and finally adding in the four remaining pieces of M. We will follow a similar procedure, but we will use spaces with nontrivial fundamental groups.

Start with a group C and groups A, B, D and E which each contain C as a proper subgroup. We will assume that the intersection of any two of these groups is C. Let $H = A *_C B$ and let $K = D *_C E$. Think of C as corresponding to the point w, the groups A and B as corresponding to the two halves of l, and the groups D and E as corresponding to the two halves of m. Thus H corresponds to l and K corresponds to m. Let L_0 denote $H *_C K$, which corresponds to $\pi_1(N)$. Now the four components of $\mathrm{fr}(N)$ naturally have corresponding groups which are $A *_C D$, $D *_C B$, $B *_C E$ and $E *_C A$, each of which is a subgroup of L_0. Denote these by L_1, L_2, L_3 and L_4 respectively. For $i = 1, 2, 3, 4$, pick a group G_i which properly contains L_i, and think of the G_i's as corresponding to the components of $M - N$. Now we define the group G to be the fundamental group of the graph of groups Γ which is a tree with a vertex v_0 with associated group L_0, with four edges attached to v_0 which carry the L_i's, and with the four other vertices carrying the G_i's. Let v_i denote the vertex which carries G_i.

To understand this construction topologically, one needs to build a space with fundamental group G which mimics the structure of our initial example M. We pick spaces M_A, M_B, M_C, M_D, and M_E with fundamental groups A, B, C, D and E respectively, such that each space contains M_C and the intersection of any two equals M_C. Let Z denote the union of these spaces, so that $\pi_1(Z) = L_0$. Then for each $i \geqslant 1$, we choose a space M_i with fundamental group L_i, take its product with the unit interval and glue one end to Z using the inclusion of L_i into L_0. Finally, for each G_i, we choose a space with fundamental group G_i and glue the other end of $M_i \times I$ to it using the inclusion of L_i in G_i. The resulting space M has fundamental group G. Further, its structure clearly yields splittings σ and τ over H and K respectively. For σ is the splitting of G obtained by "cutting along" $M_A \cup M_B$, and τ is the splitting of G obtained by "cutting along" $M_D \cup M_E$. Consider the pre-image $\widetilde{Z}$ of Z in the universal cover $\widetilde{M}$ of M. If G is finitely generated, then it is easy to see that the pretree constructed combinatorially in chapter 3 is the same as the pretree of components of $\widetilde{Z}$. (The main point to notice is that the stabiliser of a component of $\widetilde{Z}$ equals L_0, which also equals the stabiliser of the corresponding CCC. Thus Z is in 'good position'.) It follows that σ and τ have intersection number 1, and that Γ is their regular neighbourhood in G.

By making interesting choices of the groups involved in the above construction, one can give many interesting examples. Here is an important example which explains why we have spent so much time considering splittings over non-finitely generated subgroups.

Example 6.11. — We give here an example of a finitely presented group G which splits over finitely presented subgroups H and K, such that the algebraic regular neighbourhood Γ of these splittings has an edge group and a vertex group which is not finitely generated. We do this by making choices of the groups involved in the construction of Example 6.10. The edge group L_1 and the vertex group G_1 of the regular neighbourhood are not finitely generated.

In Example 6.10, choose $C = F_\infty$, the free group of countably infinite rank, choose B and E to be F_2, the free group of rank 2, and choose A and D to be $F_2 * C$. The inclusions of C in A and D are the obvious ones. The inclusions of C in B and E can be chosen in any reasonable way. A good example would be to map the i–the basis element of C to $u^{-i}vu^i$, where u and v are the basis elements of F_2. Thus $H = A *_C B$ and $K = D *_C E$ are each isomorphic to $(F_2 * C) *_C F_2$ which is simply F_4, the free group of rank 4. In order to complete the construction of the group G, we need to choose the groups G_i. Recall that the groups L_1, L_2, L_3 and L_4 are respectively isomorphic to $A *_C D$, $D *_C B$, $B *_C E$ and $E *_C A$. This means that L_1 is not finitely generated, though the remaining L_i's are finitely generated. Further, L_2 and L_4 are each isomorphic to F_4. The group $L_3 = B *_C E = F_2 *_C F_2$ is finitely generated but is not finitely presented. It is trivial to choose finitely presented groups G_2 and G_4 which properly contain L_2 and L_4 respectively. We can use Higman's Embedding Theorem [**23**] to find a finitely presented group G_3 which contains L_3. Finally, we choose G_1 to be any group of the form $P *_Q L_1$, where Q is finitely generated and P is finitely presented. Now we claim that G is finitely presented. To see this, we build up G in stages. Recall that, by construction, each of G_2, G_3 and G_4 is finitely presented. Thus the group $G_2 *_B G_3$ is finitely presented. Hence the group $(G_2 *_B G_3) *_E G_4$ is finitely presented. Denote this last group by G_5. Then $G = G_1 *_{L_1} G_5 = (P *_Q L_1) *_{L_1} G_5 = P *_Q G_5$, which is finitely presented because P and G_5 are finitely presented and Q is finitely generated.

Remark 6.12. — In the above example, the group G has subgroups which are finitely generated but not finitely presented. This raises the question of whether examples such as the above can exist for a group G which is coherent, i.e. finitely generated subgroups are finitely presented. We have no ideas about how to answer this question.

We next consider an application which strengthens a result of Niblo in [**35**] on the existence of splittings of a given group. Let H be a finitely generated subgroup of a finitely generated group G, let X be a nontrivial H–almost invariant subset of G, and consider the regular neighbourhood $\Gamma(X : G)$. Recall that Proposition 5.2 implies that $\Gamma(X : G)$ is minimal. It follows that unless it consists of a single vertex, then any edge will yield a splitting of G. We define the subgroup $S(X)$ of G, to be the stabiliser of the CCC of $\overline{E}$ which contains $\overline{X}$ in the construction of $\Gamma(X : G)$, so that $S(X)$ is the vertex group for the corresponding V_0–vertex of $\Gamma(X : G)$. Thus we immediately deduce the following result.

Corollary 6.13. — *Let G be a finitely generated group with finitely generated subgroup H, and let X be a nontrivial H–almost invariant subset of G. If $S(X)$ is not equal to G, then G splits over a subgroup of $S(X)$.*

Remark 6.14. — As $S(X)$ was defined in terms of the regular neighbourhood $\Gamma(X:G)$, Lemma 6.3 implies that if X and Y are equivalent, then $S(X) = S(Y)$.

In order to understand the implication of this, we need to consider the group $S(X)$ more carefully. Let v denote the V_0–vertex of T which corresponds to the CCC of $\overline{E}$ which contains $\overline{X}$. If X is not isolated in the set E of all translates of X and X^*, then the argument at the end of the proof of Theorem 3.8 shows that $S(X)$ is generated by H and $\{g \in G : gX \text{ crosses } X\}$. In [**35**], Niblo defined a group $T(X)$ which is the subgroup of G generated by H and $\{g \in G : gX \text{ and } X \text{ are not nested}\}$. He proved that if $T(X) \neq G$, then G splits over a subgroup of $T(X)$. Clearly $S(X)$ is a subgroup of $T(X)$ in this case, so our result implies his when X is not isolated. If X is isolated, then Theorem B.2.8 implies that G splits over a subgroup commensurable with H.

An interesting related result due to Dunwoody and Roller [**13**] is that if $S(X)$ is contained in $\mathrm{Comm}_G(H)$, then G splits over a subgroup commensurable with H even if $G = \mathrm{Comm}_G(H)$.

The above Corollary was proved by considering the regular neighbourhood $\Gamma(X:G)$ and using the fact that $S(X)$ is a vertex group. If K is another finitely generated subgroup of G, and Y is a nontrivial K–almost invariant subset of G such that $H\backslash X$ and $K\backslash Y$ have intersection number zero, then considering the regular neighbourhood $\Gamma(X, Y : G)$ yields the following result, which strengthens another result of Niblo in [**35**].

Corollary 6.15. — *Let G be a finitely generated group with finitely generated subgroups H and K. If G has a nontrivial H–almost invariant subset X and a nontrivial K–almost invariant subset Y such that $H\backslash X$ and $K\backslash Y$ have intersection number zero, then G has a minimal graph of groups decomposition with at least two edges, in which two of the edge groups are a subgroup of $S(X)$ and a subgroup of $S(Y)$.*

Now we can use this corollary to give another generalisation of Theorem 2.35. In that result, we showed that if G is a finitely generated group with splittings σ and τ over finitely generated subgroups H and K of G such that σ and τ have zero intersection number, then σ and τ are compatible. In Theorem 5.16 of this paper, we generalised Theorem 2.35 to the case of splittings over subgroups H and K which need not be finitely generated. A natural question is whether an analogous result holds for almost invariant sets which are not associated to splittings.

An equivalent formulation of Theorem 2.35 is that if X is a H–almost invariant subset of G and Y is a K–almost invariant subset of G such that each is associated to a splitting of G and $H\backslash X$ and $K\backslash Y$ have intersection number zero, then X and Y are equivalent to subsets X' and Y' of G such that the set E of all translates of X'

and Y' is nested. We will prove the following result which is the natural analogue when X and Y are not associated to splittings.

Lemma 6.16. — *Let G be a finitely generated group with finitely generated subgroups H and K. Let X be a nontrivial H–almost invariant subset of G and let Y be a nontrivial K–almost invariant subset of G such that $H\backslash X$ and $K\backslash Y$ have intersection number zero. Then X and Y are equivalent to subsets X' and Y' of G such that any translate of X' and any translate of Y' are nested.*

Remark 6.17. — We now have two generalisations of Theorem 2.35, one of which allows splittings over non-finitely generated subgroups and the other replaces splittings over finitely generated subgroups H and K by almost invariant subsets over H and K. It seems natural to ask if there is a common generalisation for arbitrary nontrivial almost invariant sets over arbitrary subgroups of G. We have no methods for answering such a question but we think that such a generalisation is unlikely.

Proof. — Let Γ denote $\Gamma(X, Y : G)$, as constructed in chapter 3. Then Γ has two V_0–vertices, one of which encloses X and the other encloses Y. Now, Lemma 4.14 shows that X and Y are equivalent to subsets X' and Y' of G such that X' is nested with respect to all the Z_s and Z_s^* and so is Y'. In particular, the claim follows. □

The theory of algebraic regular neighbourhoods is analogous to the theory of topological regular neighbourhoods, but there are situations where a slight modification of the algebraic theory seems useful. We start by discussing the topological analogue again. If we have two disjoint simple closed curves on an orientable surface M, their regular neighbourhood consists of two disjoint annuli. In the algebraic context, if one has two H–almost invariant subsets X_1 and X_2 of G each associated to the same splitting but lying in distinct G–orbits, their algebraic regular neighbourhood will have two isolated V_0–vertices. On some occasions, we would prefer to ignore this multiplicity and have the algebraic regular neighbourhood of X_1 and X_2 be the same as the regular neighbourhood of X_1 on its own. Another point arises even when there is no problem of multiplicity. Let M be a surface and X be a subsurface whose boundary consists of essential circles in M. We can regard X as a kind of regular neighbourhood of all the essential loops in X and this corresponds closely to our theory of algebraic regular neighbourhoods. However, the family of all essential loops in M which can be homotoped into X contains in addition the boundary components of X, and so our topological regular neighbourhood of this larger family will consist of X and of disjoint annuli parallel to the boundary components of X. A similar phenomenon occurs in the algebraic context. Suppose one has a family $\mathcal{F}$ of nontrivial almost invariant subsets X_λ of G whose regular neighbourhood Γ has a single V_0–vertex v. Let $\mathcal{F}'$ be obtained from $\mathcal{F}$ by adding almost invariant subsets of G associated to the splittings determined by the edges of Γ incident to v. Then every element of $\mathcal{F}'$ is enclosed by v, but the regular neighbourhood of $\mathcal{F}'$ will have a V_0–vertex corresponding to v

and isolated V_0–vertices one for each of the extra almost invariant sets. On some occasions, we would prefer not to have these extra isolated V_0–vertices.

The above discussion motivates the idea of a reduced regular neighbourhood. We will first describe how to obtain such an object from an ordinary regular neighbourhood, and will then give an abstract characterisation and a more direct construction. Let Γ be a regular neighbourhood of some family of nontrivial almost invariant subsets of a group G. Suppose that Γ has a subgraph consisting of two edges e and e' such that $e \cap e'$ is a V_1–vertex w and the other vertices of e and of e' are v and v' respectively, where $v \neq v'$. Suppose further that v' and w are isolated. Thus v' and w each has valence 2, and if e'' denotes the other edge incident to v', then the splittings associated to e, e' and e'' are all conjugate. Then we "reduce" Γ by collapsing $e \cup e'$ to a single point which will be a V_0–vertex in the new graph Γ_1. This new graph still has fundamental group G, and is still bipartite. Repeating this procedure will eventually yield a bipartite graph of groups structure Γ' for G without such subgraphs. We will say that Γ' is *reduced bipartite*. This is the reduced regular neighbourhood associated to the original regular neighbourhood Γ. Note that we can obtain Γ from Γ' by simply subdividing certain edges, which shows that Γ' is independent of the order in which reductions are made. One can give an abstract definition of a reduced regular neighbourhood as follows. Note that this definition is identical to that for an ordinary regular neighbourhood except that condition 4) of Definition 6.1 has been removed, and the fact that Γ is reduced bipartite has been added to condition 3).

Definition 6.18. — Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. Let E denote the set of all translates of the X_λ's and their complements. Then a *reduced algebraic regular neighbourhood of the X_λ 's in G* is a bipartite graph of groups structure Γ for G such that the following conditions hold:

(1) Each X_λ is enclosed by some V_0–vertex of Γ, and each V_0–vertex of Γ encloses some X_λ.

(2) If σ is a splitting of G over a subgroup H (which need not be finitely generated) such that σ does not cross any element of E, then σ is enclosed by some V_1–vertex of Γ.

(3) Γ is minimal and reduced bipartite.

(4) Any non-isolated V_0–vertex of Γ encloses some non-isolated element of E.

The construction which preceded this definition shows that if the X_λ's possess an ordinary regular neighbourhood Γ, then they also possess a reduced one Γ'. Conversely, if the X_λ's possess a reduced regular neighbourhood Γ', and if there are only finitely many G–orbits of isolated elements, then they also possess an unreduced regular neighbourhood Γ which can be constructed from Γ' by subdividing appropriate edges to add isolated V_0–vertices. However, we will be interested in families of almost invariant sets which contain infinitely many G–orbits of isolated elements, and such a family cannot have an unreduced regular neighbourhood. Thus it will be useful to have a direct approach to the construction of a reduced regular neighbourhood.

As usual, we let G be a finitely generated group with a family $\mathcal{F}$ of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G, and let E denote the set of all translates of the X_λ's and their complements. We will replace $\mathcal{F}$ by a subfamily $\mathcal{F}'$ so that the usual regular neighbourhood of $\mathcal{F}'$ will be a reduced regular neighbourhood of $\mathcal{F}$. Recall from the discussion immediately before Lemma 3.14 that if X_i and X_j are isolated elements of E and some translate of X_j is equivalent to X_i or X_i^*, we say that the G–orbits of $\overline{X_i}$ and $\overline{X_j}$ are parallel. The first step in defining $\mathcal{F}'$ is to consider the parallelism classes of isolated X_λ's. For each such class, we simply remove all but one of the X_λ's in it to obtain a new family $\mathcal{F}''$. If this family has a (unreduced) regular neighbourhood Γ'', it will not have distinct isolated V_0–vertices which enclose the same splitting, but it is possible that there is an isolated V_0–vertex v such that the splitting σ of G enclosed by v is also enclosed by some other (non-isolated) V_0–vertex. If this happens, we alter $\mathcal{F}''$ by removing the X_λ associated to σ. Repeating this process will produce a subset $\mathcal{F}'$ of $\mathcal{F}$ such that the regular neighbourhood of $\mathcal{F}'$ is also a reduced regular neighbourhood.

One consequence of the preceding discussion is that any reduced regular neighbourhood is an unreduced regular neighbourhood of a suitable family of almost invariant sets. Thus we obtain the following existence and uniqueness results as immediate consequences of the existence and uniqueness results for unreduced regular neighbourhoods.

***Theorem 6.19* (Existence of reduced algebraic regular neighbourhoods)**

Let G be a finitely generated group, and for each $1 \leqslant i \leqslant n$, let H_i be a subgroup of G, and let X_i be a nontrivial H_i–almost invariant subset of G. If H_i is not finitely generated, then we assume that X_i is associated to a splitting of G over H_i.

Then there exists a reduced algebraic regular neighbourhood of the X_i's in G.

***Theorem 6.20* (Uniqueness of reduced algebraic regular neighbourhoods)**

Let G be a finitely generated group with a family of subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ–almost invariant subset of G. If Γ_1 and Γ_2 are reduced algebraic regular neighbourhoods of the X_λ 's in G, then they are naturally isomorphic, preserving their bipartite structures.

CHAPTER 7

COENDS WHEN THE COMMENSURISER IS SMALL

The rest of this paper consists of understanding regular neighbourhoods of certain families of almost invariant subsets of a group and showing that in several interesting cases, certain infinite families of such subsets possess a regular neighbourhood. For the results about infinite families of almost invariant subsets, we will need to assume that G is finitely presented, but most of our results about finite families work without this additional restriction.

We will be interested in a one-ended, finitely generated group G and nontrivial almost invariant subsets of G which are over virtually polycyclic (VPC) subgroups. In this and the following three chapters, we will mainly be interested in the case of $VPC1$ subgroups. However many of our results are valid for VPC subgroups of length $n > 1$ assuming that G has no nontrivial almost invariant subsets over VPC subgroups of length $< n$, and in this chapter we will prove several technical results in that generality for later use. Any $VPC1$ group is virtually infinite cyclic, or equivalently two-ended, and we will use the phrase two-ended in what follows. In chapter 10, we will show that if G is finitely presented, there is a regular neighbourhood of all the equivalence classes of nontrivial almost invariant subsets of G over two-ended subgroups. It turns out that if H is a two-ended subgroup of G such that there is a nontrivial almost invariant subset over H, then the commensuriser of H in G plays an important role. We will analyse the role of the commensuriser in this and the next chapter. First we define it.

Definition 7.1. — The *commensuriser* $\mathrm{Comm}_G(H)$ of a subgroup H of a group G consists of those elements g in G such that H and H^g are commensurable.

It is easy to see that $\mathrm{Comm}_G(H)$ is a subgroup of G which contains H.

We will say that a subgroup H of G *has small commensuriser in* G if $\mathrm{Comm}_G(H)$ contains H with finite index. In this chapter we will consider a one-ended finitely generated group G and nontrivial almost invariant subsets over two-ended subgroups.

To study these, we need Bowditch's results from [**8**] as well as a non-standard accessibility result. We start by quoting a result in [**8**], but reformulated in the language of almost invariant sets (see [**43**] for a similar result for hyperbolic groups).

Proposition 7.2. — *Let G be a one-ended finitely generated group, and let X and Y be nontrivial almost invariant subsets over two-ended subgroups H and K. If X crosses Y strongly, then Y crosses X strongly and the number of coends of both H and K is two.*

We will prove a more general result in chapter 13. The following result tells us what happens when X and Y cross weakly.

Proposition 7.3. — *Let G be a one-ended finitely generated group, and let X and Y be nontrivial almost invariant subsets over two-ended subgroups H and K. If X crosses Y weakly, then H and K are commensurable.*

Proof. — As H and K are virtually infinite cyclic, either they are commensurable or $H \cap K$ is finite. We will suppose that $H \cap K$ is finite and derive a contradiction.

By Proposition 7.2, as X crosses Y weakly, we know that Y crosses X weakly. Thus one of $\delta Y \cap X^{(*)}$ is H–finite and one of $\delta X \cap Y^{(*)}$ is K–finite. By changing notation if necessary, we can arrange that $\delta Y \cap X$ is H–finite and $\delta X \cap Y$ is K–finite. As δX is H–finite and δY is K–finite, it follows that each of $\delta Y \cap X$ and $\delta X \cap Y$ is both H–finite and K–finite. Thus they are both $(H \cap K)$–finite. Now consider the coboundary $\delta(X \cap Y)$. Every edge in this coboundary meets $\delta Y \cap X$ or $\delta X \cap Y$. Hence $\delta(X \cap Y)$ is also $(H \cap K)$–finite. As $X \cap Y$ is clearly invariant under the left action of $H \cap K$, it is $(H \cap K)$–almost invariant. Now $X \cap Y$ must be infinite as X and Y cross. As we are assuming that $H \cap K$ is finite, this means that $X \cap Y$ is a nontrivial almost invariant subset of G, so that G has more than one end, which is the required contradiction. □

We now recall from [**44**] that a pair of finitely generated groups (G, H) is of *surface type* if $e(G, H') = 2$ for every subgroup H' of finite index in H and $e(G, H') = 1$ for every subgroup H' of infinite index in H. It follows that (G, H) has two coends. Conversely, suppose that (G, H) has two coends. Using the notation at the end of chapter 2, this means that $E(H)$ has two elements. Thus H has a subgroup H_1 of index at most 2 which preserves the elements of $E(H)$. It follows that $e(G, H_1) = 2$, and hence that the pair (G, H_1) is of surface type. The following result will be useful.

Proposition 7.4. — *Let G be a one-ended finitely generated group with finitely generated subgroups H and K, a nontrivial H–almost invariant subset X, and a nontrivial K–almost invariant subset Y. Suppose also that the number of coends of H in G is 2. Then Y crosses X if and only if Y crosses X strongly.*

Proof. — By replacing H by the above subgroup H_1 of index at most 2, we can assume that the pair (G, H) is of surface type. Now Proposition B.3.7 tells us that if (G, H) is a pair of finitely generated groups of surface type, X is a nontrivial H–almost invariant subset of G and Y is a nontrivial K–almost invariant subset of G, then Y crosses X if and only if Y crosses X strongly. The result follows. □

The following result summarises the above in the form which we will need.

Proposition 7.5. — *Let G be a one-ended finitely generated group and let $\{X_\lambda\}_{\lambda\in\Lambda}$ be a family of nontrivial almost invariant subsets over two-ended subgroups of G. As usual, let E denote the set of all translates of the X_λ's and their complements. Form the pretree P of cross-connected components (CCC's) of $\overline{E}$ as in the construction of regular neighbourhoods in chapter 3. Then the following statements hold:*

(1) *The crossings in a CCC of $\overline{E}$ are either all strong or are all weak.*

(2) *In a CCC with all crossings weak, the stabilisers of the corresponding elements of E are all commensurable.*

Proof. — 1) If X and Y are elements of E which cross strongly, then Propositions 7.2 and 7.4 imply not only that Y must cross X strongly, but that the same applies to any other element of E which crosses X. Hence all crossings in the CCC determined by X and Y are strong. It follows that the crossings in a CCC of $\overline{E}$ are either all strong or are all weak, as required.

2) If a CCC has weak crossing, then Proposition 7.3 implies that any two elements of this CCC have commensurable stabilisers, as required. □

We will need the following definition.

Definition 7.6. — In a minimal graph of groups decomposition Γ of a group G, we will say that a vertex v is of *finite-by-Fuchsian type*, or that the associated vertex group $G(v)$ is of *finite-by-Fuchsian type*, if $G(v)$ is a finite-by-Fuchsian group, where the Fuchsian group is not finite nor two-ended, and there is exactly one edge of Γ which is incident to v for each peripheral subgroup K of $G(v)$ and this edge carries K.

Remark 7.7. — If $G = G(v)$, then the Fuchsian quotient group corresponds to a closed orbifold. We should note that usually a Fuchsian group means a discrete group of isometries of the hyperbolic plane, but in this paper, it will be convenient to include also discrete groups of isometries of the Euclidean plane. As we are excluding finite and two-ended Fuchsian groups, the extra groups this includes are all virtually $\mathbb{Z}\times\mathbb{Z}$.

Now suppose that the family of X_λ's in Proposition 7.5 is finite. Thus their regular neighbourhood $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ exists by Theorem 6.6. Consider a V_0–vertex v of $\Gamma(\{X_\lambda\}_{\lambda\in\Lambda} : G)$ which comes from a CCC of $\overline{E}$ in which all crossing is strong. If X_λ lies in the CCC corresponding to v and is not isolated, Lemma 7.2 shows that H_λ has two coends in G. Bowditch [**8**] and Dunwoody-Swenson [**15**] have shown that the enclosing group $G(v)$ is of finite-by-Fuchsian type. Bowditch deals only with the case when H_λ has two coends in G for each $\lambda \in \Lambda$. In this case, our construction of regular neighbourhoods coincides with Bowditch's construction of enclosing groups as both use the same pretree. In fact our construction of regular neighbourhoods is suggested by Bowditch's use of pretrees in [**5**] and [**8**]. We state his result in the form that we will use later. We will also need similar results for VPC subgroups of length greater than 1, but we will simply observe that the Bowditch-Dunwoody-Swenson arguments go through in general. The following result is contained in Propositions 7.1 and 7.2 of Bowditch's paper [**8**], and in the JSJ-decomposition theorem of [**15**] but we have reformulated it using our regular neighbourhood terminology.

Theorem 7.8. — *Let G be a one-ended finitely generated group with a finite family of two-ended subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a nontrivial H_λ-almost invariant subset of G, let E denote the set of all translates of the X_λ's and their complements, and let Γ denote the regular neighbourhood of the X_λ's. Let X denote an element of E, let H denote its stabiliser, and let v denote a vertex of Γ which encloses X.*

Suppose H has two coends and that there exists an element of E which crosses X. Then the vertex group $G(v)$ is of finite-by-Fuchsian type, and H is not commensurable with a peripheral subgroup of $G(v)$.

Remark 7.9. — If Γ consists of a single vertex, so that $G = G(v)$, then G must itself be of finite-by-Fuchsian type.

In a Fuchsian group Σ, any two-ended subgroup has small commensuriser unless Σ is virtually $\mathbb{Z} \times \mathbb{Z}$. When combined with the above theorem, this yields the following result.

Corollary 7.10. — *Let G be a one-ended, finitely generated group and suppose that X is a nontrivial almost invariant subset over a two-ended subgroup H. Suppose some nontrivial almost invariant set over a two-ended subgroup K crosses X strongly. Then either G is of finite-by-Fuchsian type or H has small commensuriser in G. Further, if G has a finite normal subgroup K with Fuchsian quotient Σ, then either H has small commensuriser in G, or Σ is virtually $\mathbb{Z} \times \mathbb{Z}$.*

Later we will want to consider infinite families of such almost invariant subsets of G. We want to do this by taking increasing finite families of such sets and showing that the graphs of groups structures for G obtained in this way must stabilise. If all crossings of such subsets of G are strong, one can use Theorem 7.8 to show that this happens for homological reasons. However if weak crossings occur such arguments do not work. We will need to assume that G is finitely presented and to use variants of previous accessibility results.

Recall that, in a graph of groups decomposition, we call a vertex redundant if it has valence at most two, it is not the vertex of a loop, and each edge group includes by an isomorphism into the vertex group. Recall that a vertex is reducible if it has two incident edges, it is not the vertex of a loop, and one of the incident edge groups includes by an isomorphism into that vertex group. For a finitely presented group G, the main result of [**4**] gives a bound on the complexity of reduced graphs of groups decompositions of G with all edge groups being small groups. We will use their result to prove the following.

Theorem 7.11. — *Let G be a finitely presented group and suppose that G does not split over VPC subgroups of length less than n. For each positive integer k, let Γ_k be a graph of groups decomposition of G without redundant vertices and with all edge groups being $VPCn$, and suppose that for each k, Γ_{k+1} is a refinement of Γ_k. Then the sequence Γ_k stabilises.*

Proof. — Since G is finitely presented, Bestvina-Feighn's accessibility result [**4**] implies the theorem provided the Γ_k's are reduced. Thus we only have to bound the length of chains of splittings of G over descending subgroups (unfoldings in the language of [**36**]). This was done in [**36**] when G is finitely generated and torsion-free and $n = 1$. In [**6**], Bowditch gave a much simpler argument using tracks when G is finitely presented and $n = 1$. We give our argument which is similar to Bowditch's but which was arrived at independently.

Suppose that we have an infinite sequence of splittings of G over descending $VPCn$ subgroups H_i. Thus $\cap_{i\geqslant 1} H_i$ is $VPC(n-1)$. We fix a finite 2–complex K with fundamental group G and universal cover $\widetilde{K}$. For each $m \geqslant 1$, there is a G–tree T_m which corresponds to the first m splittings, and T_{m+1} is a refinement of T_m, i.e. there is a natural collapsing map $q_{m+1} : T_{m+1} \to T_m$. We now pick G–equivariant linear maps $p_m : \widetilde{K} \to T_m$ such that $p_m = q_{m+1}p_{m+1}$, let W_m denote the midpoints of the edges of T_m and consider $p_m^{-1}(W_m)$. This is a G–invariant pattern in $\widetilde{K}$, which projects to a finite pattern L_m in K. By construction of the maps p_m, we have $L_m \subset L_{m+1}$. Each component of L_m carries a subgroup of H_m. Since G does not split over a VPC subgroup of length less than n, $L_{m+1} - L_m$ must have at least one component C_m with stabiliser which is $VPCn$. Now Dunwoody (see [**10**]) showed that there is an upper bound on the number of non-parallel disjoint tracks one can have in K. In particular, it follows that the C_m's carry only finitely many distinct subgroups of G, and hence that the descending sequence of H_i's must stabilise, which is a contradiction. □

Let Γ be a graph of groups decomposition of a group G without redundant vertices. We will say that Γ is *maximal with respect to two-ended subgroups*, if whenever a vertex encloses a splitting over a two-ended subgroup, this splitting is already an edge splitting of Γ. This means that if we form a refinement of Γ by splitting at some vertex so that the extra edge splitting is over a two-ended group, this refinement must have a redundant vertex. The above result in particular implies the following.

Corollary 7.12. — *A one-ended, finitely presented group has maximal decompositions with respect to two-ended subgroups.*

Similarly we call a decomposition of G maximal with respect to VPC groups of length $\leqslant n$, if it cannot be refined without introducing redundant vertices by splitting at a vertex along a VPC group of length $\leqslant n$.

The proof of Theorem 7.11 applies essentially unchanged to yield the following result.

Theorem 7.13. — *Let G be a finitely presented group and let Γ_0 be a graph of groups decomposition of G which is maximal with respect to splittings over VPC groups of length $\leqslant n$. For each $k \geqslant 1$, let Γ_k be a graph of groups decomposition of G without redundant vertices, and suppose that for each $k \geqslant 0$, Γ_{k+1} is a refinement of Γ_k. Suppose further that all the edge splittings of Γ_k which are not edge splittings of Γ_0 are over $VPC(n+1)$ subgroups. Then the sequence Γ_k stabilises.*

Now consider a two-ended subgroup H of G. Let $Q(H)$ denote the collection of all almost invariant subsets of G which are over subgroups of G commensurable with H, and let $F(H)$ denote the subset of $Q(H)$ which consists of all the trivial elements. (Recall that a H–almost invariant set is trivial if it is H–finite.) Clearly if X and Y lie in $Q(H)$, then $X \cap Y$, $X + Y$ and $X \cup Y$ also lie in $Q(H)$. Thus $Q(H)$ is a subalgebra of the Boolean algebra of all subsets of G. Also $F(H)$ is an ideal in $Q(H)$. We let $B(H)$ denote the quotient Boolean algebra $Q(H)/F(H)$. Thus $B(H)$ is precisely the collection of equivalence classes of elements of $Q(H)$. Note that if H has small commensuriser so that $\mathrm{Comm}_G(H)$ contains H with finite index, then $\mathrm{Comm}_G(H)$ is also two-ended.

Theorem 7.14. — *Let G be a one-ended, finitely presented group with a two-ended subgroup H with small commensuriser. Suppose that there is a nontrivial element X of $Q(H)$ such that no nontrivial almost invariant set over a two-ended subgroup of G crosses X strongly. Then $B(H)$ is finite.*

Remark 7.15. — Proposition 7.2 shows that the conclusion of this result remains true if X does cross strongly some nontrivial almost invariant set over a two-ended subgroup.

Proof. — If there is a nontrivial almost invariant subset Y of G which is over some two-ended subgroup and crosses X, then it must do so weakly and Lemma 7.3 implies that the stabilisers of X and Y are commensurable. This implies that Y also lies in $Q(H)$.

Given a finite subset $\{U_j\}_{j\in J}$ of nontrivial elements of $Q(H)$, we can form E, T and the regular neighbourhood $\Gamma(\{U_j\}_{j\in J} : G)$, which we will denote by Γ. Consider a V_0–vertex v of T such that the corresponding CCC of $\overline{E}$ contains $\overline{U}$, for some element U of the U_j's. The other elements of E which lie in this CCC must all lie in $Q(H)$, by the preceding paragraph. Now we consider $\mathrm{Stab}(v)$, the stabiliser of v. Recall from the proof of part 2) of Theorem 3.8 that there are a finite number of elements X_i of E enclosed by v such that $\mathrm{Stab}(v)$ is generated by the stabilisers C_i of X_i and by finitely many elements g_{ij} such that $g_{ij}X_i$ crosses X_j. We know that the C_i are commensurable with H, and that $C_i^{g_{ij}}$ is commensurable with C_j since $g_{ij}X_i$ crosses X_j. Thus $\mathrm{Stab}(v)$ commensurises H and so is a subgroup of $\mathrm{Comm}_G(H)$, which we denote by C in this proof for brevity. As H has small commensuriser, C has 2 ends. Hence $\mathrm{Stab}(v)$ has 2 ends. Let e denote an edge of T incident to v, so that $\mathrm{Stab}(e)$ is a subgroup of $\mathrm{Stab}(v)$ and hence of C. If $\mathrm{Stab}(e)$ were finite, the fact that T is minimal would imply that G splits over this finite subgroup and hence has more than one end. As G is one-ended, it follows that $\mathrm{Stab}(e)$ is infinite and so has finite index in $\mathrm{Stab}(v)$. As Γ is finite, it follows that there are only finitely many edges of T incident to v. Recall that part 3) of Corollary 4.16 tells us that each nontrivial almost invariant subset X enclosed by v is determined up to equivalence by the induced partition of the edges incident to v. It follows that the elements of E enclosed by v belong to finitely many equivalence classes.

The fact that Γ has finitely many V_0–vertices implies that the V_0–vertices of T lie in finitely many G–orbits. If a V_0–vertex w of T has stabiliser which is a subgroup of C, then the stabiliser of gw satisfies the same condition if and only if g commensurises H and so lies in C. As the C–orbit of w is finite, it follows that there are only finitely many V_0–vertices of T whose stabiliser is contained in C. As for v, each such V_0–vertex can enclose only finitely many equivalence classes of elements of E.

We conclude from the above discussion that if Γ is the regular neighbourhood of a finite set of nontrivial elements of $Q(H)$, then all the edge splittings of Γ are over two-ended subgroups commensurable with H, and the V_0–vertices of Γ can only enclose finitely many equivalence classes of elements of $Q(H)$.

Now suppose that $B(H)$ is infinite. We will describe how to pick a sequence $\{U_i\}_{i\geqslant 1}$ of elements of $Q(H)$ which represent distinct elements of $B(H)$. Having chosen $U_1,\ldots,U_k$, we will form their regular neighbourhood Γ_k. As the V_0–vertices of Γ_k can only enclose finitely many equivalence classes of elements of $Q(H)$, there is U_{k+1} in $Q(H)$ which is not enclosed by any V_0–vertex of Γ_k. This implies that when we form Γ_{k+1}, it is distinct from Γ_k. Lemma 5.19 implies that each edge splitting of Γ_k is enclosed by some V_0–vertex of Γ_{k+l}, for any $l\geqslant 1$. In particular, the edge splittings of Γ_k are compatible with those of Γ_{k+l}, for any $l\geqslant 1$. Consider all the edge splittings of $\Gamma_1,\ldots,\Gamma_k$ and choose one from each conjugacy class. Let Δ_k denote the graph of groups structure for G whose edge splittings are the chosen ones. Such a graph of groups exists by Theorem 2.35. It is trivial that Δ_{k+1} is a refinement of Δ_k. Further, our construction implies that Δ_k has no redundant vertices. As the edge groups of Δ_k are two-ended, the accessibility result of Theorem 7.11 applies and tells us that the sequence $\{\Delta_k\}$ must eventually stabilise, i.e. there is N such that $\Delta_N = \Delta_n$, for all $n \geqslant N$. It follows that the V_0–vertices of Γ_N enclose every U_i. But the U_i's represent infinitely many distinct elements of $B(H)$ which contradicts the preceding paragraph. This contradiction shows that $B(H)$ must be finite as required. □

Now we know that $B(H)$ is finite, we can prove the following result.

Proposition 7.16. — *Let G be a one-ended, finitely presented group, and let H be a two-ended subgroup with small commensuriser. Let Γ denote the regular neighbourhood of the collection $B(H)$. Thus Γ is the graph of groups Γ_N above. Then one of the following cases holds:*

(1) *All V_0–vertices of Γ are isolated. In this case, Γ has at most three V_0–vertices, and each has associated group which is commensurable with H.*

(2) *There is exactly one non-isolated V_0–vertex of Γ with associated group $\mathrm{Comm}_G(H)$, and a non-zero number of isolated V_0–vertices. The non-isolated V_0–vertex encloses every element of $Q(H)$. Further each isolated V_0–vertex is joined to the non-isolated V_0–vertex by a path of length 2, such that the single V_1–vertex on this path has valence 2.*

Proof. — Choose representatives $U_1, \dots, U_k$ of the nontrivial elements of $B(H)$, so that U_i has stabiliser H_i which is commensurable with H. We will use our usual notation from chapter 3. Thus E denotes the set of all translates of $U_1, \dots, U_k$ and their complements, and T denotes the universal covering G–tree of Γ.

Let K denote the intersection of the H_i's, so that each U_i is K–almost invariant, and consider the almost invariant subsets $K\backslash U_i$ of $K\backslash G$. Note that the coboundary $\delta(K\backslash U_i)$ is finite for each i. Let Λ denote the Cayley graph of G with respect to some finite generating set, and let A denote a finite connected subcomplex of the graph $K\backslash\Lambda$ which contains every edge of each of $\delta(K\backslash U_i)$ and which carries K. Consider the inverse image Z of A in Λ. Then Z is connected and K–finite. Let $X_1, \dots, X_m$ denote the K–infinite components of the complement of Z. Since the edges of $\delta(K\backslash U_i)$ are contained in A, we see that the edges of δU_i are contained in Z, and thus each U_i is a union of some of the X_j. Hence any K–almost invariant subset of G is equivalent to a union of some of the X_j's. In particular, if a nontrivial almost invariant subset over a subgroup commensurable with H is contained in one of the X_j's then it must be H–almost equal to X_j. It follows that X_j cannot cross any element of $B(H)$. Thus any element of $B(H)$ is represented by some union of the X_j's, and those elements of E which are equivalent to some X_j must be isolated in E.

If $m > 3$, we will show that we have case 2) of the proposition, and that the number of isolated V_0–vertices of Γ equals m. Recall that the CCC's of $\overline{E}$ correspond to the V_0–vertices of T. Consider all unions of k of the X_j's, for all k such that $2 \leqslant k \leqslant (m-2)$. It is easy to check that all elements of E equivalent to such unions lie in the same CCC. For example, $X_1 \cup X_2$ and $X_1 \cup X_3$ cross each other. Thus, if we consider only those CCC's which enclose representatives of elements of $B(H)$, there is exactly one non-isolated such CCC, say v, and there are exactly m isolated such CCC's, say $v_1, \ldots, v_m$, where v_j encloses X_j. Clearly, v is invariant under $\mathrm{Comm}_G(H)$, and as its stabiliser must commensurise H by the proof of Theorem 7.14, it follows that $\mathrm{Stab}(v) = \mathrm{Comm}_G(H)$. Now Lemma 4.7 tells us that if a vertex of a G–tree encloses two almost invariant subsets of G, then it also encloses their intersection. Thus the fact that $X_1 = (X_1 \cup X_2) \cap (X_1 \cup X_3)$ implies that v encloses X_1. As Proposition 5.2 tells us that T is a minimal G–tree, we can apply part 1) of Lemma 4.9 which tells us that all the vertices on the interior of the path λ_1 in T joining v_1 to v have valence 2, and all these vertices enclose X_1. Lemma 5.4 now implies that each V_0–vertex on the interior of λ_1 is isolated. Recall that no two distinct elements of E are equivalent. This implies that X_1 can be enclosed by only one isolated V_0–vertex of T. It follows that λ_1 has no interior V_0–vertices, so that λ_1 has length 2. Similarly each v_i is joined to v by a path of length 2, such that the single V_1–vertex on this path has valence 2. This completes the proof that Γ has all the properties in case 2) of the proposition.

If $m \leqslant 3$, every element of $Q(H)$ is equivalent to some X_j or its complement. Thus we have m isolated CCC's, say $v_1, \ldots, v_m$, where v_j encloses X_j, and so Γ has m isolated V_0–vertices as in case 1) of the proposition. □

The first part of the proof of the above proposition shows that each X_j contains only one coend of H in G. Thus we have the following proposition which answers a question of Bowditch [**8**] in the finitely presented case:

Proposition 7.17. — *Suppose that G is one-ended and finitely presented and that H is a two-ended subgroup of G with small commensuriser. Then the number of coends of H in G is finite.*

Note that in either case in Proposition 7.16, there are isolated elements of E. Such elements determine splittings of G which have intersection number zero with every element of $Q(H)$. We will call such a splitting H*-canonical.* Thus we have the following result.

Corollary 7.18. — *Suppose that G is a one-ended, finitely presented group and assume that G is not of finite-by-Fuchsian type. Let H be a two-ended subgroup of G and suppose that H has small commensuriser in G. If G possesses a nontrivial almost invariant subset over a subgroup commensurable with H, then G possesses a H–canonical splitting over a subgroup commensurable with H.*

CHAPTER 8

COENDS WHEN THE COMMENSURISER IS LARGE

Recall that $Q(H)$ denotes the collection of all almost invariant subsets of G which are over subgroups of G commensurable with H, and $B(H)$ denotes the collection of equivalence classes of elements of $Q(H)$. In Proposition 7.16, we described the structure of the regular neighbourhood Γ of all nontrivial elements of $B(H)$, in the case when H has small commensuriser in G. In this chapter, we consider a one-ended finitely generated group G, and a two-ended subgroup H such that H has large commensuriser in G, i.e. H has infinite index in $\mathrm{Comm}_G(H)$. Again we want to study the structure of the regular neighbourhood Γ of all nontrivial elements of $B(H)$. However $B(H)$ may be infinite, so that it is not clear that such a regular neighbourhood exists. We will show that it does exist in chapter 10. A key point in the argument is that $B(H)$ possesses certain algebraic finiteness properties. Recall from the previous chapter that $Q(H)$ and $B(H)$ are Boolean algebras. Also $Q(H)$ is invariant under the action by left multiplication of $\mathrm{Comm}_G(H)$, and this action induces an action on $B(H)$. Thus $B(H)$ is a Boolean algebra with a natural action of $\mathrm{Comm}_G(H)$.

We will use some arguments from [**13**]. As before, the results extend to the case where H is virtually polycyclic (VPC) of any length but we will first discuss the case where H is $VPC1$, i.e. H is two-ended. We start by stating a result which is a special case of a theorem of Kropholler and Roller in [**30**] (see also [**21**]), and we sketch an argument for this special case which uses the proof of Lemma B.2.13.

Proposition 8.1. — *Let G be a finitely generated group and H a two-ended subgroup with large commensuriser. Then the number of coends of H in G is 1, 2 or infinity. The number of coends is 2 if and only if G is virtually $\mathbb{Z} \times \mathbb{Z}$.*

Proof. — If the number of coends of H in G is greater than 1, Lemma 2.40 shows that there is a subgroup K of H and a nontrivial K–almost invariant subset X of G.

If K is finite, then G has more than one end. If G were two-ended, H would be of finite index in G and so could not have large commensuriser. Thus G has infinitely many ends, and it follows that H has infinitely many coends in G.

If K is infinite, it must have finite index in H, so that $\mathrm{Comm}_G(H) = \mathrm{Comm}_G(K)$. Consider the translates of X by elements of $\mathrm{Comm}_G(H)$. If infinitely many of these are equivalent to X, then the proof of Lemma B.2.13 shows that there is a subgroup K_1 of finite index in G which contains a subgroup H_1 commensurable with H such that H_1 is normal in K_1 and $H_1\backslash K_1$ has two ends. Thus H has two coends in G, and as H is virtually infinite cyclic, G is virtually $\mathbb{Z} \times \mathbb{Z}$. If only finitely many translates of X by $\mathrm{Comm}_G(H)$ are equivalent to X, then X has infinitely many distinct such translates and it follows that H has infinitely many coends in G.

We complete the proof of the proposition by observing that if G is virtually $\mathbb{Z} \times \mathbb{Z}$, then H has 2 coends in G. □

Now we can prove the following finiteness result for the Boolean algebra $B(H)$.

Theorem 8.2. — *If G is a one-ended, finitely presented group and H is a two-ended subgroup of G, then $B(H)$ is finitely generated over* $\mathrm{Comm}_G(H)$.

Proof. — Theorem 7.14 tells us that if H has small commensuriser, then $B(H)$ is finite. Thus the result is trivial in this case. So we will assume that H has infinite index in $\mathrm{Comm}_G(H)$, and that $B(H)$ is infinite. Note that all the crossings between elements of $B(H)$ must be weak, by Corollary 7.10. The accessibility result of Theorem 7.11 tells us that there is a finite graph of groups decomposition $\mathcal{G}$ of G with all edge groups commensurable with H, such that $\mathcal{G}$ cannot be properly refined using such splittings. An alternative way of expressing this condition is to say that if G possesses a splitting over a two-ended subgroup commensurable with H which has intersection number zero with the edge splittings of $\mathcal{G}$, then this splitting is conjugate to one of these edge splittings. Let $X_1, \ldots, X_n$ denote almost invariant subsets of G associated to the edge splittings of $\mathcal{G}$, and let H_i denote the stabiliser of X_i. Let $A(H)$ denote the subalgebra of $B(H)$ generated over $\mathrm{Comm}_G(H)$ by the equivalence classes of the X_i's. We will show that $B(H) = A(H)$.

Let E denote the collection of translates of the X_i's by elements of $\mathrm{Comm}_G(H)$. Let Y be an element of $Q(H)$, such that Y is a nontrivial almost invariant subset of G over a subgroup K commensurable with H. We will show that Y crosses only finitely many elements of E. The intersection number $i(H_i\backslash X_i, K\backslash Y)$ is finite and is the number of double cosets KgH_i such that gX_i crosses Y. If gX_i crosses Y, it must do so weakly by Corollary 7.10. Now Proposition 7.3 tells us that the stabilisers of gX_i and Y are commensurable, so that H_i^g and K are commensurable. As H_i and K are commensurable, it follows that g commensurises H_i, and hence also commensurises K. If we let L_i denote $K \cap H_i$, then $g^{-1}Kg$ can be expressed as the union of cosets

$g_j(g^{-1}Kg \cap L_i)$, for $1 \leqslant j \leqslant n$. Hence

$$KgH_i = g(g^{-1}Kg)H_i = g\left(\cup_{j=1}^n g_j(g^{-1}Kg \cap L_i)\right)H_i = g\left(\cup_{j=1}^n g_jH_i\right) = \cup_{j=1}^n gg_jH_i,$$

so that KgH_i is the union of finitely many cosets gH_i of H_i. It follows that there are only finitely many translates of the X_i's which cross Y. We conclude that Y crosses only finitely many of the translates of the X_i's by G and that these must all lie in E.

If Y crosses no elements of E, then Y is enclosed by some vertex v of the universal covering G–tree T of $\mathcal{G}$. In particular, the stabiliser K of Y is a subgroup of $\mathrm{Stab}(v)$, by part 2) of Lemma 4.9. Suppose that $Y \cap \mathrm{Stab}(v)$ is a nontrivial K–almost invariant subset of $\mathrm{Stab}(v)$. Then the proof of the main theorem of [**13**] produces another nontrivial almost invariant subset U of G such that U is over a subgroup of $\mathrm{Stab}(v)$ which is commensurable with K, and $U \cap \mathrm{Stab}(v)$ is also a nontrivial almost invariant subset of $\mathrm{Stab}(v)$, and U has self-intersection number zero. As U has self-intersection number zero, Theorem B.2.8 implies that U is equivalent to an almost invariant set W which is associated to a splitting σ of G. The set U was constructed by taking successively intersections of $Y^{(*)}$ and $cY^{(*)}$ where $c \in \mathrm{Comm}_G(H) \cap \mathrm{Stab}(v)$, so that Lemma 4.7 shows that it is enclosed by v. Hence W is also enclosed by v. Thus σ has intersection number zero with the edge splittings of $\mathcal{G}$. Our choice of $\mathcal{G}$ now implies that σ must be conjugate to one of the edge splittings of $\mathcal{G}$. In particular, σ does not split $\mathrm{Stab}(v)$, so that $U \cap \mathrm{Stab}(v)$ must be a trivial almost invariant subset of $\mathrm{Stab}(v)$, which is a contradiction. This contradiction shows that if Y crosses no elements of E, then $Y \cap \mathrm{Stab}(v)$ must be a trivial almost invariant subset of $\mathrm{Stab}(v)$.

Now choose v as the basepoint of T, and define $\varphi : G \to V(T)$ by the formula $\varphi(g) = gv$. Recall from Lemma 4.14 that Y is equivalent to the set $B(Y) = \varphi^{-1}(\Sigma_v(Y)) \cup (Y \cap \varphi^{-1}(v))$. As $\varphi(e) = v$, we see that $\varphi^{-1}(v) = \mathrm{Stab}(v)$. Thus $Y \cap \varphi^{-1}(v)$ is a trivial almost invariant subset of $\mathrm{Stab}(v)$. It follows that Y is equivalent to $\varphi^{-1}(\Sigma_v(Y))$. This set is the union of some of the Z_s^*, for edges s incident to v and oriented towards v. As the Z_s^* and their coboundaries are disjoint, the coboundary of this union equals the union of the δZ_s^*. As δY is H–finite, it follows that each δZ_s^* is also H–finite. Let S denote the stabiliser of s, and recall that S must be a conjugate of some H_i. As S is infinite, so is δZ_s^*. The fact that δZ_s^* is H–finite now implies that δZ_s^* must be stabilised by an infinite subgroup of H. It follows that S and H are commensurable, so that the equivalence class of Z_s^* lies in $A(H)$. It also follows that Y is equivalent to the union of a finite number of Z_s^*. Thus the equivalence class of Y lies in the subalgebra $A(H)$ of $B(H)$.

Next suppose that Y crosses some elements of E and let v be a vertex of T some of whose incident edges have associated edge splittings crossed by Y. If $s_1, \ldots, s_k$ are these edges, we denote the almost invariant subset of G associated to s_i by Z_i for obvious typographical reasons. Then, we can express Y as the union of $Y \cap Z_i^*$, $1 \leqslant i \leqslant k$, and of $W = Y \cap \left(\bigcap_{i=1}^k Z_i\right)$. Now Lemma 4.7 shows that W is enclosed by v, so that the equivalence class of W lies in the subalgebra $A(H)$ of $B(H)$, by the preceding paragraph. Moreover, each of $Y \cap Z_i^*$ crosses a smaller number of elements of E than does Y and so by induction we see that the equivalence class of each $Y \cap Z_i^*$ also lies in the subalgebra $A(H)$ of $B(H)$. This now implies that the equivalence class of Y itself lies in $A(H)$ and completes the proof that $B(H)$ is finitely generated over $\mathrm{Comm}_G(H)$. □

In the preceding proof, we referred several times to the arguments of Dunwoody and Roller in [**13**]. It will also be convenient to state a result which combines one of the main results of that paper with results from [**15**].

Theorem 8.3. — *Let G be a one-ended, finitely generated group which does not split over VPC subgroups of length $< n$, and let H be a $VPCn$ subgroup with large commensuriser, such that $e(G, H) \geqslant 2$. Then G splits over some subgroup commensurable with H.*

Note that the proof of Theorem 8.2 shows that $B(H)$ is generated by almost invariant sets associated to splittings of G over subgroups commensurable with H, so that Theorem 8.3 follows. However, this is not a new proof of Theorem 8.3 as we used the main result of [**13**] in our proof of Theorem 8.2.

We call a nontrivial element X of $Q(H)$ *special* if $X \cap \mathrm{Comm}_G(H)$ is H–finite. A splitting of G over a subgroup commensurable with H will be called *special* if one of the associated almost invariant subsets of G is a special element of $Q(H)$. Recall that a splitting of G over a subgroup commensurable with H is H–canonical if it has zero intersection number with every element of $Q(H)$.

Proposition 8.4. — *Let G be a one-ended, finitely presented group which is not virtually $\mathbb{Z} \times \mathbb{Z}$, with a two-ended subgroup H which has large commensuriser. If $Q(H)$ has a nontrivial special element, then G has an H–canonical special splitting over a subgroup commensurable with H.*

Remark 8.5. — If H has small commensuriser, then every element of $Q(H)$ is special, so the result follows immediately from Corollary 7.18.

Proof. — The proof uses details from the arguments of [**13**] with slight improvements from [**15**]. Note that as H has large commensuriser and G is not virtually $\mathbb{Z} \times \mathbb{Z}$, Corollary 7.10 tells us that an element of $Q(H)$ cannot cross strongly any nontrivial almost invariant set over a two-ended subgroup.

The proof in [**13**] shows that if X is a special element of $Q(H)$, then there is a splitting of G over a subgroup commensurable with H such that one of the almost invariant sets associated with the splitting is contained in X (and thus special). We recall some of the argument. By modifying X by a H–finite set, we can arrange that X does not intersect $\mathrm{Comm}_G(H)$. In this case, Lemma 5 of [**13**] tells us that if $c \in \mathrm{Comm}_G(H)$ and X and cX are nested, then either $cX = X$ or $X \cap cX = \varnothing$. Now [**13**] produces a new element Y of $Q(H)$ which is a finite intersection of some

$c_i X$, $c_i \in \mathrm{Comm}_G(H)$, such that the translates of Y by all elements of G are almost nested. This is enough to produce a splitting of G with associated almost invariant set equivalent to Y. There is an improvement in [**15**] (see section 3, last paragraph on page 622), which shows there is a subset Z of Y which is equivalent to Y such that the translates of Z are actually nested. This implies that Z is associated to a splitting of G over a subgroup commensurable with H. Thus, if we start with an element of $Q(H)$ which is disjoint from $\mathrm{Comm}_G(H)$, there is a subset which is also an element of $Q(H)$ and is associated to a splitting of G. Of course, this new element of $Q(H)$ is also disjoint from $\mathrm{Comm}_G(H)$.

Let X_1 denote an element of $Q(H)$ which is disjoint from $\mathrm{Comm}_G(H)$, and is associated to a splitting of G. If X_1 contains an element Y_1 of $Q(H)$ which is not equivalent to X_1, the preceding paragraph yields a subset X_2 of Y_1, which is also an element of $Q(H)$ and is associated to a splitting of G. We repeat this process to obtain a descending sequence of inequivalent elements X_i of $Q(H)$, each of which is disjoint from $\mathrm{Comm}_G(H)$, and is associated to a splitting σ_i of G.

Suppose that this sequence stops after finitely many steps. Then we will have found an element X of $Q(H)$ which is disjoint from $\mathrm{Comm}_G(H)$ and is associated to a splitting of G, such that any subset of X which lies in $Q(H)$ is equivalent to X. Thus X is a minimal element in $B(H)$. Any such minimal element cannot be crossed by any element in $Q(H)$, so it determines a H–canonical splitting of G, as required.

If the sequence does not stop after finitely many steps, we will obtain an infinite sequence X_i. We will show that this leads to a contradiction. Let H_i denote the stabiliser of X_i. We claim $H_2 \subset H_1$. Let c be an element of H_2. As $cX_2 = X_2$, we see that $cX_1 \cap X_1$ is non-empty. As X_1 is associated to a splitting, all its translates are nested. In particular, as c lies in $\mathrm{Comm}_G(H_1)$, Lemma 5 of [**13**] tells us that we have $cX_1 = X_1$ or $cX_1 \cap X_1 = \varnothing$. As $cX_1 \cap X_1$ is non-empty, we must have $cX_1 = X_1$, so that c lies in H_1. It follows that $H_2 \subset H_1$. Next we claim that X_1 does not cross any translate of X_2, so that the splittings σ_1 and σ_2 have intersection number zero. For suppose that X_1 and gX_2 cross each other. We pointed out at the start of this proof that they must cross weakly because X_1 lies in $Q(H)$, so that Proposition 7.3 tells us that their stabilisers must be commensurable. As X_1 and X_2 have commensurable stabiliser, it follows that g lies in $\mathrm{Comm}_G(H)$. As before this means that $gX_1 = X_1$ or $gX_1 \cap X_1 = \varnothing$. But the fact that X_1 and gX_2 cross implies that $X_1 \cap gX_2$ is not empty, so that $X_1 \cap gX_1$ is not empty. It follows that $gX_1 = X_1$, so that $gX_2 \subset X_1$, which contradicts our assumption that X_1 and gX_2 cross. Similarly, $H_{i+1} \subset H_i$, for all $i \geqslant 1$, and X_i does not cross any translate of X_j, for any i and j. Thus the splittings σ_i are all compatible. Hence we obtain an infinite sequence of graphs of groups decompositions of G each refining the previous one. As the σ_i's are distinct, this contradicts the accessibility result in Theorem 7.11. This contradiction shows that the sequence of X_i's cannot be infinite, so that we obtain a minimal X as in the preceding paragraph, and hence obtain a H–canonical splitting of G. □

The argument above shows that even though the number of coends of H in G is infinite, if $Q(H)$ has a special element, then there are elements of $Q(H)$ which contain a finite number of coends of H with the complementary set containing all the other coends. Of course, there may not be any special elements in $Q(H)$.

Now we will consider the following construction. Let G be a one-ended, finitely presented group, pick one representative for each element of $B(H)$ and let E denote the set of all translates of this collection and their complements. We will consider the CCC's of $\overline{E}$. As usual, these form a pretree, but we will not be able to prove that this is discrete until chapter 10.

Proposition 8.6. — *Let G be a one-ended, finitely presented group and suppose that G is not virtually $\mathbb{Z} \times \mathbb{Z}$. Let H be a two-ended subgroup of G such that H has large commensuriser and there are nontrivial H–almost invariant subsets of G. Consider the CCC's of $\overline{E}$ which consist of elements of $Q(H)$. There is exactly one such CCC which is infinite, and there are finitely many (possibly zero) other CCC's all of which are isolated. The infinite CCC has stabiliser equal to* $\mathrm{Comm}_G(H)$, *and encloses every element of $Q(H)$.*

Remark 8.7. — In the case when G is virtually $\mathbb{Z} \times \mathbb{Z}$, the CCC's of $\overline{E}$ which consist of elements of $Q(H)$ form a single CCC. This is discussed at the start of the proof of Theorem 10.1.

Proof. — Theorem 8.2 shows that there are a finite number of compatible splittings $\sigma_1, \ldots \sigma_n$ of G each over a subgroup of G commensurable with H, such that the equivalence classes of the associated almost invariant subsets $\{X_i, X_i^* : 1 \leqslant i \leqslant n\}$ generate $B(H)$ over $\mathrm{Comm}_G(H)$. In fact the proof shows more. It shows that any element of $B(H)$ is represented by some finite union of translates by $\mathrm{Comm}_G(H)$ of the X_i's. As G is not virtually $\mathbb{Z} \times \mathbb{Z}$, Proposition 8.1 tells us that H has infinitely many coends in G. As the translates of X_i and X_i^* are nested, the proof of part 2) of Proposition 7.16 shows that the collection of all finite unions of more than one of these sets forms a single CCC. It further shows that this CCC also encloses the X_i's themselves. It follows that this CCC has stabiliser equal to $\mathrm{Comm}_G(H)$. Finally, there are isolated CCC's corresponding to those X_i's which are isolated (if any). □

CHAPTER 9

CANONICAL DECOMPOSITIONS OVER TWO-ENDED GROUPS WHEN COMMENSURISERS ARE SMALL

In this chapter and the next, we will find canonical decompositions of a one-ended, finitely presented group G which are analogous to the topological JSJ-decomposition of an atoroidal 3–manifold. Our approach is similar to an unpublished approach of Scott [**39**] for proving the classical results on the JSJ-decomposition. The idea is to enclose all nontrivial almost invariant subsets of G which are over two-ended subgroups. However, instead of enclosing all at once, we form regular neighbourhoods of larger and larger finite families and use accessibility to show that the sequence obtained in this way must stabilise. This is where we use finite presentability. The result is a regular neighbourhood of all equivalence classes of nontrivial almost invariant subsets of G over two-ended subgroups. (See Definition 6.4.) Since we enclose all nontrivial almost invariant sets over two-ended groups, the decompositions that we obtain are unique and are invariant under automorphisms of the group.

The preceding two chapters contain the crucial pattern for obtaining our canonical decompositions. We want to form a regular neighbourhood of an infinite family of nontrivial almost invariant subsets of G. The first step is to show that the cross-connected components are of two types, those which contain only strong crossings and those which contain only weak crossings. The structure of the strong crossing components is handled by the techniques of Bowditch [**8**] and of Dunwoody and Swenson [**15**]. If commensurisers are small, the structure of weak crossing components is easy to describe using regular neighbourhoods as in chapter 7. If the commensuriser of a subgroup H of G is large, we use the fact that the Boolean algebra $B(H)$ is finitely generated over $\mathrm{Comm}_G(H)$, as in chapter 8. This used the arguments of Dunwoody and Roller [**13**] for a special case of the annulus theorem. To obtain our canonical decompositions, the only remaining difficulty is to show that the pretree which we construct from the cross-connected components is discrete. This is clear in the case when all commensurisers are small, and is proved in general using again the fact that $B(H)$ is finitely generated over $\mathrm{Comm}_G(H)$. This is what we will do in this and the next chapter. We will use the same strategy in the more general cases which we consider later.

Let Γ be a minimal graph of groups decomposition of G. Recall from Definition 7.6 that a vertex v of Γ is of finite-by-Fuchsian type if $G(v)$ is a finite-by-Fuchsian group, where the Fuchsian group is not finite nor two-ended, and there is exactly one edge of Γ which is incident to v for each peripheral subgroup K of $G(v)$ and this edge carries K. We will also need the following definition.

Definition 9.1. — Let Γ be a minimal graph of groups decomposition of a group G. A vertex v of Γ is *simple*, if whenever X is a nontrivial almost invariant subset of G over a two-ended subgroup such that X is enclosed by v, then X is associated to an edge splitting of Γ.

Note that if v is simple, then part 1) of Lemma 4.9 implies that X is associated to an edge splitting for an edge of Γ which is incident to v. Also note that it is not possible for v to be both simple and of finite-by-Fuchsian type. For if a Fuchsian group is not finite nor two-ended, the corresponding 2-orbifold has non-peripheral loops representing elements of infinite order, and any such yields a nontrivial almost invariant subset of G over a two-ended subgroup which is not associated to an edge splitting of Γ.

The above definition is the analogue of the topological fact that if $V(M)$ denotes the characteristic submanifold of a 3–manifold M, then any essential annulus in M which lies in $M - V(M)$ is homotopic to a covering of an annulus in the characteristic family $\mathcal{T}$. We could have defined simple differently by insisting that the condition applies only to almost invariant subsets which are associated to a splitting. Call this condition "simple for splittings". This would correspond to a topological definition of simple which applies only to embedded annuli. Unfortunately, the word simple is used in both these senses in the literature of 3–manifold theory. Note that a vertex v of a graph of groups Γ can be of finite-by-Fuchsian type and be simple for splittings. For example, this occurs when $G(v)$ is the fundamental group of a thrice punctured sphere.

If one applies Definition 9.1 to splittings of G, one sees that if v is simple, then Γ cannot be properly refined by splitting at v using a splitting of G over a two-ended subgroup.

Finally, note that v need not be simple even if $G(v)$ is two-ended, despite the fact that such a group does not admit any splitting over a two-ended subgroup. Again this is clear from consideration of topological examples. Consider a solid torus component Σ of $V(M)$ such that Σ has at least four frontier annuli. Then Σ will contain essential embedded annuli which are not homotopic to covers of the frontier annuli. This example shows that v need not even be simple for splittings.

In this chapter, we will assume that whenever H is a two-ended subgroup of G and $e(G, H) \geqslant 2$, then H has small commensuriser. The class of groups for which this holds includes all word hyperbolic groups. For any group G, we will let $\mathcal{F}_1$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a two-ended subgroup. The subscript 1 is because a two-ended group is $VPC1$. Now we can state the main result of this chapter which is our version of the JSJ-decomposition for this class of groups.

Theorem 9.2. — *Let G be a one-ended, finitely presented group such that whenever H is a two-ended subgroup and $e(G, H) \geqslant 2$, then H has small commensuriser. Let $\mathcal{F}_1$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a two-ended subgroup.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$. Further each V_0-vertex v of $\Gamma(\mathcal{F}_1 : G)$ satisfies one of the following conditions:

(1) v is isolated.

(2) $G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some two-ended subgroup H, such that $e(G, H) \geqslant 2$.

(3) v is of finite-by-Fuchsian type.

$\Gamma(\mathcal{F}_1 : G)$ consists of a single vertex if and only if either $\mathcal{F}_1$ is empty or G is of finite-by-Fuchsian type.

Remark 9.3. — In cases 1) and 2), $G(v)$ is two-ended. If v is of type 2) but not isolated, we will say that v is *of small commensuriser type.*

Proof. — We start by picking a representative for each element of $\mathcal{F}_1$, subject to the condition that if A and B are elements of $\mathcal{F}_1$ such that $B = gA$, for some g in G, then the representatives X and Y chosen for A and B must satisfy $Y = gX$. As $\mathcal{F}_1$ is countable, it can be expressed as the union of an ascending sequence of finite subsets $\mathcal{E}_i$, for $i \geqslant 1$. As each $\mathcal{E}_i$ is finite, it has a regular neighbourhood $\Gamma_i = \Gamma(\mathcal{E}_i : G)$. We will choose the sequence $\mathcal{E}_i$ carefully, and then show that the sequence Γ_i must stabilise eventually. The resulting graph of groups structure Γ for G will be the required regular neighbourhood $\Gamma = \Gamma(\mathcal{F}_1 : G)$.

We will start with some choice of the $\mathcal{E}_i$'s and then will modify this sequence inductively. We will continue to denote the modified sets by $\mathcal{E}_i$. We modify $\mathcal{E}_1$ as follows. Let E_1 denote the collection of all the translates of the chosen representatives of $\mathcal{E}_1$. Consider a H–almost invariant subset X of G which lies in E_1. If X crosses weakly some element of $\mathcal{F}_1$, Lemma 7.3 tells us that this element lies in $B(H)$. Theorem 7.14 tells us that $B(H)$ is finite. Thus, by enlarging $\mathcal{E}_1$, we can arrange that whenever E_1 contains a H–almost invariant subset of G which crosses weakly some element of $\mathcal{F}_1$, then $\mathcal{E}_1$ contains each element of $B(H)$. This is the final version of $\mathcal{E}_1$.

The sequence $\mathcal{E}_i$ is constructed inductively starting from $\mathcal{E}_1$. Having constructed $\mathcal{E}_i$, we first check whether $\mathcal{E}_i = \mathcal{F}_1$. If it does, then $\Gamma(\mathcal{E}_i : G)$ is the required regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$. Otherwise, we need to construct $\mathcal{E}_{i+1}$. As the new $\mathcal{E}_i$

is still finite, there is some index $j > i$ such that $\mathcal{E}_j$ properly contains the new $\mathcal{E}_i$. By replacing our sequence of subsets of $\mathcal{F}_1$ by a subsequence, we can suppose that $j = i+1$. Now we enlarge $\mathcal{E}_{i+1}$ in the same way in which we enlarged $\mathcal{E}_1$.

In order to show that the sequence Γ_i stabilises, we first need to describe the V_0–vertices of Γ_i. Recall that Γ_i can be constructed using the collection E_i of all the translates of the chosen representatives of $\mathcal{E}_i$. Proposition 7.5 shows that the non-isolated CCC's of $\overline{E_i}$ are of two types. In a given CCC, the crossings are either all strong or all weak. In the strong crossing case, Theorem 7.8 shows that the corresponding V_0–vertex of Γ_i is of finite-by-Fuchsian type. Suppose that E_i has an element X which is H–almost invariant and crosses some other element of E_i weakly. Then Proposition 7.5 shows that the CCC of $\overline{E_i}$ which contains $\overline{X}$ contains only elements whose stabiliser is commensurable with H. Now the proof of Proposition 7.16 tells us that the CCC of $\overline{E_i}$ which contains $\overline{X}$ has stabiliser equal to $\mathrm{Comm}_G(H)$, and encloses every element of E_i whose stabiliser is commensurable with H. As $\mathrm{Comm}_G(H)$ contains H with finite index, it is two-ended. It follows that every V_0–vertex group of Γ_i is either two-ended or of finite-by-Fuchsian type, and hence that every edge group of Γ_i is two-ended. Lemma 5.19 shows that the edge splittings of Γ_i are compatible with those of Γ_j for every $j > i$. Now we proceed as at the end of the proof of Theorem 7.14. Consider all the edge splittings of $\Gamma_1, \dots, \Gamma_k$ and choose one from each conjugacy class. Let Δ_k denote the graph of groups structure for G whose edge splittings are the chosen ones. Such a graph of groups exists by Theorem 2.35. It is trivial that Δ_{k+1} is a refinement of Δ_k. Further, our construction implies that Δ_k has no redundant vertices. As the edge groups of Δ_k are two-ended, the accessibility result of Theorem 7.11 applies and tells us that the sequence Δ_k must eventually stabilise, i.e. there is N such that $\Delta_N = \Delta_n$, for all $n \geqslant N$. It follows that the sequence Γ_i must stabilise, as required. We call this final graph of groups Γ.

By construction, Γ is the required regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$. Each V_0–vertex group of Γ is isolated, of small commensuriser type or of finite-by-Fuchsian type because this holds for each Γ_i. Finally, Γ will consist of a single vertex if and only if either $\mathcal{F}_1$ is empty or the representatives of $\mathcal{F}_1$ lie in a single CCC. In the second case, the vertex group of Γ must be two-ended or of finite-by-Fuchsian type. As the vertex group is G and we assumed that G is one-ended, it follows that $\Gamma(\mathcal{F}_1 : G)$ consists of a single vertex if and only if either $\mathcal{F}_1$ is empty or G is of finite-by-Fuchsian type. □

In order to understand $\Gamma(\mathcal{F}_1 : G)$ in more detail, our next result lists some properties which follow almost immediately from the above theorem and the properties of an algebraic regular neighbourhood.

Theorem 9.4. — *Let G be a one-ended, finitely presented group such that whenever H is a two-ended subgroup and $e(G, H) \geqslant 2$, then H has small commensuriser. Let $\mathcal{F}_1$*

denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a two-ended subgroup.

Then the regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$ is a minimal bipartite graph of groups decomposition of G with the following properties:

(1) *Each V_0–vertex v of $\Gamma(\mathcal{F}_1 : G)$ satisfies one of the following conditions:*

(a) *v is isolated.*

(b) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some two-ended subgroup H, such that $e(G, H) \geqslant 2$.*

(c) *v is of finite-by-Fuchsian type.*

(2) *the edge groups of $\Gamma(\mathcal{F}_1 : G)$ are two-ended.*

(3) *any element of $\mathcal{F}_1$ is enclosed by some V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$, and each V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$ encloses such a subset of G. In particular, any splitting of G over a two-ended subgroup is enclosed by some V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$.*

(4) *if X is a nontrivial almost invariant subset of G over a finitely generated subgroup H, and if X does not cross any element of $\mathcal{F}_1$, then X is enclosed by a V_1–vertex of $\Gamma(\mathcal{F}_1 : G)$.*

(5) *if X is a H–almost invariant subset of G associated to a splitting of G over H, and if X does not cross any element of $\mathcal{F}_1$, then X is enclosed by a V_1–vertex of $\Gamma(\mathcal{F}_1 : G)$.*

(6) *the V_1–vertices of $\Gamma(\mathcal{F}_1 : G)$ are simple.*

(7) *If Γ_1 and Γ_2 are minimal bipartite graphs of groups structures for G which satisfy conditions 3) and 5) above, then they are isomorphic provided that there is a one-to-one correspondence between their isolated V_0–vertices, and that any non-isolated V_0–vertex of Γ_1 or Γ_2 encloses some non-isolated element of $\mathcal{F}_1$.*

(8) *The graph of groups $\Gamma(\mathcal{F}_1 : G)$ is invariant under the automorphisms of G.*

(9) *The edge splittings of $\Gamma(\mathcal{F}_1 : G)$ are precisely the canonical splittings of G over two-ended subgroups. Hence if G is not of finite-by-Fuchsian type and $\mathcal{F}_1$ is non-empty, then G has a canonical splitting over a two-ended subgroup.*

Proof. — Part 1) holds by Theorem 9.2. Part 2) holds because every edge of $\Gamma(\mathcal{F}_1 : G)$ has one end at a V_0–vertex, and each such vertex group is two-ended or of finite-by-Fuchsian type.

By construction, every element of $\mathcal{F}_1$ is enclosed by some V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$, and every nontrivial almost invariant subset of G over a two-ended subgroup is equivalent to some element of $\mathcal{F}_1$. Thus part 3) follows at once. Parts 4) and 5) follow from Proposition 5.7.

To see that part 6) holds, suppose that X is a nontrivial almost invariant subset of G over a two-ended subgroup which is enclosed by some V_1–vertex v of $\Gamma(\mathcal{F}_1 : G)$. Part 3) tells us that X is enclosed by some V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$. Now it follows from part 1) of Lemma 4.9 that X is associated to an edge splitting of $\Gamma(\mathcal{F}_1 : G)$. This implies that v is simple, so that every V_1–vertex of $\Gamma(\mathcal{F}_1 : G)$ is simple as required.

Part 7) is exactly the uniqueness result for regular neighbourhoods stated in Theorem 6.7. For conditions 3) and 5) of Theorem 9.4 ensure that Γ_1 and Γ_2 each satisfy conditions 1) and 2) of our definition of a regular neighbourhood, Definition 6.1. And the two provisos in the statement of part 7) ensure that Γ_1 and Γ_2 each satisfy conditions 4) and 5) of Definition 6.1.

For part 8), consider any automorphism α of G. Then α induces a natural G–invariant action on $\mathcal{F}_1$. If we denote by P the pretree of CCC's of this collection, where we choose one representative for each element of $\mathcal{F}_1$ as in Definition 6.4, then α induces a G–invariant automorphism of P and thus defines a simplicial automorphism of the tree T. Thus there is an induced automorphism of $\Gamma(\mathcal{F}_1 : G)$. (See [**2**] for a discussion of automorphisms of graphs of groups.)

To prove part 9), we start by observing that it is clear that every edge splitting of $\Gamma(\mathcal{F}_1 : G)$ is canonical, i.e. it has zero intersection number with any element of $\mathcal{F}_1$, because any element of $\mathcal{F}_1$ is enclosed by some vertex of $\Gamma(\mathcal{F}_1 : G)$. It remains to show that these are the only canonical splittings of G over two-ended subgroups. Let σ denote a canonical splitting of G over a two-ended subgroup H, and let X denote the associated H–almost invariant subset of G. As X has intersection number zero with every element of $\mathcal{F}_1$, Proposition 5.7 implies that it is enclosed by some V_1–vertex of T, the universal covering G–tree of $\Gamma(\mathcal{F}_1 : G)$. But X is also enclosed by a V_0–vertex of T. Now part 1) of Lemma 4.9 implies that σ is conjugate to an edge splitting of $\Gamma(\mathcal{F}_1 : G)$ as required. Finally, if G is not of finite-by-Fuchsian type and $\mathcal{F}_1$ is non-empty, Theorem 9.2 implies that $\Gamma(\mathcal{F}_1 : G)$ does not consist of a single vertex. Thus $\Gamma(\mathcal{F}_1 : G)$ has at least one edge and so G has a canonical splitting over a two-ended subgroup. □

At this point we need to discuss further a point about the topological JSJ-decomposition which we mentioned in chapter 1. Recall that the frontier of the characteristic submanifold $V(M)$ of a 3–manifold M is not quite the same as the canonical family $\mathcal{T}$ of annuli and tori in M. Some of the components of $\mathcal{T}$ may appear twice in the frontier of $V(M)$. This means that we can obtain two slightly different graphs of groups structures for $\pi_1(M)$, one graph being dual to $\mathcal{T}$ and the other dual to the frontier of $V(M)$. The algebraic decomposition $\Gamma(\mathcal{F}_1 : G)$ which we obtained in Theorem 9.2 is closer to the second case. However, there is a natural algebraic object which corresponds to the first case also, and this can be defined without any regular neighbourhood theory. Namely consider the family of all conjugacy classes of canonical splittings of G over two-ended subgroups. Any finite subset of this family will be compatible and so will determine a graph of groups structure for G, by Theorem B.2.5. The accessibility result of Theorem 7.11 implies that if we take an ascending sequence of such finite families of splittings, the resulting sequence of graphs of groups structures will stabilise. Thus G has only finitely many conjugacy classes of canonical splittings and they determine a natural graph of groups structure

Γ' for G. Note that Γ' need not be bipartite. The following result gives the connection between $\Gamma(\mathcal{F}_1 : G)$ and Γ'. Recall that if a minimal graph of groups structure Γ for G has a redundant vertex, we can remove it by replacing the two incident edges by a single edge. If Γ is finite, repeating this will yield a graph of groups structure with no redundant vertices.

Theorem 9.5. — *Let G be a one-ended, finitely presented group such that whenever H is a two-ended subgroup and $e(G,H) \geqslant 2$, then H has small commensuriser. Suppose G is not of finite-by-Fuchsian type and possesses a nontrivial almost invariant subset over some two-ended subgroup H of G, so that G has a canonical splitting. Let Γ' denote the graph of groups structure for G determined by a maximal family of non-conjugate canonical splittings of G. Then*

(1) *The graph of groups structure Γ' for G is obtained from $\Gamma(\mathcal{F}_1 : G)$ by removing all redundant vertices as above.*

(2) *The vertex groups of Γ' are each simple, two-ended or of finite-by-Fuchsian type.*

(3) *Any element of $\mathcal{F}_1$ is enclosed by some vertex of Γ' which is either two-ended or of finite-by-Fuchsian type. In particular, any splitting of G over a two-ended subgroup is enclosed by such a vertex of Γ'.*

Remark 9.6. — It follows from this result that Γ' is closely related to the reduced algebraic regular neighbourhood of the collection of all nontrivial almost invariant subsets of G which are over a two-ended subgroup.

Proof. — Let Γ_1 denote the graph of groups structure obtained from $\Gamma(\mathcal{F}_1 : G)$ by removing all the redundant vertices. Thus the edge splittings of Γ_1 are exactly those of $\Gamma(\mathcal{F}_1 : G)$, but now distinct edges of Γ_1 have non-conjugate splittings of G associated. Part 9) of Theorem 9.4 implies that there is a bijection between the edge splittings of Γ' and of Γ_1, so that these graphs of groups structures for G are isomorphic as required. This proves part 1). Now parts 2) and 3) follow immediately from Theorems 9.2 and 9.4. □

Remark 9.7. — In the case of word hyperbolic groups, the graph of groups $\Gamma(\mathcal{F}_1 : G)$ obtained in Theorem 9.2 is similar to that obtained by Bowditch in [**5**], but it may differ from that in [**8**] when $\Gamma(\mathcal{F}_1 : G)$ has isolated V_0–vertices corresponding to splittings over two-ended subgroups H which have at most three coends. This point is closely related to the discussion at the end of chapter 1 of the difference between the characteristic submanifold $V(M)$ of a 3–manifold M and the submanifold $V'(M)$. In particular, if $V(M)$ has a solid torus component W whose frontier consists of a single annulus of degree 2 or 3 in W, then this annulus determines a splitting of $G = \pi_1(M)$ over an infinite cyclic subgroup H which has two or three coends in G. Another difference between the decompositions is that ours has redundant vertices corresponding to canonical splittings which may not appear in Bowditch's decomposition. However, such redundant vertices do not appear in the reduced version of our construction.

The decomposition Γ' described in Theorem 9.5 also differs from that obtained by Sela in [**49**] for similar reasons. He further decomposes some of the vertex groups not of finite-by-Fuchsian type, and seems to take unfoldings of some of the edge splittings considered here. Moreover some vertex groups which are of finite-by-Fuchsian type in our terminology (for example, pairs of pants) may be counted as simple in his decomposition. Conditions 3), 4) and 5) of Theorem 9.4 are not in either of their results.

There is yet another graph of groups structure for G which is also natural and is similar to Γ'. We define a splitting σ of G to be *splitting-canonical* if it has zero intersection number with every splitting of G over a two-ended subgroup. As for Γ', there can be only finitely many conjugacy classes of splitting-canonical splittings of G which are over two-ended subgroups, and these yield a natural graph of groups structure Γ'' for G, whose edge splittings are these splittings of G. The concepts of canonical and splitting-canonical are in general different. The following example demonstrates this, and simultaneously gives some insight into the properties of our regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$.

Example 9.8. — Let A be a finitely presented group which is one-ended, admits no splitting over a two-ended subgroup, and is not finite-by-Fuchsian. For example, A could be the fundamental group of a closed hyperbolic 3–manifold. Let B be an infinite cyclic group, let $d \geqslant 4$ be an integer, and let dB denote the subgroup of index d in B. Let C denote an infinite cyclic subgroup of A and let G be constructed from A and B by identifying C with dB. Thus $G = A *_C B$, and G is one-ended and finitely presented. Note also that if A is word hyperbolic, then so is G. Suppose that d is a composite number, say $d = mn$. Then $G = A *_C mB *_{mB} B$. We let σ_m denote the splitting of G over mB given by $G = (A *_C mB) *_{mB} B$. Thus the original splitting of G over C is σ_d. The lemma below shows that if p is prime and $d = p^2$, then the splitting σ_p of G over pB is splitting-canonical but not canonical.

Lemma 9.9. — *Let A, B, C, G and d be as in the above example. If p is prime and $d = p^2$, then the splitting σ_p of G over pB is splitting-canonical but not canonical.*

Proof. — We start by considering the construction of G for general values of d. Let T be the G–tree determined by the given splitting σ_d of G over C, and let v be the vertex with stabiliser B. Then v has valence d in T. Let b denote a generator of B, and let $s_0, \ldots, s_{d-1}$ denote the edges of T which are incident to v, labeled so that $bs_i = s_{i+1}$, where the suffix is regarded as being modulo d, so that $s_d = s_0$. As usual, we choose v as the basepoint of T, and define $\varphi : G \to V(T)$ by the formula $\varphi(g) = gv$. We also orient the edges $s_0, \ldots, s_{d-1}$ towards v. Let Z_i denote the almost invariant subset of G associated to s_i. As v has finite valence, Corollary 4.16 implies that the almost invariant subsets of G which are enclosed by v are all equivalent to some union of the Z_i's and their complements. As each Z_i is the union of all the remaining Z_j^*'s, it follows that any almost invariant subset of G which is enclosed by v is equivalent to some union of the Z_i^*'s. As the Z_i^*'s are disjoint, we can apply the proof of Proposition 7.16. If $d \geqslant 4$, we have Case 2) of that proposition so that the regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$ is the graph of groups with a single edge given by the splitting $G = A *_C B$. The vertex which carries B is a V_0–vertex of small commensuriser type. If d equals 2 or 3, then we have Case 1) of that proposition so that $\Gamma(\mathcal{F}_1 : G)$ has graph corresponding to $G = A *_C C *_C B$, and the vertex carrying C is the only V_0–vertex.

If $d = mn$, then, as discussed in the above example, we have the splitting σ_m of G over mB given by $G = (A *_C mB) *_{mB} B$. This induces a refinement T_m of T, in which the vertex v of valence $d = mn$ is replaced by a tree which is the cone on points $v_1, \ldots, v_m$, and the d edges which were attached to v are now attached to the vertices $v_1, \ldots, v_m$ with n of these edges being attached to each v_j. A similar replacement occurs for each translate of v, so that T_m is a G–tree. This shows that if we pick some Z_i, then the union of its translates by mB is associated to a conjugate of the splitting σ_m of G. It is also easy to see that, if $m, n \geqslant 2$, then σ_m is never canonical. For if X denotes the union of the translates of Z_0^* by mB, we let Y denote $Z_0^* \cup Z_1^*$, and clearly X crosses Y.

Now suppose that G has a splitting over some subgroup H, such that an associated H–almost invariant subset X of G is enclosed by v. As X is equivalent to some union of the Z_i^*'s, and each Z_i^* has stabiliser dB, it follows that $dB \subset H \subset B$. Thus we can write $d = mn$ and $H = mB$. As the splitting over H and the given splitting over C are compatible, we can split T at v and at all its translates, to obtain a new G–tree which must be isomorphic to the above G–tree T_m. It follows that the splitting over H is conjugate to the splitting σ_m over mB.

Now pick a prime p and let $m = n = p$, so that $d = p^2$. The above discussion shows that if an almost invariant subset X is associated to a splitting σ of G and is enclosed by v, then, up to equivalence, either X is a single Z_i or X is the orbit of a single Z_i under the action of the subgroup pB of B. Thus σ must be conjugate either to the original splitting σ_d over C, or to the splitting σ_p over pB. In particular, it follows that the splitting σ_p is splitting-canonical. As σ_p is not canonical, the result follows. □

CHAPTER 10

CANONICAL DECOMPOSITIONS OVER TWO-ENDED GROUPS WHEN COMMENSURISERS ARE LARGE

Again we consider a one-ended, finitely presented group G and nontrivial almost invariant subsets over two-ended subgroups. This time, we do not assume that our two-ended subgroups have small commensurisers. This leads to two additional difficulties. When we form the regular neighbourhood of a finite number of nontrivial almost invariant subsets of G which are over two-ended subgroups, the edge groups of the regular neighbourhood may no longer be two-ended. As pointed out in the introduction, this happens in the topological situation. For if one wants to enclose essential annuli in a 3–manifold M, one may obtain Seifert fibre space components of $V(M)$ whose frontier has toral components. In the case of general finitely presented groups, the edge groups may be even more complicated. Thus if we proceed, as in the previous chapter, to take regular neighbourhoods of larger and larger finite collections of almost invariant subsets over two-ended subgroups, we will not be able to show that our construction stabilises. Hence we are forced to consider directly regular neighbourhoods of infinite families of almost invariant sets. This leads to the additional problem of showing that the pretrees which appear in the regular neighbourhood construction are discrete. For this, Theorem 8.2 on the finite generation of certain Boolean algebras plays a key role.

The main result of this chapter is our version of the JSJ-decomposition for arbitrary finitely presented groups with one end.

Theorem 10.1. — *Let G be a one-ended, finitely presented group, and let $\mathcal{F}_1$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a two-ended subgroup.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$. Each V_0–vertex v of $\Gamma(\mathcal{F}_1 : G)$ satisfies one of the following conditions:

(1) *v is isolated.*

(2) *v is of finite-by-Fuchsian type.*

(3) *$G(v)$ is the full commensuriser* $\mathrm{Comm}_G(H)$ *for some two-ended subgroup H, such that $e(G,H) \geqslant 2$.*

$\Gamma(\mathcal{F}_1 : G)$ consists of a single vertex if and only if $\mathcal{F}_1$ is empty, or G itself satisfies one of conditions 2) or 3) above.

Remark 10.2. — We will say that a V_0–vertex in case 3) above is *of commensuriser type*, if v is not isolated nor of finite-by-Fuchsian type, and is *of large commensuriser type*, if in addition H has large commensuriser.

Note that even if G is finitely presented, Example 11.1 shows that a vertex group of commensuriser type need not be finitely generated.

Proof. — We start by picking a representative for each element of $\mathcal{F}_1$, subject to the condition that if A and B are elements of $\mathcal{F}_1$ such that $B = gA$, for some g in G, then the representatives X and Y chosen for A and B must satisfy $Y = gX$.

Before proceeding further, we consider the very special case when G is virtually $\mathbb{Z} \times \mathbb{Z}$. The equivalence classes of nontrivial almost invariant subsets of $\mathbb{Z} \times \mathbb{Z}$ which are over two-ended subgroups correspond to all the simple closed curves on the torus. It follows that the collection of all the chosen representatives of elements of $\mathcal{F}_1$ is cross-connected, so that the required regular neighbourhood exists and consists of a single V_0–vertex with associated group G. As G is finite-by-Fuchsian, this proves the theorem in this case. In the following we will assume that G is not virtually $\mathbb{Z} \times \mathbb{Z}$.

Next we let $\mathcal{E}_0$ denote the subset of $\mathcal{F}_1$ whose elements are represented by almost invariant subsets of G which are over subgroups with small commensuriser. The proof of Theorem 9.2 applies to show that the regular neighbourhood construction of chapter 3 works to yield $\Gamma_0 = \Gamma(\mathcal{E}_0 : G)$. In what follows, we will express $\mathcal{F}_1$ as an ascending sequence of subsets $\mathcal{E}_i$ of $\mathcal{F}_1$, for $i \geqslant 0$, and show that each $\mathcal{E}_i$ has a regular neighbourhood $\Gamma_i = \Gamma(\mathcal{E}_i : G)$. Finally, we will show that the sequence Γ_i must eventually stabilise. The resulting graph of groups structure Γ for G will be the required regular neighbourhood $\Gamma = \Gamma(\mathcal{F}_1 : G)$.

By our definition of $\mathcal{E}_0$, any element of $\mathcal{F}_1 - \mathcal{E}_0$ will be represented by a H–almost invariant subset X of G, such that H is two-ended and has large commensuriser. From the collection of all such subgroups of G, we choose one group from each conjugacy class and denote the chosen groups by H_j, $j \geqslant 1$. We choose $\mathcal{E}_{i+1}$ to be the union of $\mathcal{E}_i$ and all translates by G of elements of $B(H_{i+1})$. Clearly the union of the $\mathcal{E}_i$'s equals $\mathcal{F}_1$.

Let E_0 denote the collection of the chosen representatives of all the elements of $\mathcal{E}_0$. Denote H_1 by H, so that $\mathcal{E}_1$ is the union of $\mathcal{E}_0$ and all translates by G of elements of $B(H)$. We let E_1 denote the collection of the chosen representatives of all the elements of $\mathcal{E}_1$. Theorem 8.2 tells us that $B(H)$ has a finite system of generators when we regard $B(H)$ as a Boolean algebra over $\mathrm{Comm}_G(H)$. We let $X_1, \ldots, X_n$ be the chosen representatives of this system of generators. Thus any element of $B(H)$

can be represented by taking finite sums of finite intersections of the X_i's and their complements. The proof of Proposition 8.6 shows that any isolated element of $B(H)$ is a translate of some X_i by an element of $\mathrm{Comm}_G(H)$. We will also need to consider the union of E_0 with all the isolated elements of E_1. We denote this union by E'_0. Thus E'_0 consists of E_0 together with all the translates by G of those X_i's which are isolated in E_1.

In order to show that $\mathcal{E}_1$ has a regular neighbourhood $\Gamma_1 = \Gamma(\mathcal{E}_1 : G)$, we will need to use some more facts from chapter 8. Recall that G is finitely presented, that H is a two-ended subgroup of G with large commensuriser and that G has a nontrivial H–almost invariant subset. As we are assuming that G is not virtually $\mathbb{Z}\times\mathbb{Z}$, Proposition 8.1 tells us that the number of coends of H in G is infinite. Proposition 8.6 tells us that the corresponding cross-connected components (CCC's) consist of translates of a finite number of isolated almost invariant sets and a single CCC H_∞ which consists of an infinite number of almost invariant sets (we will call this 'the infinite CCC' corresponding to H). We saw that any nontrivial K–almost invariant set with K commensurable with H is either isolated or in H_∞. Note also that the stabiliser of H_∞ is $\mathrm{Comm}_G(H)$.

Let P_0 denote the pretree of CCC's of $\overline{E_0}$, let P_1 denote the pretree of CCC's of $\overline{E_1}$, and let P'_0 denote the pretree of CCC's of $\overline{E'_0}$. We know that P_0 is discrete and want to show that P_1 is discrete. We note that the natural maps $P_0 \to P'_0 \to P_1$ are injective.

Recall that the proof that P_0 is discrete, depended crucially on the discreteness of E_0, which holds because E_0 consists of translates of a finite family of almost invariant sets. Recall that discreteness of a partially ordered set F means that, for any U, $V \in F$, there are only finitely many $Z \in F$ such that $U \leqslant Z \leqslant V$. (See Lemma 3.1). The description of E'_0 implies that this also is discrete, so that P'_0 is discrete.

To show that P_1 is discrete, we need to show that if A and C are distinct CCC's of $\overline{E_1}$, then there are only finitely many CCC's B of $\overline{E_1}$ such that ABC. We know that there are only finitely many finite CCC's between A and C, because these come from elements of E'_0, which is discrete. So it remains to show that there are only finitely many infinite CCC's B between A and C. By construction, the only infinite CCC's in E_1 are translates of the infinite CCC H_∞ corresponding to the commensurability class of H. We choose elements U and W of E_1, such that $\overline{U} \in A$ and $\overline{W} \in C$. Suppose that $B = H_\infty$. The definition of betweenness for CCC's implies that there exists an element $\overline{V}$ of H_∞ such that $U < V < W$. Lemma 10.3 below shows that we can choose V in a special way. It shows that there is an almost invariant subset X of G which is a translate of one of the X_i's by an element of $\mathrm{Comm}_G(H)$, such that $U \leqslant X \leqslant W$. Note that X will represent an element of $B(H)$. If B is a translate gH_∞ of H_∞ such that ABC, then H_∞ lies between $g^{-1}A$ and $g^{-1}C$, and applying Lemma 10.3 to this situation yields an almost invariant subset X of G as above except that

$U \leqslant gX \leqslant W$. Thus gX is a translate by an element of $g\,\mathrm{Comm}_G(H)$ of one of the X_i's. As only finitely many translates of the X_i's can lie between U and W, it follows that there are only finitely many translates of H_∞ between A and C. Thus P_1 is a discrete pretree.

Similar arguments yield an inductive proof that the pretree P_i is discrete, for all $i \geqslant 0$, so that each $\mathcal{E}_i$ has a regular neighbourhood $\Gamma_i = \Gamma(\mathcal{E}_i : G)$, as required.

In order to show that the sequence Γ_i stabilises, we first need to describe the V_0–vertices of Γ_i. The results of chapters 7 and 8 show that these are of three types, isolated, finite-by-Fuchsian type, and commensuriser type, which is where a V_0–vertex carries the group $\mathrm{Comm}_G(H)$ for some two-ended subgroup H of G such that $e(G, H) \geqslant 2$. Further, our construction implies that each Γ_i has i vertices of large commensuriser type. Now we claim that each V_0–vertex of large commensuriser type encloses a splitting of G over a two-ended subgroup of G. For consider the V_0–vertex v determined by the infinite CCC H_∞. Theorem 8.3 tells us that G splits over some subgroup commensurable with H. As the almost invariant subset associated to this splitting must lie in $Q(H)$, and the proof of Lemma 8.6 shows that every element of $Q(H)$ is enclosed by H_∞, it follows that v encloses a splitting of G over a two-ended subgroup, as claimed. Thus we can refine Γ_i by splitting at each V_0–vertex of large commensuriser type using such a splitting, to obtain a new graph of groups structure Γ_i' for G. This construction means that if we let $f(i)$ denote the number of those edge splittings of Γ_i' which are over a two-ended group, then $f(i)$ is strictly increasing. Now Theorem 7.11 implies that the sequence Γ_i' must stabilise and hence that the sequence Γ_i must stabilise, as required. We call this final graph of groups Γ. By construction, Γ is the required regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$. Each V_0–vertex group of Γ satisfies one of the three conditions in the statement of the theorem because this holds for each Γ_i.

Finally, Γ will consist of a single vertex if and only if either $\mathcal{F}_1$ is empty or the representatives of $\mathcal{F}_1$ lie in a single CCC. In the second case, the vertex group of Γ must satisfy one of the three conditions in the statement of the theorem. As the vertex group is G and we assumed that G is one-ended, it follows that $\Gamma(\mathcal{F}_1 : G)$ consists of a single vertex if and only if $\mathcal{F}_1$ is empty, or G is of finite-by-Fuchsian type, or G equals $\mathrm{Comm}_G(H)$ for some two-ended subgroup H, such that $e(G, H) \geqslant 2$. □

Now we prove the following result which was used in the above proof. Recall that $X_1, \ldots, X_n$ are the chosen representatives of elements of $B(H)$ which generate $B(H)$ over $\mathrm{Comm}_G(H)$.

Lemma 10.3. — *Let A, B and C be distinct CCC's of $\overline{E_1}$, such that B equals the infinite CCC H_∞, and B lies between A and C. Let U and W be almost invariant subsets of G such that $\overline{U} \in A$ and $\overline{W} \in C$. Then there is an almost invariant subset X of G which is a translate of one of the X_i's by an element of* $\mathrm{Comm}_G(H)$, *such that $U \leqslant X \leqslant W$.*

Proof. — As B lies between A and C, there is $\overline{V} \in B$ with $U < V < W$. In particular, V represents an element of $B(H)$. As $B(H)$ is generated over $\mathrm{Comm}_G(H)$ by the X_i's, any element of $B(H)$ can be represented as a finite sum of finite intersections of translates of the X_i's and their complements. We will need to prove the following two claims.

Claim 10.4. — *If Y and Z represent elements of $B(H)$ and $U < (Y + Z) < W$, then either U or W represents an element of $B(H)$, or one of the four sets Y, Y^*, Z, Z^* lies between U and W.*

Claim 10.5. — *If Y and Z represent elements of $B(H)$ and $U < (Y \cap Z) < W$, then either U or W represents an element of $B(H)$, or one of the four sets Y, Y^*, Z, Z^* lies between U and W.*

Here is the argument assuming these claims. Starting from an expression of V as a finite sum of finite intersections of the X_i's and their complements, we apply one of these two claims. If one of the four sets Y, Y^*, Z, Z^* lies between U and W, we again apply one of these two claims, and repeat this process as long as possible. This process will eventually stop, at which point either we will see that U or W represents an element of $B(H)$, or we will find an almost invariant subset X of G which is a translate of one of the X_i's by an element of $\mathrm{Comm}_G(H)$, such that $U < X < W$. Such a set X completes the proof of the lemma. If U represents an element of $B(H)$, then U is enclosed by the CCC $B = H_\infty$. As U is also enclosed by the CCC A which is distinct from B, part 1) of Lemma 4.9 shows that U is equivalent to Z_s for some edge s of T which is incident to H_∞. In particular, U represents an isolated element of $B(H)$, which implies that U is equivalent to a translate of some X_i by an element of $\mathrm{Comm}_G(H)$. In this case, we find the required set X by simply choosing $X = U$. Similarly if W represents an element of $B(H)$, we can choose $X = W$. Thus in all cases, we have found the required set X.

It remains to prove Claims 10.4 and 10.5.

To motivate our proof of Claim 10.4, consider the special case when $U \subset (Y+Z) \subset W$, and Y and Z are each nested with respect to U and to W. The group G is divided into the four disjoint subsets $Y^{(*)} \cap Z^{(*)}$. As U is nested with respect to Y and Z, either U is contained in one of these four sets or U^* is contained in one of these four sets. Now $Y + Z = (Y \cap Z^*) \cup (Y^* \cap Z)$, so it follows that either $U \subset Y \cap Z^*$ or $U \subset Y^* \cap Z$. Similarly, as $W^* \subset (Y + Z)^* = (Y \cap Z) \cup (Y^* \cap Z^*)$, it follows that we must have $W^* \subset Y \cap Z$ or $W^* \subset Y^* \cap Z^*$. For each of the four possibilities, we will then obtain one of the four inclusions $U \subset Y^{(*)} \subset W$ or $U \subset Z^{(*)} \subset W$ as required. Here is the formal proof of Claim 10.4.

Suppose that neither U nor W represents an element of $B(H)$. Thus neither U nor W is enclosed by H_∞. As Y lies in $B(H)$, it is enclosed by H_∞. As it is enclosed by a different CCC from U or W, it cannot cross U or W. Similarly, Z cannot cross U

or W. Now consider the inequality $U < (Y + Z)$. This is equivalent to the statement that $U \cap (Y + Z)^*$ is small. As $(Y + Z)^* = (Y \cap Z) \cup (Y^* \cap Z^*)$, it follows that $U \cap (Y \cap Z)$ and $U \cap (Y^* \cap Z^*)$ are each small. Hence $U^* \cap (Y \cap Z)$ and $U^* \cap (Y^* \cap Z^*)$ are each large (i.e. not small). This implies that each of $U^* \cap Y$, $U^* \cap Z$, $U^* \cap Y^*$ and $U^* \cap Z^*$ is large. As U does not cross Y or Z, we know that one of the four sets $U^{(*)} \cap Y^{(*)}$ is small and that one of the four sets $U^{(*)} \cap Z^{(*)}$ is small. It follows that one of the two sets $U \cap Y^{(*)}$ is small and that one of the two sets $U \cap Z^{(*)}$ is small. If $U \cap Y$ is small, then $U \cap (Y \cap Z^*)$ is also small. As we already know that $U \cap (Y \cap Z)$ and $U \cap (Y^* \cap Z^*)$ are small, it follows that we have $U \leqslant (Y^* \cap Z)$. Similar arguments apply in the other three cases. We conclude that $U \leqslant (Y^* \cap Z)$ or $U \leqslant (Y \cap Z^*)$.

Similar arguments using the inequality $(Y + Z) < W$ show that $W^* \leqslant (Y \cap Z)$ or $W^* \leqslant (Y^* \cap Z^*)$. In each of these four cases, one sees that one of the four sets Y, Y^*, Z, Z^* is either equal to U or W or lies between U and W. Thus either U or W represents an element of $B(H)$, or one of the four sets Y, Y^*, Z, Z^* lies between U and W, as required.

The proof of Claim 10.5 is somewhat simpler. Suppose that $U < (Y \cap Z) < W$. Much as in the proof of Claim 10.4, one can show that as W does not cross Y or Z, we have $W^* \leqslant (Y \cap Z^*)$, $W^* \leqslant (Y^* \cap Z)$ or $W^* \leqslant (Y^* \cap Z^*)$. It follows that we have one of $Y \leqslant W$ or $Z \leqslant W$. Thus we have $U \leqslant Y \leqslant W$ or $U \leqslant Z \leqslant W$. As above, it follows that either U or W represents an element of $B(H)$, or one of the two sets Y and Z lies between U and W, as required. □

In order to understand $\Gamma(\mathcal{F}_1 : G)$ in more detail, our next result lists some properties which follow almost immediately from the above theorem and the properties of an algebraic regular neighbourhood.

Theorem 10.6. — *Let G be a one-ended, finitely presented group, and let $\mathcal{F}_1$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a two-ended subgroup.*

Then the regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$ is a minimal bipartite graph of groups decomposition of G with the following properties:

(1) *each V_0–vertex v of $\Gamma(\mathcal{F}_1 : G)$ satisfies one of the following conditions:*

(a) *v is isolated.*

(b) *v is of finite-by-Fuchsian type.*

(c) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some two-ended subgroup H, such that $e(G, H) \geqslant 2$.*

Further, if H is a two-ended subgroup of G such that $e(G, H) \geqslant 2$, and if H has large commensuriser, then $\Gamma(\mathcal{F}_1 : G)$ will have a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$.

(2) *If an edge of $\Gamma(\mathcal{F}_1 : G)$ is incident to a V_0–vertex of type a) or b) above, then it carries a two-ended group.*

(3) *any representative of an element of* $\mathcal{F}_1$ *is enclosed by some* V_0*–vertex of* $\Gamma(\mathcal{F}_1 : G)$*, and each* V_0*–vertex of* $\Gamma(\mathcal{F}_1 : G)$ *encloses such a subset of* G*. In particular, any splitting of* G *over a two-ended subgroup is enclosed by some* V_0*–vertex of* $\Gamma(\mathcal{F}_1 : G)$*.*

(4) *if* X *is a nontrivial almost invariant subset of* G *over a finitely generated subgroup* H*, and if* X *does not cross any element of* $\mathcal{F}_1$*, then* X *is enclosed by a* V_1*–vertex of* $\Gamma(\mathcal{F}_1 : G)$*.*

(5) *if* X *is a* H*–almost invariant subset of* G *associated to a splitting of* G *over* H*, and if* X *does not cross any element of* $\mathcal{F}_1$*, then* X *is enclosed by a* V_1*–vertex of* $\Gamma(\mathcal{F}_1 : G)$*.*

(6) *the* V_1*–vertices of* $\Gamma(\mathcal{F}_1 : G)$ *are simple. (See Definition 9.1.)*

(7) *If* Γ_1 *and* Γ_2 *are minimal bipartite graphs of groups structures for* G *which satisfy conditions 3) and 5) above, then they are isomorphic provided that there is a one-to-one correspondence between their isolated* V_0*–vertices, and that any non-isolated* V_0*–vertex of* Γ_1 *or* Γ_2 *encloses some non-isolated element of* $\mathcal{F}_1$*.*

(8) *The graph of groups* $\Gamma(\mathcal{F}_1 : G)$ *is invariant under the automorphisms of* G*.*

(9) *The canonical splittings of* G *over two-ended subgroups are precisely those edge splittings of* $\Gamma(\mathcal{F}_1 : G)$ *which are over two-ended subgroups. This includes, but need not be limited to, all those edges of* $\Gamma(\mathcal{F}_1 : G)$ *which are incident to* V_0*–vertices whose associated groups are of type a) or b) above.*

Proof. — The description of the possible types of V_0–vertex given in part 1) follows from Theorem 10.1. Further, if H is a two-ended subgroup of G such that $e(G,H) \geqslant 2$, and if H has large commensuriser, then either G is virtually $\mathbb{Z} \times \mathbb{Z}$, or $e(G,H)$ is infinite. In the first case, $\Gamma(\mathcal{F}_1 : G)$ consists of a single vertex labeled G, and in the second case, Proposition 8.6 shows that $\Gamma(\mathcal{F}_1 : G)$ will have a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$. Thus in either case, $\Gamma(\mathcal{F}_1 : G)$ will have a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$.

The proofs for parts 2)-9) are the same as for parts 2)-9) of Theorem 9.4. □

At this point, we note the connection between the above results and the Algebraic Annulus Theorem [**15**] for finitely generated groups. (See also [**43**] and [**5**] for the case of word hyperbolic groups.) The proof we give below is for finitely presented groups only and is not essentially different from that given by Dunwoody and Swenson in [**15**]. We include the argument here for completeness only. In the topological context, one can deduce the Annulus Theorem from the JSJ-decomposition in much the same way. Clearly regular neighbourhood theory is not essential for the proof of the Algebraic Annulus Theorem. Nor can it yield a proof for groups which are not finitely presented.

***Theorem 10.7* (Algebraic Annulus Theorem).** — *Let* G *be a one-ended, finitely presented group. If* G *has a two-ended subgroup* H *such that* $e(G,H) \geqslant 2$*, then either* G *splits over some two-ended subgroup or* G *is of finite-by-Fuchsian type.*

Proof. — The assumption implies that the set $\mathcal{F}_1$ in Theorems 10.1 and 10.6 is nonempty. Applying Theorem 10.1 yields the regular neighbourhood $\Gamma = \Gamma(\mathcal{F}_1 : G)$. If Γ consists of a single vertex, then G is of finite-by-Fuchsian type. Otherwise, each V_0–vertex of Γ has at least one incident edge. Any edge incident to a V_0–vertex of type a) or b) carries a two-ended group and so yields a splitting of G over such a group. If Γ has no such V_0–vertices, then it must have a V_0–vertex v of type c), and Theorem 8.3 shows that G splits over some two-ended subgroup. The result follows. □

A key point about the preceding arguments was that we considered all nontrivial almost invariant subsets of G over two-ended subgroups and did not restrict attention to those which are associated to splittings. However, now we have Theorem 10.1, it is quite easy to deduce the existence of a regular neighbourhood of this smaller collection of almost invariant subsets. The result we obtain is the following.

Theorem 10.8. — *Let G be a one-ended, finitely presented group, and let $\mathcal{S}_1$ denote the collection of equivalence classes of all almost invariant subsets of G which are associated to a splitting of G over a two-ended subgroup.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{S}_1 : G)$. Each V_0–vertex v of $\Gamma(\mathcal{S}_1 : G)$ satisfies one of the following conditions:

(1) v is isolated.

(2) v is of finite-by-Fuchsian type.

(3) $G(v)$ contains a two-ended subgroup H which it commensurises, such that $e(G, H) \geqslant 2$.

If $\Gamma(\mathcal{S}_1 : G)$ consists of a single vertex, then either $\mathcal{S}_1$ is empty, or G itself satisfies one of conditions 2) or 3) above.

Remark 10.9. — Note that even if G is finitely presented, Example 11.1 shows that a vertex group of type 3) need not be finitely generated. Note also that if G commensurises a two-ended subgroup H such that $e(G, H) \geqslant 2$, then Example 10.10 shows that $\Gamma(\mathcal{S}_1 : G)$ need not consist of a single vertex. This is in contrast with the situation of Theorem 10.6.

Proof. — We start from the regular neighbourhood $\Gamma(\mathcal{F}_1 : G)$ obtained in Theorem 10.1. Now we know that the construction of chapter 3 works, we can consider this construction directly. As in the proof of Theorem 10.1, we start by picking a representative for each element of $\mathcal{F}_1$, subject to the condition that if A and B are elements of $\mathcal{F}_1$ such that $B = gA$, for some g in G, then the representatives X and Y chosen for A and B must satisfy $Y = gX$. This determines the set E of all translates of these subsets of G, and we let S denote the subset of E consisting of almost invariant subsets of G which are associated to a splitting. When we replace E by S, we want to describe how the CCC's and their stabilisers alter. We claim that each G–orbit of CCC's of $\overline{E}$ is the union of a finite number of G–orbits of CCC's of $\overline{S}$. Given this, it is not difficult to verify that the pretree determined by S must be discrete, so that the regular neighbourhood $\Gamma(\mathcal{S}_1 : G)$ exists. The reason for our claim is simply that otherwise, some V_0–vertex of $\Gamma(\mathcal{F}_1 : G)$ would enclose an infinite number of non-conjugate compatible splittings of G over two-ended subgroups obtained by picking a splitting from each CCC of $\overline{S}$, and this would contradict the accessibility result of Theorem 7.11. □

Now we can give the examples referred to in Remark 10.9. These examples are to demonstrate that if G commensurises a two-ended subgroup H such that $e(G, H) \geqslant 2$, then $\Gamma(\mathcal{S}_1 : G)$ need not consist of a single vertex.

Example 10.10. — Let $G_{p,q} = A *_C B$, where A and B are both infinite cyclic and C has index p in A and index q in B. Thus $G_{p,q}$ centralises, and hence commensurises, the two-ended subgroup C. If $p, q \geqslant 2$, then G splits over C, so that $e(G, C) \geqslant 2$. If, in addition, we exclude the case $p = q = 2$, then we claim that $\Gamma(\mathcal{S}_1 : G_{p,q})$ does not consist of a single vertex. If p and q are both prime, then part 1) of Lemma 10.11 below tells us that there is only one splitting of $G_{p,q}$ over a two-ended subgroup, up to conjugacy, which implies that $\Gamma(\mathcal{S}_1 : G_{p,q})$ is the graph of groups associated to $G = A *_C C *_C B$, where the vertex carrying C is the only V_0–vertex. If either p or q is composite, then there are other splittings of $G_{p,q}$ over a two-ended subgroup. For example, if $q = rs$, then $G = A *_C rB *_{rB} B$, so that G splits over rB. However, part 2) of Lemma 10.11 shows that the given splitting of $G_{p,q}$ over C is isolated among all the splittings of $G_{p,q}$ over two-ended subgroups, so that $\Gamma(\mathcal{S}_1 : G_{p,q})$ has more than one V_0–vertex, and in particular does not consist of a single vertex.

These examples are closely related to some topological examples discussed in chapter 1. Let $W_{p,q}$ denote the 3–manifold which is obtained by gluing two solid tori along an annulus A which has degree p in one solid torus and degree q in the other. Thus $G_{p,q} = \pi_1(W_{p,q})$. If $p, q \geqslant 2$ and we exclude the case $p = q = 2$, then the annulus A is the only embedded essential annulus in $W_{p,q}$, up to isotopy. Note that this holds whether or not p and q are prime. The extra splittings discussed in the previous paragraph in the case when p or q is composite cannot be represented by an embedded annulus.

On the other hand, $W_{2,2}$ is filled by essential embedded annuli, and this implies that $\Gamma(\mathcal{S}_1 : G_{2,2})$ consists of a single vertex.

Lemma 10.11. — *Let $G_{p,q} = A *_C B$, where A and B are both infinite cyclic and C has index p in A and index q in B, where $p, q \geqslant 2$.*

(1) *If p and q are prime and we exclude the case $p = q = 2$, then this is the only splitting of $G_{p,q}$ over a two-ended subgroup, up to conjugacy.*

(2) *If we exclude the case $p = q = 2$, then the given splitting of G over C is isolated among all the splittings of $G_{p,q}$ over two-ended subgroups.*

Proof. — Denote the given splitting of G over C by σ, and suppose that $G_{p,q}$ has a splitting τ over a two-ended subgroup H. As $G_{p,q}$ is torsion free, H must be infinite cyclic. Let T denote the G–tree determined by the splitting τ, and let c denote a generator of C. As C is central in $G_{p,q}$, either c fixes every vertex of T, or c fixes no vertex of T. In the first case C also fixes every edge of T and so is contained in H. In the second case, C has an axis l in T, which must be G–invariant, so it follows that $l = T$. As H stabilises one edge of l, it must act trivially on T, so that $G_{p,q}$ is isomorphic to $\mathbb{Z} \times \mathbb{Z}$ or to the fundamental group of the Klein bottle, both of which are clearly impossible. (Note that if $p = q = 2$, then $G_{p,q}$ *is* isomorphic to the fundamental group of the Klein bottle. This is why we excluded this case.) We conclude that C is contained in H. Now we consider the quotient group $G_{p,q}/C$. The splitting σ of $G_{p,q}$ induces the free product splitting $\mathbb{Z}_p * \mathbb{Z}_q$ of $G_{p,q}/C$, which we denote by σ'. The splitting τ of $G_{p,q}$ induces a splitting τ' of $G_{p,q}/C$ over the finite group H/C. Note that as $\mathbb{Z}_p * \mathbb{Z}_q$ has finite abelianisation, τ' cannot be a HNN extension.

If H/C is trivial, we use the fact that any free product of two freely indecomposable groups, neither of which is $\mathbb{Z}$ has a unique free product splitting up to conjugacy. This can be seen by considering the two G–trees corresponding to the two splittings, as the vertex stabilisers must be the same in both trees. It follows that if H/C is trivial, then σ' and τ' are conjugate, and hence that σ and τ are conjugate.

Now we suppose that H/C is nontrivial and denote it by K. Note that K is finite cyclic. Suppose that the splitting τ' of $\mathbb{Z}_p * \mathbb{Z}_q$ is $P *_K Q$. The Kuros Subgroup Theorem implies that K is conjugate to a subgroup of $\mathbb{Z}_p$ or $\mathbb{Z}_q$.

If p and q are both prime, it follows that K is conjugate to $\mathbb{Z}_p$ or $\mathbb{Z}_q$. By conjugating τ, and interchanging p and q if needed, we can suppose that $K = \mathbb{Z}_q$. As P and Q each contain $K = \mathbb{Z}_q$, the Kuros Subgroup Theorem implies that each has $\mathbb{Z}_q$ as a free factor. Further as $P \neq K \neq Q$, there are nontrivial groups L and M such that $P *_K Q = (K * L) *_K (K * M) = K * L * M$. This is impossible, as $\mathbb{Z}_p * \mathbb{Z}_q$ cannot be expressed as a free product of three nontrivial groups. This contradiction shows that when p and q are prime, H/C must be trivial. Now the preceding paragraph shows that σ and τ must be conjugate. As τ is arbitrary, this completes the proof of part 1) of the lemma.

Next we consider general values of p and q. We will show that σ and τ are compatible splittings of $G_{p,q}$. As τ is arbitrary this will show that σ is isolated among all the splittings of $G_{p,q}$ over two-ended subgroups, as required. Suppose that $K \subset \mathbb{Z}_q$. The Subgroup Theorem applied to $\mathbb{Z}_p$ and $\mathbb{Z}_q$ implies that each is conjugate into P or Q. As they generate $\mathbb{Z}_p * \mathbb{Z}_q$, one must be conjugate into P and one into Q. Thus we can suppose that $\mathbb{Z}_p$ is conjugate into P, and $\mathbb{Z}_q$ is conjugate into Q. Now the Subgroup Theorem applied to P and Q implies that $P \cong \mathbb{Z}_p * K * L$, for some group L, and $Q \cong \mathbb{Z}_q * M$, for some group M. It follows that

$\mathbb{Z}_p * \mathbb{Z}_q = P *_K Q \cong (\mathbb{Z}_p * K * L) *_K (\mathbb{Z}_q * M) = (\mathbb{Z}_p * L) * (\mathbb{Z}_q * M) = \mathbb{Z}_p * \mathbb{Z}_q * L * M$, which implies that L and M are trivial. Hence $P \cong \mathbb{Z}_p * K$ and $Q \cong \mathbb{Z}_q$, so that $\mathbb{Z}_p * \mathbb{Z}_q = P *_K Q \cong (\mathbb{Z}_p * K) *_K \mathbb{Z}_q$. Thus we have a graph of groups structure for $\mathbb{Z}_p * \mathbb{Z}_q$ with two edges, whose associated edge splittings are conjugate to σ' and τ'. This induces a graph of groups structure for $G_{p,q}$ with two edges, whose associated edge splittings are σ and τ, which shows that σ and τ are compatible, as required. $\square$

CHAPTER 11

EXAMPLES

We start this chapter with some specific examples, and then give some more general ones.

Our first example is of a one-ended, finitely presented group G such that the regular neighbourhood $\Gamma(G)$ of equivalence classes of all nontrivial almost invariant subsets of G which are over two-ended subgroups has a V_0–vertex of commensuriser type with a non-finitely generated vertex group.

Example 11.1. — We start by showing that there exists a one-ended, finitely presented group A which has an infinite cyclic subgroup H such that $\mathrm{Comm}_A(H)$ is not finitely generated. To construct such a group, we take a free group F of countably infinite rank, and an infinite cyclic group H. As F embeds in F_2, the free group of rank 2, we can embed $F \times H$ in $E = (F \times H) *_F F_2$. Clearly E is finitely generated and recursively presented, and $\mathrm{Comm}_E(H) = F \times H$. Now we embed E in a finitely presented group L using Higman's Embedding Theorem [**23**]. He first constructs a certain finitely presented group K and then constructs L as a HNN extension with vertex group $K \times E$. It is clear from the construction that $\mathrm{Comm}_L(H) = K \times \mathrm{Comm}_E(H) = K \times F \times H$, which is not finitely generated. If L is one-ended, we take $A = L$. Otherwise, the accessibility result of Dunwoody in [**12**] implies that L can be expressed as the fundamental group of a graph of groups with all edge groups finite, and all vertex groups having zero or one end. The vertex groups must then be finitely presented. Now one of these vertex groups must contain $K \times F \times H$, and this is the required one-ended, finitely presented group A.

Now let C denote $K \times F$ so that $\mathrm{Comm}_A(H) = C \times H$, let D denote any nontrivial finitely presented group, and let B denote $C * D$. We define $G = A *_{C \times H} (B \times H)$. As $B = C * D$, we can also write $G = A *_H (D * H)$, so that G is finitely presented and splits over H. Now $\mathrm{Comm}_G(H) = B \times H$ which is not finitely generated. As G splits over H and $\mathrm{Comm}_G(H)$ is not finitely generated, it follows that $\Gamma(G)$ has a V_0–vertex with associated group $\mathrm{Comm}_G(H)$.

If one can choose A so as not to split over any two-ended subgroup, then the graph of groups $\Gamma(G)$ consists of a single edge which induces the decomposition $G = A *_{C\times H} (B \times H)$.

It is natural to ask what can be said about the edge groups of $\Gamma(G)$ which are incident to a V_0–vertex of commensuriser type. We have already pointed out that if the commensuriser vertex group is not finitely generated, then some incident edge group must also be not finitely generated. In the case of 3–manifolds, the components of $V(M)$ which correspond to a V_0–vertex of commensuriser type are Seifert fibre spaces, and each frontier component is a vertical annulus or torus. In particular, the incident edge groups all contain the normal subgroup H carried by a regular fibre of the Seifert fibre space. However the following example shows that this does not hold in general.

Example 11.2. — Let K and L be free groups of rank at least 2, let H be an infinite cyclic group and let A and B be groups which properly contain K and L respectively. Then define $G = A *_K (K \times H) *_H (H \times L) *_L B$. Then $\mathrm{Comm}_G(H) = (K * L) \times H$. As G splits over H, it follows that $\Gamma(G)$ has a V_0–vertex of commensuriser type with associated group $\mathrm{Comm}_G(H)$. If we choose A and B to be one-ended, it is easy to see that G is also one-ended. For if G splits over a finite subgroup, the one-endedness of A implies that A must be conjugate into a vertex group G_1 of G. In particular K is conjugate into G_1 which implies that H, and then L must also be conjugate into G_1. As B contains L and is one-ended, it follows that B is also conjugate into G_1, which contradicts the assumption that G has a splitting.

If we assume that A has no two-ended subgroups D with $e(A, D) \geqslant 2$ and similarly for B, then we claim that $\Gamma(G)$ is the graph of groups given by

$$G = A *_K [(K * L) \times H] *_L B.$$

Assuming this, then it is clear that neither of the edge groups of the two edges incident to the commensuriser vertex of $\Gamma(G)$ contains H.

To justify the above claim about $\Gamma(G)$, we need to show that if C is a two-ended subgroup of G such that $e(G, C) \geqslant 2$, then C is conjugate to a subgroup of H, and any nontrivial C–almost invariant subset of G is enclosed by the commensuriser vertex of the above graph of groups. This can be shown topologically. Pick compact spaces with fundamental groups A, B, K and L, and use them to form a compact space with fundamental group G. Then consider a covering space with fundamental group C.

Next we give a specific example, which puzzled us for many years before we understood the theory of regular neighbourhoods. This is related to the problem of unfolding of splittings over two-ended subgroups, which appeared to make it very difficult to produce a truly canonical algebraic JSJ-decomposition. Our work in this paper solves this problem by showing how to enclose all splittings over two-ended subgroups simultaneously.

Example 11.3. — Let A and B be finitely generated groups, and let C and D be infinite cyclic subgroups of A and B respectively. Let nD denote the subgroup of D of index n. Let G denote the group $A *_{C=6D} B$, and let σ_6 denote this splitting of G over $6D$. For $k = 1$, 2 or 3, let A_k denote $A *_{C=6D} kD$. Then we can also express G as $A_k *_{kD} B$, for $k = 1$, 2 or 3. Let σ_1, σ_2 and σ_3 denote these three splittings of G over D, $2D$ and $3D$ respectively. The two splittings σ_2 and σ_3 of G must have non-zero intersection number, because otherwise they would be compatible by Theorem 2.35, and it is easy to see that this is impossible. We prove in the lemma below that the regular neighbourhood of σ_2 and σ_3 in G is the graph of groups Γ with two edges given by $A *_{C=6D} D *_D B$. This graph has a single V_0–vertex carrying D and two other vertices which are V_1–vertices. One way to prove this would be to check that the conditions in Definition 6.1 hold. Certainly the vertex of Γ which carries D does enclose each of the splittings σ_2 and σ_3. To see this for σ_2, observe that the graph of groups Γ_2 given by $G = A *_{C=6D} 2D *_{2D} D *_D B$ is the required refinement of Γ, and similarly for σ_3. Also Γ is minimal and the condition on isolated V_0–vertices is vacuous, because neither σ_2 nor σ_3 is isolated. But it is not easy to verify directly that Γ satisfies Condition 2) of the definition. Instead, we will directly consider the construction of the regular neighbourhood of σ_2 and σ_3 in G, which we gave in chapter 3. This is possible in this case, because we can directly understand the connections between the almost invariant sets associated to the four different splittings described above.

Lemma 11.4. — *If A, B, C, D, G and σ_k, for $k = 1$, 2, 3 or 6, are all as in Example 11.3, then the regular neighbourhood of σ_2 and σ_3 in G is the graph of groups Γ with two edges given by $A *_{C=6D} D *_D B$. This graph has a single V_0–vertex carrying D and two other vertices which are V_1–vertices.*

Proof. — For $k = 1$, 2, or 3, we let Z_k denote one of the standard almost invariant subsets of G associated to the splitting σ_k. We let Z denote one of the standard almost invariant subsets of G associated to the original splitting σ_6. Finally, let d denote a generator of D. Consider the G–tree T_2 determined by the graph of groups Γ_2. Let v be a vertex with stabiliser $2D$. There is one edge s incident to v with stabiliser $2D$ and three other edges incident to v each with stabiliser $6D$. If t denotes one of these three edges, then the other two equal d^2t and d^4t. We orient s towards v and t away from v, and pick any basepoint for T_2. We choose $Z_2 = Z_s$ and $Z = Z_t$. It is now immediate that $Z_2 = Z \cup d^2Z \cup d^4Z$. Similarly, considering Γ_3, shows that $Z_3 = Z \cup d^3Z$. We claim that the CCC v_0 of $\overline{E}$ which contains $\overline{Z_2}$ and $\overline{Z_3}$ consists precisely of $\overline{Z_2}$, $d^3\overline{Z_2}$, $\overline{Z_3}$, $d^2\overline{Z_3}$ and $d^4\overline{Z_3}$. Clearly these must all lie in v_0, and we are claiming that no other translates of Z_2 cross Z_3, and that no other translates of Z_3 cross Z_2. Assuming this, it follows that the stabiliser of v_0 is simply the group generated by the stabilisers of Z_2 and Z_3, namely $2D$ and $3D$, which is exactly D. Further, we claim that for any element a of $A-D$, the CCC's v_0 and av_0 are adjacent in the pretree, and that the analogous statement holds for any element of $B-D$. It follows that the regular neighbourhood of σ_2 and σ_3 in G has a single V_0–vertex with associated group D and has a V_1–vertex with associated group A and a V_1–vertex with associated group B. As these groups generate G and the regular neighbourhood is a minimal graph of groups, it follows that the regular neighbourhood must be the graph Γ described above.

Now we prove the first claim about the composition of the CCC v_0 of $\overline{E}$ which contains $\overline{Z_2}$ and $\overline{Z_3}$. This requires showing that no translate of Z_2 by an element of $G-D$ can cross Z_3. It will then follow immediately that no translate of Z_3 by an element of $G-D$ can cross Z_2. Let g be an element of $G-D$. We need to show that gZ_2 and Z_3 are nested. Considering T_2 shows that we have either $gZ_2^{(*)} \subset Z^*$ or $gZ_2^{(*)} \subset kZ^{(*)}$, where $kZ^{(*)}$ is not Z or Z^*, so that $k \notin D$. Considering T_3 shows that $Z^* \subset Z_3^*$ and $kZ^{(*)} \subset Z_3$, when $k \notin D$, so that gZ_2 and Z_3 are nested, as required. Hence the CCC v_0 contains the five elements claimed. Note that as v_0 contains $\overline{Z_2}$ and $\overline{Z_3}$, every CCC is a translate of v_0 by some element of G.

Next we prove that av_0 is adjacent to v_0, for any element a of $A-D$. If some CCC gv_0 lies between v_0 and av_0, then there is an element X of v_0 such that gX lies between $\overline{Z_2}$ and $a\overline{Z_2}$. It is easy to see that this is impossible, using the fact that a stabilises a vertex of T_2 adjacent to v. We prove similarly that bv_0 is adjacent to v_0, for any element b of $B-D$. This completes the proof of the lemma. $\square$

Now we come to the first of our general examples.

Example 11.5. — Let G be the fundamental group of an orientable Haken manifold. Recall from the discussion in chapter 1 that for our purposes we will consider the submanifold $V'(M)$ of M rather than the characteristic submanifold $V(M)$. The V_0–vertices of the decomposition Γ of chapter 10 applied to G essentially correspond to the peripheral components of $V'(M)$. However we get extra V_0–vertices corresponding to most of the annulus components in the frontier of the peripheral components of $V'(M)$. In fact, if S is an annulus component of the frontier of a peripheral component W of $V'(M)$, we get an extra V_0–vertex corresponding to S except in the case when W is homeomorphic to $S \times I$. To see this, observe that the peripheral components of $V'(M)$ are filled by essential annuli. (In most cases, these annuli can be chosen to be all embedded.) Moreover, we showed in [**45**], that the frontier components of $V(M)$ induce splittings of G which are 1–canonical. The non-peripheral components of $V'(M)$ do not enclose nontrivial almost invariant sets over two-ended subgroups, since this would give a splitting of a Seifert fibre space over a two-ended subgroup relative to its boundary. Thus the V_0–vertices of Γ correspond to the peripheral components of $V'(M)$ together with the extra annuli mentioned above. Note the reduced version of Γ does not have these extra V_0–vertices.

We will obtain an algebraic analogue of the whole submanifold $V'(M)$ in chapter 13.

We next compare the decompositions of arbitrary finitely presented groups obtained by Bowditch in section 15 of [**8**] with ours. The construction of Bowditch in [**8**] is in terms of axes and does not seem to give the enclosing properties that we described in the previous chapters. Also it is not clear whether his decomposition is independent of the axis chosen. It seems that different choices of axes give the same V_0–vertex groups, but the edge groups may be different. We give an example where the decomposition given in the previous chapter may differ from that in [**8**].

Example 11.6. — Start with a one-ended hyperbolic group K which does not have any splittings over two-ended subgroups and let L be the HNN-extension obtained by identifying two non-conjugate infinite cyclic subgroups H_1, H_2 of K. Let M be the product of a one-ended group with an infinite cyclic group H. Let G be the group obtained by amalgamating L and M along H_1 and H. If we take the axis corresponding to the final decomposition, Bowditch's construction again yields the decomposition $L *_{H_1=H} M$. However in our decomposition, there are two more V_0–vertices corresponding to the two splittings over H_i, $i = 1, 2$. This is more canonical, since it takes care of all splittings of G over virtually cyclic groups.

Recall from Example 11.1 that $\mathrm{Comm}_G(H)$ need not be finitely generated. In such a case, it is not at all clear how the edge groups in Bowditch's decomposition and ours correspond. It is possible that the decompositions obtained by Bowditch for big enough axes are the same as ours. Even if this is possible, it is not clear how to prove the enclosing property for almost invariant sets without going through some work similar to that in this paper.

Finally, we give an example from 3–manifold topology which motivates some of our later work. This is Example 2.13 of our paper [**45**].

Example 11.7. — Let F_1 and F_2 denote two compact surfaces each with at least two boundary components. Let Σ_i denote $F_i \times S^1$, let T_i denote a boundary component of Σ_i, and construct a 3–manifold M from Σ_1 and Σ_2 by gluing T_1 to T_2 so that the given fibrations by circles do not match. Let T denote the torus $\Sigma_1 \cap \Sigma_2$. Then T is a canonical torus in M, and the characteristic submanifold $V(M)$ of M has two components which are a copy of Σ_1 and a copy of Σ_2. Let H denote the subgroup of $G = \pi_1(M)$ carried by T. Then there is a splitting σ of G over H which has non-zero intersection number with the splitting τ determined by T.

The splitting σ is constructed as follows. Let G_i denote $\pi_1(\Sigma_i)$, and let C_i denote the subgroup of G_i carried by T_i. The starting point of our construction is that if F is a compact surface with at least two boundary components, and if S denotes a boundary circle of F, then S carries an infinite cyclic subgroup of $\pi_1(F)$ which is a free factor of $\pi_1(F)$. Now it is easy to give a splitting of $\pi_1(F)$ over $\pi_1(S)$, and hence a splitting of G_i over C_i. If each $\pi_1(F_i)$ is free of rank at least 3, then we can write $G_i = A_i *_{C_i} B_i$. If we let A denote the subgroup of G generated by A_1 and A_2, i.e. $A = A_1 *_H A_2$, and define B similarly, then we can express G as $A *_H B$, and it is easy to see that this splitting σ of G has non-zero intersection number with the splitting τ of G determined by T. If $\pi_1(F_i)$ has rank 2, then we can write $G_i = A_i *_{C_i}$ and a similar construction can be made.

The point of this example is the following. The fact that T is topologically canonical means that any essential annulus or torus in M has zero intersection number with T. But this example shows that the splitting τ determined by T is not algebraically 2–canonical. Now the natural next step after the results of the previous chapters would be to attempt to define an algebraic analogue of the characteristic submanifold of a 3–manifold to be the regular neighbourhood of all nontrivial almost invariant subsets of G which are over subgroups isomorphic to $\mathbb{Z}$ or to $\mathbb{Z} \times \mathbb{Z}$. Suppose that this can be done and consider the case when G is the fundamental group of the manifold M in the above example. Let Γ denote the regular neighbourhood. The fact that τ crosses another splitting over $\mathbb{Z} \times \mathbb{Z}$ implies that it cannot be an edge splitting of Γ, so that clearly Γ would not be the same as the topological JSJ-decomposition of M. In fact, we do not know whether there is such a regular neighbourhood. However, our results in [**45**] imply that the topological JSJ-decomposition is a reduced algebraic regular neighbourhood of all nontrivial almost invariant subsets over $\mathbb{Z}$ and of all 1–canonical nontrivial almost invariant subsets over $\mathbb{Z} \times \mathbb{Z}$, and this is what we will generalise in later chapters.

CHAPTER 12

CANONICAL DECOMPOSITIONS OVER VPC GROUPS OF A GIVEN LENGTH

As stated at the beginning of chapter 7, the analogues of the results of chapters 7 to 10 go through for almost invariant sets over $VPCn$ groups assuming that G does not have nontrivial almost invariant sets over VPC groups of length $< n$ (the analogue of Proposition 7.2 is Proposition 13.3, which we prove in the next chapter). Note that the results of [**13**] and [**15**] imply that the condition that G does not have nontrivial almost invariant sets over VPC groups of length $< n$ is equivalent to the condition that G does not split over such a subgroup. We will use these two conditions interchangeably.

We will need the following definitions.

Definition 12.1. — Let Γ be a minimal graph of groups decomposition of a group G. A vertex v of Γ is of *VPC–by–Fuchsian type* if $G(v)$ is a VPC-by-Fuchsian group, where the Fuchsian group is not finite nor two-ended, and there is exactly one edge of Γ which is incident to v for each peripheral subgroup K of $G(v)$ and this edge carries K. If the length of the normal VPC subgroup of $G(v)$ is n, we will say that $G(v)$ is of *$VPCn$–by–Fuchsian type.*

Note that if $G = G(v)$, then the Fuchsian quotient group corresponds to a closed orbifold. Note also that if $G(v)$ is of $VPCn$–by–Fuchsian type, then each peripheral subgroup of $G(v)$ is $VPC(n+1)$.

Definition 12.2. — Let Γ be a minimal graph of groups decomposition of a group G. A vertex v of Γ is *n–simple*, if whenever X is a nontrivial almost invariant subset of G over a VPC subgroup of length at most n such that X is enclosed by v, then X is associated to an edge splitting of Γ.

The following are the results which we obtain. To prove these results, we will need generalisations of the results in chapter 8 on the Boolean algebra $B(H)$ to the case when H is $VPCn$. Such results can be proved by the same methods. This will then allow us to handle the V_0–vertices of large commensuriser type in our regular neighbourhood. As in chapters 9 and 10, we first state the basic existence result for the appropriate regular neighbourhood and then list its properties. Recall that we use the term length instead of Hirsch length, for brevity.

Theorem 12.3. — *Let G be a one-ended, finitely presented group which does not split over VPC groups of length $< n$, and let $\mathcal{F}_n$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over $VPCn$ groups.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma_n = \Gamma(\mathcal{F}_n : G)$. Each V_0–vertex v of Γ_n satisfies one of the following conditions:

(1) *v is isolated, so that $G(v)$ is $VPCn$.*

(2) *$G(v)$ is of $VPC(n-1)$–by–Fuchsian type.*

(3) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some $VPCn$ subgroup H, such that $e(G,H) \geqslant 2$.*

Γ_n consists of a single vertex if and only if $\mathcal{F}_n$ is empty, or G itself satisfies one of the above three conditions.

Remark 12.4. — When $n = 1$, the decomposition Γ_1 is precisely the decomposition obtained in Theorem 10.1.

Now we list the properties of Γ_n.

Theorem 12.5. — *Let G be a one-ended, finitely presented group which does not split over VPC groups of length $< n$, and let $\mathcal{F}_n$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over $VPCn$ subgroups.*

Then the regular neighbourhood $\Gamma_n = \Gamma(\mathcal{F}_n : G)$ is a minimal bipartite graph of groups decomposition of G with the following properties:

(1) *each V_0–vertex v of Γ_n satisfies one of the following conditions:*

(a) *v is isolated, so that $G(v)$ is $VPCn$.*

(b) *$G(v)$ is of $VPC(n-1)$–by–Fuchsian type.*

(c) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some $VPCn$ subgroup H, such that $e(G,H) \geqslant 2$.*

Further, if H is a $VPCn$ subgroup such that $e(G,H) \geqslant 2$, and if H has large commensuriser, then Γ_n will have a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$.

(2) *If an edge of Γ_n is incident to a V_0–vertex of type a) or b) above, then it carries a $VPCn$ group.*

(3) *any representative of an element of $\mathcal{F}_n$ is enclosed by some V_0–vertex of Γ_n, and each V_0–vertex of Γ_n encloses such a subset of G. In particular, any splitting of G over a $VPCn$ subgroup is enclosed by some V_0–vertex of Γ_n.*

(4) *if X is a nontrivial almost invariant subset of G over a finitely generated subgroup H, and if X does not cross any element of $\mathcal{F}_n$, then X is enclosed by a V_1–vertex of Γ_n.*

(5) *if X is a H–almost invariant subset of G associated to a splitting of G over H, and if X does not cross any element of $\mathcal{F}_n$, then X is enclosed by a V_1–vertex of Γ_n.*

(6) *the V_1–vertices of Γ_n are n–simple. (See Definition 12.2.)*

(7) *If Γ_1 and Γ_2 are minimal bipartite graphs of groups structures for G which satisfy conditions 3) and 5) above, then they are isomorphic provided that there is a one-to-one correspondence between their isolated V_0–vertices, and that any non-isolated V_0–vertex of Γ_1 or Γ_2 encloses some non-isolated element of $\mathcal{F}_1$.*

(8) *The graph of groups Γ_n is invariant under the automorphisms of G.*

(9) *The n–canonical splittings of G over a $VPCn$ subgroup are precisely those edge splittings of Γ_n which are over such a subgroup. This includes, but need not be limited to, all those edges of Γ_n which are incident to V_0–vertices whose associated groups are of type a) or b) above.*

We noted in the introduction, that, for $n \geqslant 2$, the special case of Theorem 12.3 when G is a Poincaré duality group of dimension $n+1$ recovers the results of Kropholler in [**27**]. We will briefly discuss this. First we need to show that if G is a Poincaré duality group of dimension $n+1$, then G does not split over VPC groups of length $< n$, so that Theorem 12.3 applies. This follows from Theorem A of Kropholler and Roller in [**31**], which shows that if H is a VPC subgroup of G such that G splits over H, then H has length n. Now we consider the V_0–vertices of Γ_n. Those of isolated type and of $VPC(n-1)$–by–Fuchsian type correspond to vertices of Kropholler's decomposition, but his results do not mention vertices of commensuriser type. In order to complete our discussion, we consider the V_0–vertices of Γ_n of this type. If H is a PDn–subgroup of G, it has 2 coends in G. Hence Proposition 7.4 implies that nontrivial almost invariant subsets of G which are over PDn–subgroups of G cannot cross weakly. If $n = 1$, it follows from the proofs of Lemmas 7.16 and 8.6 that if v is a V_0–vertex group of commensuriser type, so that $G(v) = \mathrm{Comm}_G(H)$, then $G(v)$ contains the incident edge groups with index at most 3. If $n > 1$, we obtain the same conclusion by using generalisations of these two lemmas. Thus $G(v)$ is a finite torsion free extension of the PDn-group H. This implies that $G(v)$ is itself a PDn-group, and the same holds for the incident edge groups. Now Kropholler and Roller proved in Corollary A2 of [**29**] that if a $PD(n+1)$-group G splits over a PDn–subgroup K, and if K is contained in a PDn–subgroup L of G, then K has index at most 2 in L. It follows that either v is isolated, or that it has valence 1 and $G(v)$ contains the incident edge group K as a subgroup of index 2. This reconciles our decomposition with that of Kropholler. In a later paper, we plan to discuss this in more detail. We also plan to show how our ideas can be used to recover and generalise Kropholler's results for the case of PDn–pairs.

As in chapter 10.1, we can deduce from Theorem 12.3 that one can also form a regular neighbourhood of only those almost invariant subsets which are associated to splittings. This is the result we obtain.

Theorem 12.6. — *Let G be a one-ended, finitely presented group which does not split over VPC groups of length $< n$, and let $\mathcal{S}_n$ denote the collection of equivalence classes of all almost invariant subsets which are associated to splittings of G over $VPCn$ groups.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{S}_n : G)$. Each V_0–vertex v of $\Gamma(\mathcal{S}_n : G)$ satisfies one of the following conditions:

(1) *v is isolated, so that $G(v)$ is $VPCn$.*

(2) *$G(v)$ is of $VPC(n-1)$–by–Fuchsian type.*

(3) *$G(v)$ contains a $VPCn$ subgroup H, which it commensurises, such that $e(G, H) \geqslant 2$.*

If $\Gamma(\mathcal{S}_n : G)$ consists of a single vertex, then either $\mathcal{S}_n$ is empty, or G itself satisfies one of conditions 2) or 3) above.

Remark 12.7. — Note that even if G is finitely presented, Example 11.1 shows that a vertex group of type 3) need not be finitely generated. Note also that if G commensurises a $VPCn$ subgroup H, such that $e(G, H) \geqslant 2$, then Example 12.8 below shows that $\Gamma(\mathcal{S}_n : G)$ need not consist of a single vertex. This is in contrast with the situation of Theorem 12.5.

Example 12.8. — This is the natural generalisation to higher rank of Example 10.10. We let $G^n_{p,q} = A *_C B$, where A and B are both free abelian of rank n and C has index p in A and index q in B. Thus $G^n_{p,q}$ centralises, and hence commensurises, the $VPCn$ subgroup C. If $p, q \geqslant 2$, then $G^n_{p,q}$ splits over C, so that $e(G, C) \geqslant 2$. Note that $G^n_{p,q}$ is not VPC, except in the case $p = q = 2$. If we exclude the case $p = q = 2$, then we claim that $\Gamma(\mathcal{S}_n : G^n_{p,q})$ does not consist of a single vertex.

For simplicity, we present the proof of this claim on the assumption that the quotients A/C and B/C are cyclic. This allows us to apply directly results used in Example 10.10. However, the claim holds without this assumption, and its proof needs only a very minor generalisation of the argument.

To prove this claim, we argue as in the proof of Lemma 10.11, which deals with the case when $n = 1$. Now suppose that $n > 1$. Let σ denote the given splitting of $G^n_{p,q}$ over C, and suppose that $G^n_{p,q}$ has a splitting τ over a $VPCn$ subgroup H. The first paragraph of the proof of Lemma 10.11 implies that, as C is central in $G^n_{p,q}$, either it is contained in H, or H is normal in $G^n_{p,q}$ and the quotient group is isomorphic to $\mathbb{Z}$ or to $\mathbb{Z}_2 * \mathbb{Z}_2$. The second case would imply that $G^n_{p,q}$ was VPC, which is only possible in the case $p = q = 2$, which we are excluding. It follows that the splittings σ and τ of $G^n_{p,q}$ induce splittings σ' and τ' of $G^n_{p,q}/C$, where σ' is the free product splitting $\mathbb{Z}_p * \mathbb{Z}_q$, and τ' splits over the finite group H/C.

Now we proved in Lemma 10.11 that, if p and q are prime, then σ' and τ' are conjugate splittings of $\mathbb{Z}_p * \mathbb{Z}_q$. This implies that σ and τ are conjugate splittings of $G^n_{p,q}$. Thus, up to conjugacy, $G^n_{p,q}$ has a unique splitting over a $VPCn$ subgroup. We also proved for general values of p and q, that σ' and τ' are compatible splittings of $G^n_{p,q}$. As τ is arbitrary, this implies that σ is isolated among all splittings of $G^n_{p,q}$ over a $VPCn$ subgroup. As in Example 10.10, this implies that $\Gamma(\mathcal{S}_n : G^n_{p,q})$ does not consist of a single vertex, as claimed.

CHAPTER 13

CANONICAL DECOMPOSITIONS OVER VPC GROUPS OF TWO SUCCESSIVE LENGTHS

As usual, let G denote a one-ended, finitely presented group. In this chapter, we examine the problem of enclosing almost invariant subsets of G over VPC subgroups of two successive lengths n and $n+1$. Note that VPC groups of length at most 2 are virtually abelian. As discussed at the end of chapter 11 in the case when $n = 1$, we should not expect to be able to enclose all almost invariant subsets of G over VPC groups of lengths 1 and 2. However, if G is the fundamental group of a Haken 3–manifold M, and $V'(M)$ denotes the submanifold of M which we discussed in chapter 1, then our results in [**45**] imply that the graph of groups structure for G determined by $V'(M)$ is an algebraic regular neighbourhood of all nontrivial almost invariant subsets over $\mathbb{Z}$ and of all 1–canonical almost invariant subsets over $\mathbb{Z} \times \mathbb{Z}$. (Note that a k–canonical almost invariant set is always nontrivial, by Definition 2.36.) Thus, in the case $n = 1$, we will show for general G that one can enclose all nontrivial almost invariant subsets over two-ended subgroups together with all 1–canonical almost invariant subsets over $VPC2$ subgroups. This essentially corresponds to the classical JSJ-decomposition. If $n > 1$, we will need to assume that G does not have any nontrivial almost invariant subsets over VPC groups of length $< n$, and we will then show that one can enclose all nontrivial almost invariant subsets of G over $VPCn$ subgroups together with all n–canonical almost invariant subsets over $VPC(n+1)$ subgroups.

Here is a more detailed statement of what we will do in this chapter. Assume that G does not have any nontrivial almost invariant subsets over VPC groups of length $< n$. Recall from Theorem 12.3, that $\mathcal{F}_n$ denotes the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over $VPCn$ groups, and Γ_n denotes the regular neighbourhood $\Gamma(\mathcal{F}_n : G)$ of $\mathcal{F}_n$ in G. Let T_n denote the universal covering G–tree of Γ_n. Now we enlarge $\mathcal{F}_n$ to the set $\mathcal{F}_{n,n+1}$ which consists of $\mathcal{F}_n$ together with the equivalence classes of all n–canonical almost invariant subsets of G which are over a $VPC(n+1)$ subgroup. We will construct the regular neighbourhood

$\Gamma_{n,n+1} = \Gamma(\mathcal{F}_{n,n+1} : G)$ by refining Γ_n by splitting at some of its V_1–vertices. This is natural because a n–canonical almost invariant subset of G over a $VPC(n+1)$ subgroup does not cross any element of $\mathcal{F}_n$, and so must be enclosed by some V_1–vertex of T_n, by part 1 of Proposition 5.7. Note that if two such subsets cross, they must be enclosed by the same V_1–vertex of T_n.

Before starting on the main argument, we prove two facts about VPC groups which we will use frequently. These results are not new. See, for example, [**30**] and [**15**].

Lemma 13.1. — *Let K be a $VPC(n+1)$ group, with a subgroup L of length n. Then the number of coends of L in K is 2 and moreover there is a subgroup L' of finite index in L such that L' has infinite index in its normaliser.*

Proof. — The number of coends of L in K is 2 since both are virtual Poincaré duality groups. Now, without loss of generality, we can assume that they are both Poincaré duality groups. In order to find the required subgroup L', we use the fact that L can be separated from elements of $K - L$. If K and L are both orientable, then the number of ends of the pair (K, L) is 2. It follows from Scott's Theorem 4.1 in [**40**] on ends of pairs of groups that some subgroup K' of finite index in K splits over L. Now the fact that the number of ends of the pair (K', L) is 2 implies that L is normal in K' and $L\backslash K'$ has 2 ends. In particular, L has infinite index in its normaliser, so that L' equals L in this case. If L is not orientable, we let L' denote the orientable subgroup of index two, and the preceding argument shows that L' has infinite index in its normaliser, as required. □

The same argument shows the following result.

Lemma 13.2. — *Suppose that L is a $VPCn$ group and that L is a subgroup of two $VPC(n+1)$ groups K_1 and K_2. Then there is a subgroup L' of finite index in L and subgroups L_1 and L_2 of finite index in K_1 and K_2 respectively so that $L'\backslash L_1$ and $L'\backslash L_2$ are both infinite cyclic.*

Next we need to consider crossing of n–canonical almost invariant subsets of G which are over $VPC(n+1)$ subgroups. First we consider strong crossing of such subsets. The following proposition and proof are suggested by the symmetry of crossings proved in [**43**].

Proposition 13.3. — *Let G be a one-ended, finitely generated group, and let X and Y be n–canonical almost invariant subsets of G over $VPC(n+1)$ subgroups H and K. If X crosses Y strongly, then Y crosses X strongly and the number of coends in G of both H and K is 2.*

Remark 13.4. — Proposition 7.2 is the special case of this result when $n = 0$.

Proof. — Let Δ be the Cayley graph of G with respect to some finite system of generators. As $H\backslash\delta X$ is finite, there is a finite connected subcomplex of $H\backslash\Delta$ which contains $H\backslash\delta X$ and carries the group H. The pre-image of this subcomplex in Δ is a connected subcomplex C which contains δX and is H–finite. Similarly, there is a connected subcomplex D of Δ which contains δY and is K–finite. Since X crosses Y strongly, there are points of δX, and hence of C, in Y and in Y^* which are outside any d–neighbourhood of δY. Thus the projection of C into $K\backslash\Delta$ has at least two ends, so that $e(H, H\cap K)\geqslant 2$. It follows that $H\cap K$ is a $VPCn$ group.

Lemma 13.2 tells us that we can find a subgroup L of finite index in $H\cap K$ so that L has infinite index in its normalisers in both H and K. We claim that $L\backslash\Delta$ has one end. For suppose that there is a nontrivial L–almost invariant subset Z of G. Since X and Y are n–canonical, we have one of the four inequalities $Z^{(*)}\leqslant X^{(*)}$ and one of the four inequalities $Z^{(*)}\leqslant Y^{(*)}$. By appropriately replacing some of X, Y and Z by their complements, we may arrange that $Z\leqslant X$ and $Z\leqslant Y$, which implies that $Z\leqslant X\cap Y$. As L has infinite index in its normaliser $N_H(L)$ in H, there is an infinite cyclic subgroup J of $N_H(L)/L$ which acts freely on $L\backslash\Delta$. As X crosses Y strongly, the orbit of any point of $L\backslash\Delta$ under the action of J contains points on each side of $L\backslash\delta Y$ which are arbitrarily far from $L\backslash\delta Y$. As $L\backslash\delta Z$ is finite, there is an element j of J such that $j(L\backslash\delta Z)$ is contained in $L\backslash Y^*$. Thus there is an element h of $N_H(L)$ such that $h\delta Z$ is contained in Y^*. Thus we have one of the inclusions $hZ^{(*)}\subset Y^*$. Suppose that $hZ\subset Y^*$. As $Z\leqslant X$ and $hX=X$, we have $hZ\leqslant X\cap Y^*$. As h normalises L, it follows that $Z\cup hZ$ is a nontrivial L–almost invariant subset of G. As $Z\leqslant Y$ and $hZ\leqslant Y^*$, this set crosses Y which contradicts our assumption that Y is n–canonical. If $hZ^*\subset X\cap Y^*$, we consider $Z\cup hZ^*$ instead, to complete the proof of the claim that $L\backslash\Delta$ has one end.

Recall that H and K are $VPC(n+1)$, that $H\cap K$ and hence L have length n, and that L has infinite index in its normaliser. Thus both the covers $L\backslash\Delta\to H\backslash\Delta$ and $L\backslash\Delta\to K\backslash\Delta$ have two-ended covering groups. Let C_L and D_L be the images of C and D respectively in $L\backslash\Delta$, so that both C_L and D_L are two-ended.

Our hypothesis implies that $H\backslash\Delta$ has more than one end. Now we claim that $H\backslash\Delta$ has only two ends. For suppose that $H\backslash\Delta$ has more than two ends. Since $e(H, H\cap K)=2$, the image of D in $H\backslash\Delta$ has two ends. Thus an end of $H\backslash\Delta$ is free of the image of D. Choose a compact set separating this end from the image of D and let M be an infinite component of the complement of this compact set which does not intersect the image of D. Let N be a component of the pre-image of M in $L\backslash\Delta$, so that N is disjoint from D_L. Since $L\backslash\Delta$ has only one end, the coboundary of N is not finite. Thus the stabiliser of δN is an infinite subgroup of the covering transformation group of the cover $L\backslash\Delta\to H\backslash\Delta$. As this covering group is two-ended, the stabiliser of δN must have finite index in it. Thus δN and C_L lie in a bounded neighbourhood of each other. Recall that as X crosses Y strongly, C has points in at least two components of $\Delta - D$, which are arbitrarily far from D. Hence C_L has points in at least two components of $(L\backslash\Delta)-D_L$, which are arbitrarily far from D_L. Thus there are points of δN, and hence of N, in at least two components of $(L\backslash\Delta) - D_L$, which are arbitrarily far from D_L. This is a contradiction since D_L and N are disjoint and are both connected. It follows that the number of ends of $H\backslash\Delta$ is 2, as claimed.

Now Y crosses X since crossing is symmetric, and Proposition B.3.7 shows that Y crosses X strongly. We can repeat the above argument for subgroups of finite index in H and K to conclude that both H and K have two coends in G, as required. $\square$

We now consider n–canonical almost invariant subsets of G which cross weakly.

Proposition 13.5. — *Let G be a one-ended, finitely generated group without nontrivial almost invariant subsets over VPC groups of length $< n$. Let H and K be $VPC(n+1)$ subgroups of G, and let X and Y be nontrivial n–canonical subsets of G over H and K respectively. Suppose that X crosses Y weakly. Then H and K are commensurable.*

Remark 13.6. — Proposition 7.3 is the special case of this result when $n = 0$.

Proof. — By Proposition 13.3, we know that Y also crosses X weakly. Thus one of $\delta Y \cap X^{(*)}$ is H–finite and one of $\delta X \cap Y^{(*)}$ is K–finite. By replacing X and Y by their complements as needed, we can arrange that $\delta Y \cap X$ is H–finite and $\delta X \cap Y$ is K–finite. As δX is H–finite and δY is K–finite, it follows that $\delta Y \cap X$ and $\delta X \cap Y$ are both H–finite and K–finite. Thus they are $(H \cap K)$–finite. Now consider $X \cap Y$. Every edge in $\delta(X \cap Y)$ lies in $\delta Y \cap X$ or $\delta X \cap Y$. It follows that $\delta(X \cap Y)$ is also $(H \cap K)$–finite, so that $X \cap Y$ is $(H \cap K)$–almost invariant.

If H and K are not commensurable, then $H \cap K$ has infinite index in both H and K. In particular, $H \cap K$ has length $\leqslant n$. Suppose first that $H \cap K$ has length $< n$. As G has no nontrivial almost invariant subsets over VPC subgroups of length $< n$, it follows that $X \cap Y$ is $(H \cap K)$–finite, and hence H–finite, which contradicts our hypothesis that X and Y cross. Now suppose that $H \cap K$ has length n, so that $e(H, H \cap K) = 2$. As in the proof of Proposition 13.3, we can translate $X \cap Y$ by an element h of H such that h commensurises $H \cap K$ and $h(X \cap Y)$ is contained in $X \cap Y^*$. Thus $(X \cap Y) \cup h(X \cap Y)$ is a nontrivial almost invariant set over a subgroup commensurable with $H \cap K$ which crosses Y. This contradicts the assumption that Y is n–canonical. Hence H and K are commensurable as required. $\square$

Note that the proof of Proposition 13.3 did not use the hypothesis that G has no nontrivial almost invariant subsets over VPC groups of length $< n$. But the proof of Proposition 13.5 used this hypothesis in an essential way. It was used to exclude the case when $H \cap K$ has length $< n$. If this occurs, we might not be able to construct a nontrivial almost invariant set which crosses Y, because there need not be any suitable elements of G which commensurise $H \cap K$.

Remark 13.7. — We note that if X and Y are n–canonical and their stabilisers are commensurable, then $X \cap Y$, $X + Y$, $X \cup Y$ are again n–canonical.

These results show that as in Proposition 7.5, we have:

Proposition 13.8. — *Let G be a one-ended, finitely generated group and let $\{X_\lambda\}_{\lambda\in\Lambda}$ be a family of nontrivial almost invariant subsets over $VPCn$ subgroups and of n–canonical almost invariant subsets over $VPC(n+1)$ subgroups. As usual, let E denote the set of all translates of the X_λ's and their complements. Form the pretree P of cross-connected components (CCC's) of $\overline{E}$ as in the construction of regular neighbourhoods in chapter 3. Then the following statements hold:*

(1) *The crossings in a CCC of $\overline{E}$ are either all strong or are all weak.*

(2) *In a CCC with all crossings weak, the stabilisers of the corresponding elements of E are all commensurable. In a CCC with all crossings strong, the stabilisers of the corresponding elements of E have 2 coends in G.*

Recall that T_n denotes the universal covering G–tree of the regular neighbourhood Γ_n of $\mathcal{F}_n$, the equivalence classes of all nontrivial almost invariant subsets of G over $VPCn$ subgroups. Recall also that if X is a n–canonical H–almost invariant subset of G, then it is enclosed by some V_1–vertex v of T_n. In our construction of $\Gamma_{n,n+1}$, the commensuriser of H will play an important role, just as in the construction of Γ_n. In particular, it will be important to know that the commensurisers of H in $\mathrm{Stab}(v)$ and in G are equal. This is the last part of the following proposition.

Proposition 13.9. — *Let G be a one-ended, finitely presented group without nontrivial almost invariant subsets over VPC groups of length $< n$. Let X and Y be n–canonical subsets of G over $VPC(n+1)$ groups, and let H and K denote the stabilisers of X and Y respectively.*

(1) *Then X is enclosed by a unique V_1–vertex v_X of T_n.*

(2) *If H and K are commensurable, then v_X and v_Y are equal.*

(3) *Let v denote the V_1–vertex of T_n which encloses X. Then $\mathrm{Comm}_{\mathrm{Stab}(v)}(H) = \mathrm{Comm}_G(H)$.*

Proof. — Recall that a n–canonical almost invariant subset of G over a $VPC(n+1)$ subgroup must be enclosed by some V_1–vertex of T_n, by Proposition 5.7. Now part 1) is the special case of part 2) obtained when Y equals X, so we will prove part 2). Let v_1 and v_2 be V_1–vertices of T_n which enclose X and Y respectively, and suppose that v_1 and v_2 are distinct. Then there is a V_0–vertex v separating v_1 from v_2. This implies that there is a nontrivial almost invariant subset Z of G over a $VPCn$ group L such that Z is enclosed by v and $\overline{Z}$ lies between $\overline{X}$ and $\overline{Y}$. Note that the stabilisers of X and Y are $VPC(n+1)$ groups and the stabiliser of Z is a VPC group of length n. Since H and K are commensurable, δX and δY lie in a bounded neighbourhood of each other. In $L\backslash\Delta$, where Δ is the Cayley graph of G, the images of δX and of δY are non-compact and lie essentially on different sides of the image $L\backslash\delta Z$ which is compact. This contradicts the fact that δX and δY lie in a bounded neighbourhood of each other. This contradiction shows that v_1 must equal v_2, as required.

For the third part of the proposition, let g be an element of $\mathrm{Comm}_G(H)$. Then gX is an almost invariant subset of G over H^g, which is commensurable with H, and gX is enclosed by the V_1–vertex gv. Part 2) of the proposition tells us that $gv = v$, showing that we must have $\mathrm{Comm}_{\mathrm{Stab}(v)}(H) = \mathrm{Comm}_G(H)$, as required. □

Propositions 13.3 and 13.5 are the generalisations to higher lengths of Propositions 7.2 and 7.3. We will also need generalisations of some of the other results from chapters 7 and 8. We will sketch what is involved.

In [**13**] and [**15**], the algebraic Torus Theorem is proved for a finitely generated group G which has a nontrivial almost invariant subset over a $VPCn$ subgroup, on the assumption that G has no such subsets over a VPC subgroup of length $\leqslant n-1$. Their work dealt with the case when $n \geqslant 1$. Stallings structure theorem for groups with more than one end [**52**][**53**] deals with the case when $n = 0$. We will need to apply arguments from [**13**] and [**15**] to cases where G may have nontrivial almost invariant subsets over a VPC subgroup of length $\leqslant n-1$. Here are two of the results which we will need. The first is essentially in step 2 of [**13**] (see also [**35**] and Proposition B.2.14).

Theorem 13.10. — *Let G be a finitely generated group and let X be a nontrivial almost invariant subset over a $VPCn$ group H. Suppose that whenever gX crosses X, then g commensurises H. Then G splits over a subgroup commensurable with H.*

The arguments for the second result are essentially in [**15**]. Their results in section 4 of that paper are only for splittings and are not formulated in terms of regular neighbourhoods. However their methods do apply to almost invariant subsets in general, and we formulate the result in terms of regular neighbourhoods. This is the generalisation to higher lengths of Theorem 7.8 of this paper.

Theorem 13.11. — *Let G be a one-ended finitely generated group with a finite family of $VPCn$ subgroups $\{H_\lambda\}_{\lambda\in\Lambda}$. For each $\lambda \in \Lambda$, let X_λ denote a $(n-1)$–canonical H_λ–almost invariant subset of G, let E denote the set of all translates of the X_λ's and their complements, and let Γ denote the regular neighbourhood of the X_λ's.*

Let X denote an element of E, let H denote its stabiliser, and let v denote a vertex of Γ which encloses X. Suppose that each H_λ has two coends in G and that there exists an element of E which crosses X. Then the vertex group $G(v)$ is of $VPC(n-1)$–by–Fuchsian type, and H is not commensurable with a peripheral subgroup of $G(v)$.

We note that both the above results imply that G splits over a $VPCn$ subgroup. Theorem 13.10 yields a splitting over a subgroup commensurable with H. Theorem 13.11 yields a splitting over each peripheral subgroup of $G(v)$. (But no such subgroup is commensurable with H.) Thus, as in our proof of Theorem 10.7, the algebraic

Annulus Theorem, these two results yield a proof of the algebraic Torus Theorem. More importantly from our point of view, these two results lead naturally to relative versions of the algebraic Torus Theorem. In Theorem 13.10, suppose there is some nontrivial almost invariant set Y over a finitely generated subgroup of G which does not cross X or any of its translates. Then Y also does not cross the almost invariant sets associated with the splitting given by this theorem. This is because the almost invariant set associated to the splitting is obtained by taking successive intersections of translates of X or X^*. In Theorem 13.11, suppose there is some nontrivial almost invariant set Y over a finitely generated subgroup of G which does not cross any element of E. Then Proposition 5.7 shows that Y is enclosed by some V_1–vertex of Γ, and hence does not cross the almost invariant sets associated with any of the splittings over peripheral subgroups of $G(v)$.

In our construction of the regular neighbourhood $\Gamma_{n,n+1} = \Gamma(\mathcal{F}_{n,n+1} : G)$ from the regular neighbourhood $\Gamma_n = \Gamma(\mathcal{F}_n : G)$, we use the fact that n–canonical almost invariant sets are enclosed by some V_1–vertex v of Γ_n. In the case when G is finitely presented and we are considering n–canonical almost invariant subsets of G which cross weakly, we will use our accessibility results, Theorems 7.11 and 7.13, and Theorem 13.10 in two different ways. Suppose that we are in the situation of Proposition 13.5 or Proposition 14.2. If H has small commensuriser in G then the number of coends of H in G is finite. If H has large commensuriser in G, then the Boolean algebra $B(H)$ is finitely generated over $\mathrm{Comm}_G(H)$. The proofs are as before using the second accessibility result, Theorem 7.13.

Now we state the main results of this chapter. As usual we state first the existence result.

Theorem 13.12. — *Let G be a one-ended, finitely presented group which does not split over VPC subgroups of length $< n$, and let $\mathcal{F}_{n,n+1}$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a $VPCn$ subgroup, together with the equivalence classes of all n–canonical almost invariant subsets of G which are over a $VPC(n+1)$ subgroup.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma_{n,n+1} = \Gamma(\mathcal{F}_{n,n+1} : G)$.

Each V_0–vertex v of $\Gamma_{n,n+1}$ satisfies one of the following conditions:

(1) *v is isolated, so that $G(v)$ is VPC of length n or $n+1$.*

(2) *v is of VPC-by-Fuchsian type, where the VPC group has length $n-1$ or n.*

(3) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some VPC subgroup H of length n or $n+1$, such that $e(G,H) \geqslant 2$.*

$\Gamma_{n,n+1}$ consists of a single vertex if and only if $\mathcal{F}_{n,n+1}$ is empty, or G itself satisfies one of the above three conditions.

Now we list the properties of the decomposition $\Gamma_{n,n+1}$.

Theorem 13.13. — *Let G be a one-ended, finitely presented group which does not split over VPC subgroups of length $< n$, and let $\mathcal{F}_{n,n+1}$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a $VPCn$ subgroup, together with the equivalence classes of all n–canonical almost invariant subsets of G which are over a $VPC(n+1)$ subgroup.*

Then the regular neighbourhood $\Gamma_{n,n+1} = \Gamma(\mathcal{F}_{n,n+1} : G)$ is a minimal bipartite graph of groups decomposition of G with the following properties:

(1) *each V_0–vertex v of $\Gamma_{n,n+1}$ satisfies one of the following conditions:*

(a) *v is isolated, so that $G(v)$ is VPC of length n or $n+1$.*

(b) *v is of VPC–by–Fuchsian type, where the VPC group has length $n-1$ or n.*

(c) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some VPC subgroup H of length n or $n+1$, such that $e(G,H) \geqslant 2$.*

Further if H is a $VPCn$ subgroup such that $e(G,H) \geqslant 2$, and if H has large commensuriser, then $\Gamma_{n,n+1}$ will have a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$. The same holds if H is $VPC(n+1)$, so long as there exists a n–canonical almost invariant subset of G over H.

(2) *If an edge of Γ is incident to a V_0–vertex of type a) or b) above, then it carries a VPC group of length n or $n+1$, as appropriate.*

(3) *any representative of an element of $\mathcal{F}_{n,n+1}$ is enclosed by some V_0–vertex of $\Gamma_{n,n+1}$, and each V_0–vertex of $\Gamma_{n,n+1}$ encloses such a subset of G. In particular, any splitting of G over a $VPCn$ subgroup, and any n–canonical splitting of G over a $VPC(n+1)$ subgroup is enclosed by some V_0–vertex of $\Gamma_{n,n+1}$.*

(4) *if X is a nontrivial almost invariant subset of G over a finitely generated subgroup H, and if X does not cross any element of $\mathcal{F}_{n,n+1}$, then X is enclosed by a V_1–vertex of $\Gamma_{n,n+1}$.*

(5) *if X is a H–almost invariant subset of G associated to a splitting of G over H, and if X does not cross any element of $\mathcal{F}_{n,n+1}$, then X is enclosed by a V_1–vertex of $\Gamma_{n,n+1}$.*

(6) *the V_1–vertices of $\Gamma_{n,n+1}$ are $(n+1)$–simple. In particular, $\Gamma_{n,n+1}$ cannot be further refined by splitting at a V_1–vertex along a VPC group of length $\leqslant n+1$.*

(7) *If Γ_1 and Γ_2 are minimal bipartite graphs of groups structures for G which satisfy conditions 3) and 5) above, then they are isomorphic provided that there is a one-to-one correspondence between their isolated V_0–vertices, such that corresponding vertices have stabilisers of the same length, and that any non-isolated V_0–vertex of Γ_1 or Γ_2 encloses some non-isolated element of $\mathcal{F}_{n,n+1}$.*

(8) *The graph of groups $\Gamma_{n,n+1}$ is invariant under the automorphisms of G.*

(9) *For $k = n$ or $n+1$, the k–canonical splittings of G over a $VPCk$ subgroup are precisely those edge splittings of $\Gamma_{n,n+1}$ which are over such a subgroup. This includes, but need not be limited to, all those edges of $\Gamma_{n,n+1}$ which are incident to V_0–vertices whose associated groups are of types a) or b) above.*

Proof. — The only new point which arises is in the proof of part 6). Suppose that X is a nontrivial almost invariant subset of G over a VPC subgroup H of length at most $n+1$ such that X is enclosed by a V_1–vertex of $\Gamma_{n,n+1}$. This implies that X crosses no element of $\mathcal{F}_{n,n+1}$, so that X crosses no nontrivial almost invariant subset of G over a VPC subgroup of length at most n. It follows that X represents an element of $\mathcal{F}_{n,n+1}$, and so is enclosed by some V_0–vertex of $\Gamma_{n,n+1}$. It follows from part 1) of Lemma 4.9 that X is associated to an edge splitting of $\Gamma_{n,n+1}$, as required. □

If we specialise to the case $n = 1$, and apply this result to the fundamental group of a Haken 3–manifold M, then the V_0–vertices of $\Gamma_{1,2}$ essentially correspond to the components of the submanifold $V'(M)$, which we discussed in chapter 1. The only difference is that $\Gamma_{1,2}$ has extra V_0–vertices corresponding to most of the components of the frontier of $V'(M)$. In fact, if S is a component of the frontier of a component W of $V'(M)$, we get an extra V_0–vertex corresponding to S except in the case when W is homeomorphic to $S \times I$. To see this, observe that the peripheral components of $V'(M)$ have enough immersions of the annulus to make them cross-connected, and the interior components of $V'(M)$ have enough immersions of the torus to make them cross-connected. Moreover, we showed in [**45**], that the frontier components of $V'(M)$ induce splittings of G which are all 1–canonical. This is similar to the discussion in chapter 11 for the case of the canonical decomposition obtained in chapter 10. Note that, as before, the reduced version of $\Gamma_{1,2}$ does not have these extra V_0–vertices.

We already saw from Example 11.7 that if there is a regular neighbourhood of *all* the nontrivial almost invariant subsets of G which are over VPC subgroups of length 1 or 2, then it cannot be a refinement of Γ_1. Thus there may be nontrivial almost invariant subsets of G which are over $VPC2$ subgroups but are not enclosed by any V_0–vertex of $\Gamma_{1,2}$. However, we prove below that the stabiliser of such an almost invariant subset of G must be 'almost' conjugate into a V_0–vertex of $\Gamma_{1,2}$.

In order to discuss V_0–vertices, it will be helpful to introduce some new language. For the graph of groups $\Gamma_{n,n+1}$, there is a natural idea of the level of a V_0–vertex v. If v appears only after refining Γ_n, then v has level $n+1$. Otherwise v has level n. In this second case, it is natural to think of v as belonging to Γ_n, in some sense, but as there is no map from Γ_n to $\Gamma_{n,n+1}$, this is not very precise. A more accurate way to describe the level of a V_0–vertex is to use the projection map from $\Gamma_{n,n+1}$ to Γ_n, which is part of the definition of a refinement. This map sends each edge of $\Gamma_{n,n+1}$ either to an edge or to a V_1–vertex of Γ_n. Then a V_0–vertex of $\Gamma_{n,n+1}$ has level $n+1$ if it is sent to a V_1–vertex, and has level n otherwise.

Proposition 13.14. — *Let G be a one-ended, finitely presented group which does not split over VPC subgroups of length $< n$, and let $\Gamma_{n,n+1}$ denote the regular neighbourhood of the previous theorem. Let X be any nontrivial almost invariant subset of G over a $VPC(n+1)$ subgroup H. Then either X represents an element of $\mathcal{F}_{n,n+1}$, and so is enclosed by some V_0–vertex of $\Gamma_{n,n+1}$, or some subgroup of finite index in H is* *conjugate into the vertex group of a V_0–vertex of $\Gamma_{n,n+1}$ which is of large commensuriser type and level n. In the second case, G has a $VPCn$ subgroup A, such that $e(G,A) \geqslant 2$, and there is a V_0–vertex v of $\Gamma_{n,n+1}$ which is of large commensuriser type such that $G(v) = \mathrm{Comm}_G(A)$, and some subgroup of finite index in H is conjugate to a subgroup of $G(v)$ which contains A.*

Proof. — If X does not cross any nontrivial almost invariant subset Y over a $VPCn$ subgroup K, then X is n–canonical, and so represents an element of $\mathcal{F}_{n,n+1}$. Otherwise X crosses such a set Y. If X crosses Y strongly, then the first paragraph of the proof of Proposition 13.3 shows that $H \cap K$ must have length n and hence be of finite index in K. Now Lemma 13.1 tells us that a subgroup of finite index in H commensurises K. It follows that K has large commensuriser so that $\Gamma_{n,n+1}$ has a V_0–vertex group which equals $\mathrm{Comm}_G(K)$, and so contains a subgroup of finite index in H as required. We can take the group A to be $H \cap K$.

Now suppose that X crosses Y weakly. If Y also crosses X weakly, then the first paragraph of the proof of Proposition 13.5 shows that one of $X^{(*)} \cap Y^{(*)}$, say W, is almost invariant over $H \cap K$. As X and Y cross, W will be a nontrivial almost invariant set over $H \cap K$. It follows that the length of $H \cap K$ cannot be less than n, since G does not have any nontrivial almost invariant subsets over VPC subgroups of length less than n. Thus $H \cap K$ has length n, and we can apply the arguments in the preceding paragraph. Note that X does not cross W. We simply need the fact that H contains $H \cap K$.

The only remaining case is when X crosses Y weakly and Y crosses X strongly. In this case, let $L = H \cap K$. By replacing K by a subgroup of finite index, we may assume that L is normal in K and that $L\backslash K$ is infinite cyclic. Since X crosses Y weakly one of $\delta X \cap Y^{(*)}$ is K–finite. We will assume that $\delta X \cap Y$ is K–finite. Now by again replacing K by a subgroup of finite index, we may assume that for a generator k of $L\backslash K$, the translates of $\delta X \cap Y$ by the powers of k do not intersect. We choose k so that $k(\delta X \cap Y) \subset X \cap Y$, and consider the set $Z = X \cap Y \cap kX^*$. This set is almost invariant over L. As L has length $< n$, any L–almost invariant subset of G is trivial. Thus Z is L–finite. In particular, Z lies within a finite distance of δY. As $\cup_{i \geqslant 1} k^i Z = X \cap Y$ and k preserves δY, it follows that $X \cap Y$ also lies within a finite distance of δY, contradicting the hypothesis that X and Y cross. This contradiction completes the proof. □

Finally, as in chapters 10 and 12, it follows that one can also form a regular neighbourhood of only those almost invariant subsets which are associated to splittings. This is the result we obtain.

Theorem 13.15. — *Let G be a one-ended, finitely presented group which does not split over VPC subgroups of length $< n$, and let $\mathcal{S}_{n,n+1}$ denote the collection of equivalence classes of all almost invariant subsets which are associated to a splitting of G over* *a $VPCn$ subgroup, together with the equivalence classes of all n–canonical almost invariant subsets which are associated to a splitting of G over a $VPC(n+1)$ subgroup.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{S}_{n,n+1} : G)$.

Each V_0–vertex v of $\Gamma(\mathcal{S}_{n,n+1} : G)$ satisfies one of the following conditions:

(1) v is isolated, so that $G(v)$ is VPC of length n or $n+1$.

(2) v is of VPC-by-Fuchsian type, where the VPC group has length $n-1$ or n.

(3) $G(v)$ contains a VPC subgroup H of length n or $n+1$, which it commensurises, such that $e(G,H) \geqslant 2$.

If $\Gamma(\mathcal{S}_{n,n+1} : G)$ consists of a single vertex, then either $\mathcal{S}_{n,n+1}$ is empty, or G itself satisfies one of the above three conditions.

Remark 13.16. — Note that even if G is finitely presented, Example 11.1 shows that a vertex group of type 3) need not be finitely generated. Note also that if G itself satisfies the third condition, then the example below shows that $\Gamma(\mathcal{S}_{n,n+1} : G)$ need not consist of a single vertex. This is in contrast with the situation of Theorem 13.13.

Example 13.17. — This is the natural generalisation of Example 12.8. As in that example, we let $G^n_{p,q} = A *_C B$, where A and B are both free abelian of rank n and C has index p in A and index q in B. Thus $G^n_{p,q}$ centralises, and hence commensurises, the $VPCn$ subgroup C. If $p, q \geqslant 2$, then G splits over C, so that $e(G,C) \geqslant 2$. If we exclude the case $p = q = 2$, we will show that $\Gamma(\mathcal{S}_{n,n+1} : G^n_{p,q})$ does not consist of a single vertex. Let σ denote the given splitting of $G^n_{p,q}$ over C. Now suppose that $G^n_{p,q}$ has a n–canonical splitting over a $VPC(n+1)$ subgroup H. Then Theorem 5.16 implies that this splitting is compatible with the splitting σ, so that, in particular, H must be isomorphic to a subgroup of A or B, which is impossible. We conclude that there are no n–canonical splittings of $G^n_{p,q}$ over a $VPC(n+1)$ subgroup. Thus $\mathcal{S}_{n,n+1}$ equals the set $\mathcal{S}_n$ of Example 12.8, so that $\Gamma(\mathcal{S}_{n,n+1} : G^n_{p,q})$ equals $\Gamma(\mathcal{S}_n : G^n_{p,q})$ which does not consist of a single vertex.

CHAPTER 14

CANONICAL DECOMPOSITIONS OVER VIRTUALLY ABELIAN GROUPS

Here is a summary of what we achieved in the last chapter. Consider a one-ended, finitely presented group G which does not have any nontrivial almost invariant subsets over VPC groups of length $< n$. Let Γ_n denote the regular neighbourhood of $\mathcal{F}_n$, the equivalence classes of all nontrivial almost invariant subsets of G over $VPCn$ subgroups, and let $\mathcal{F}_{n,n+1}$ denote $\mathcal{F}_n$ together with the equivalence classes of all n–canonical almost invariant subsets of G which are over a $VPC(n+1)$ subgroup. We showed that there is a regular neighbourhood $\Gamma_{n,n+1}$ of $\mathcal{F}_{n,n+1}$ which is a refinement of Γ_n obtained by splitting Γ_n at some of its V_1–vertices.

The natural next step would be to let $\mathcal{F}_{n,n+1,n+2}$ denote $\mathcal{F}_{n,n+1}$ together with the equivalence classes of all $(n+1)$–canonical almost invariant subsets of G which are over a $VPC(n+2)$ subgroup, and show that $\mathcal{F}_{n,n+1,n+2}$ has a regular neighbourhood $\Gamma_{n,n+1,n+2}$ which is a refinement of $\Gamma_{n,n+1}$ obtained by splitting $\Gamma_{n,n+1}$ at some of its V_1–vertices. However, the following example for the case $n = 1$ shows that this cannot be done following the pattern of the previous results. On the other hand, we will show in this chapter, that such refinements always exist if we restrict our attention to almost invariant sets over virtually abelian groups, and that the process can be repeated up to any given rank. This seems to indicate that there may be geometric differences between splittings over VPC groups and virtually abelian groups.

We should emphasise that the following example does not show that $\mathcal{F}_{1,2,3}$ has no regular neighbourhood. We do not know whether this is true. This example shows simply that if $\mathcal{F}_{1,2,3}$ has a regular neighbourhood, then it cannot be constructed from $\Gamma_{1,2}$ by splitting at V_1–vertices.

Example 14.1. — This is an example of a one-ended group G with incommensurable polycyclic subgroups H and K of length 3, and 2–canonical almost invariant sets X and Y over H and K respectively which cross weakly. Thus Proposition 13.5 cannot be generalised and there is no hope of enclosing X and Y in a V_0–vertex group with stabiliser equal to the commensuriser of H or K.

We start with an extension of $\mathbb{Z} \times \mathbb{Z}$ by $\mathbb{Z}$ which is given by an automorphism of $\mathbb{Z} \times \mathbb{Z}$ with no real eigenvalues. This gives us a polycyclic group H and we denote a lift of $\mathbb{Z}$ into H by C_1. Note that H is the fundamental group of a closed 3–manifold M which is a bundle over the circle with fibre the torus. Our choice of H implies that any polycyclic subgroup of length 2 is contained in the normal $\mathbb{Z} \times \mathbb{Z}$. We let K denote a second copy of H and let C_2 denote the subgroup of K corresponding to C_1. Let L denote the fundamental group of a hyperbolic surface F with one boundary component and denote the subgroup corresponding to ∂F by C_3. Now we amalgamate H, K and L along the C_i's to obtain the desired group G, and denote by C the identified copies of the C_i's. Thus G is the fundamental group of a space Z which is the union of two copies of M and the surface F. Let Γ denote the associated graph of groups structure for G, which is a tree with four vertices carrying the subgroups C, H, K and L. Consider the subgroups $H *_C K$, $H *_C L$, $K *_C L$ of G. Then G can be obtained from the first two groups by amalgamating over H. We let X denote one of the standard H–almost invariant subsets of G associated to this splitting. Similarly, the first and third groups give an amalgamated free product decomposition of G over K. We let Y denote the corresponding K–almost invariant subset of G.

Clearly H and K are not commensurable in G. Also it is clear that the above splittings of G over H and K are not compatible, so that X and Y must cross. As $H \cap K = C$, and $e(H, C) = e(K, C) = 1$, the splittings cannot cross strongly, so that X and Y must cross weakly. It remains to show that X and Y are 2–canonical. We will do this by showing that they are 1–canonical and that G has no nontrivial almost invariant subsets over any $VPC2$ subgroups.

We claim that if W is a nontrivial almost invariant subset of G over a two-ended subgroup A, then W is enclosed by the vertex of Γ which carries L. Assuming this, it follows that X and Y are 1–canonical as required. The claim can be seen by simply considering the covering space of Z corresponding to a two-ended subgroup A. It follows that the regular neighbourhood of all the nontrivial almost invariant subsets of G over two-ended subgroups has a single V_0–vertex of finite-by-Fuchsian type with associated group L, has no V_0–vertices of commensuriser type and has three isolated V_0–vertex groups which carry C. Collapsing the edges which carry C will yield Γ.

Now any length two polycyclic subgroup of G is conjugate into H or K, and any such subgroup of H or K is contained in the normal $\mathbb{Z} \times \mathbb{Z}$ subgroup of H or of K. It is now easy to check, by considering the covering space of Z corresponding to the fibre torus of a copy of M, that G has no nontrivial almost invariant subsets over $VPC2$ subgroups. It follows that the splittings of G over H and K that we considered above are 2–canonical, as required.

The above example shows that the process of refining our algebraic analogues of the JSJ-decomposition is not possible over VPC groups for more than two successive lengths even if we take i–canonical sets over VPC groups of length $(i + 1)$ at each stage. However, it is possible for virtually abelian groups and we will indicate the necessary changes to the arguments.

Two of the crucial properties we needed to obtain canonical decompositions in the previous chapter were contained in the following two propositions which we reproduce here for the reader's convenience.

Proposition 13.3. — *Let G be a one-ended, finitely generated group, and let X and Y be n–canonical subsets of G over $VPC(n+1)$ subgroups H and K. If X crosses Y strongly, then Y crosses X strongly and the number of coends in G of both H and K is* 2.

Proposition 13.5. — *Let G be a one-ended, finitely generated group without nontrivial almost invariant subsets over VPC groups of length $< n$. Let H and K be $VPC(n+1)$ subgroups of G, and let X and Y be nontrivial n–canonical subsets of G over H and K respectively. Suppose that X crosses Y weakly. Then H and K are commensurable.*

Note that in Proposition 13.3 we only needed the almost invariant sets to be n–canonical whereas in Proposition 13.5 we excluded the existence of nontrivial almost invariant sets over VPC groups of length $< n$. The example at the end of the previous chapter showed that the analogue of Proposition 13.5 is not true in general. However, the following analogue holds when we restrict our attention to virtually abelian subgroups of G.

Proposition 14.2. — *Let G be a one-ended, finitely generated group, and let H and K be virtually abelian subgroups of G of rank $n+1$. Let X and Y be almost invariant subsets of G over H and K respectively which are n–canonical with respect to abelian groups. Suppose that X crosses Y weakly. Then H and K are commensurable.*

Proof. — Our argument is based on the proof of Proposition 13.5. As in the first part of that proof, we know that X and Y cross each other weakly and that $X \cap Y$ is $(H \cap K)$–almost invariant.

If H and K are not commensurable, then $H \cap K$ has infinite index in both H and K. In particular, $H \cap K$ has rank $\leqslant n$. As H is virtually abelian, there is h in H of infinite order which commutes with a subgroup of $H \cap K$ of finite index. Thus h commensurises $H \cap K$. Further, by replacing h by a suitable power, we can arrange that $h(X \cap Y)$ is contained in $X \cap Y^*$, as in the proof of Proposition 13.3. Thus $(X \cap Y) \cup h(X \cap Y)$ is a nontrivial almost invariant subset of G over a subgroup of H commensurable with $H \cap K$ which crosses Y. This contradicts the assumption that Y is n–canonical with respect to abelian groups. Hence H and K are commensurable as required. □

Note that in the proof of Proposition 13.5, we proceeded essentially as above in the case when $H \cap K$ had length n, but we eliminated the possibility that $H \cap K$ had length $< n$ by using the assumption that G had no nontrivial almost invariant subsets over VPC subgroups of length $< n$. In the case above, G may have such subsets. Instead we used the assumption that H and K are virtually abelian, and applied the same argument as when $H \cap K$ had length n.

With this proposition available, and the other results discussed in the previous chapter, there is no difficulty in extending the main decomposition theorems to almost invariant sets over virtually abelian groups up to any rank. We will need the following definitions. In this chapter, it will be convenient to use the notation VA for a virtually abelian group of finite rank, and VAn for a virtually abelian group of rank n.

Definition 14.3. — Let Γ be a minimal graph of groups decomposition of a group G. A vertex v of Γ is of *VA–by–Fuchsian type* if $G(v)$ is an extension of a VA group by a Fuchsian group, where the Fuchsian group is not finite nor two-ended, and there is exactly one edge of Γ which is incident to v for each peripheral subgroup K of $G(v)$ and this edge carries K. If the rank of the normal VA subgroup of $G(v)$ is n, we will say that $G(v)$ is of *VAn–by–Fuchsian type.*

Note that if $G = G(v)$, then the Fuchsian quotient group corresponds to a closed orbifold. Note also that if $G(v)$ is of VAn–by–Fuchsian type, then each peripheral subgroup of $G(v)$ is $VPC(n+1)$ but need not be $VA(n+1)$. See Example 14.7.

Definition 14.4. — Let Γ be a minimal graph of groups decomposition of a group G. A vertex v of Γ is *n–simple for abelian groups*, if whenever X is a nontrivial almost invariant subset of G over a VA subgroup of rank at most n such that X is enclosed by v, then X is associated to an edge splitting of Γ.

The results we obtain follow. The proof consists of starting with the graph of groups structure $\Gamma_{1,2}$ described in the previous chapter, and then using the methods of that chapter to repeatedly refine it by splitting at V_1–vertices. We will need to use Proposition 14.2 in place of Proposition 13.5 because we are now in the virtually abelian case. As usual, we state the existence result and then list the properties of the decomposition obtained.

Theorem 14.5. — *Let G be a one-ended, finitely presented group. Let $\mathcal{F}_{1,2,\ldots,n}$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a virtually abelian subgroup of rank i, for $1 \leqslant i \leqslant n$, and are $(i-1)$–canonical with respect to abelian groups.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma_{1,2,\ldots,n} = \Gamma(\mathcal{F}_{1,2,\ldots,n} : G)$.

Each V_0–vertex v of $\Gamma_{1,2,\ldots,n}$ satisfies one of the following conditions:

(1) *v is isolated, so that $G(v)$ is VA of rank $\leqslant n$.*

(2) *$G(v)$ is of VA-by-Fuchsian type, where the VA group has rank $k < n$.*

(3) *$G(v)$ is the full commensuriser* $\mathrm{Comm}_G(H)$ *for some VA subgroup H of rank at most n, such that $e(G,H) \geqslant 2$.*

$\Gamma_{1,2,\ldots,n}$ consists of a single vertex if and only if $\mathcal{F}_{1,2,\ldots,n}$ is empty, or G itself satisfies one of the above three conditions.

Now we list the properties of $\Gamma_{1,2,\ldots,n}$.

Theorem 14.6. — *Let G be a one-ended, finitely presented group. Let $\mathcal{F}_{1,2,\ldots,n}$ denote the collection of equivalence classes of all nontrivial almost invariant subsets of G which are over a virtually abelian subgroup of rank i, for $1 \leqslant i \leqslant n$, and are $(i-1)$–canonical with respect to abelian groups.*

Then the regular neighbourhood $\Gamma_{1,2,\ldots,n} = \Gamma(\mathcal{F}_{1,2,\ldots,n} : G)$ is a minimal bipartite graph of groups decomposition of G with the following properties:

(1) *each V_0–vertex v of $\Gamma_{1,2,\ldots,n}$ satisfies one of the following conditions:*

(a) *v is isolated, so that $G(v)$ is VA of rank $\leqslant n$.*

(b) *$G(v)$ is of VAk–by–Fuchsian type, where $k < n$.*

(c) *$G(v)$ is the full commensuriser $\mathrm{Comm}_G(H)$ for some VA subgroup H of rank at most n, such that $e(G,H) \geqslant 2$.*

Further if H is a VA subgroup of G of rank $k \leqslant n$ such that $e(G,H) \geqslant 2$, if H has large commensuriser and if there exists a nontrivial H–almost invariant subset of G which is $(k-1)$–canonical with respect to abelian groups, then $\Gamma_{1,2,\ldots,n}$ has a V_0–vertex v such that $G(v) = \mathrm{Comm}_G(H)$.

(2) *If an edge of $\Gamma_{1,2,\ldots,n}$ is incident to a V_0–vertex of type a) or b) above, then it carries a VPC group of some length at most n.*

(3) *any representative of an element of $\mathcal{F}_{1,2,\ldots,n}$ is enclosed by some V_0–vertex of $\Gamma_{1,2,\ldots,n}$, and each V_0–vertex of $\Gamma_{1,2,\ldots,n}$ encloses such a subset of G. In particular, if $1 \leqslant i \leqslant n$, then any $(i-1)$–canonical splitting of G over a VAi subgroup is enclosed by some V_0–vertex of $\Gamma_{1,2,\ldots,n}$.*

(4) *if X is a nontrivial almost invariant subset of G over a finitely generated subgroup H, and if X does not cross any element of $\mathcal{F}_{1,2,\ldots,n}$, then X is enclosed by a V_1–vertex of $\Gamma_{1,2,\ldots,n}$.*

(5) *if X is a H–almost invariant subset of G associated to a splitting of G over H, and if X does not cross any element of $\mathcal{F}_{1,2,\ldots,n}$, then X is enclosed by a V_1–vertex of $\Gamma_{1,2,\ldots,n}$.*

(6) *the V_1–vertices of $\Gamma_{1,2,\ldots,n}$ are n–simple. In particular, $\Gamma_{1,2,\ldots,n}$ cannot be further refined by splitting at a V_1–vertex along a virtually abelian subgroup of rank at most n.*

(7) *If Γ_1 and Γ_2 are minimal bipartite graphs of groups structures for G which satisfy conditions 3) and 5) above, then they are isomorphic provided that there is a one-to-one correspondence between their isolated V_0–vertices, such that corresponding vertices have stabilisers of the same rank, and that any non-isolated V_0–vertex of Γ_1 or Γ_2 encloses some non-isolated element of $\mathcal{F}_{1,2,\ldots,n}$.*

(8) *The graph of groups $\Gamma_{1,2,\ldots,n}$ is invariant under the automorphisms of G.*

(9) *For $k \leqslant n$, the splittings of G over a VAk subgroup which are k–canonical with respect to abelian groups, are precisely those edge splittings of $\Gamma_{1,2,\dots,n}$ which are over such a subgroup.*

In case 1b), where v is a V_0–vertex such that $G(v)$ is of VAk–by–Fuchsian type, the edge groups of v need not be $VA(k+1)$. This may seem somewhat surprising but is another aspect of the general phenomenon that when one encloses almost invariant sets over subgroups of a certain type, the edge groups obtained need not be of the same type. Here is a specific example which is closely related to an example we gave at the end of chapter 1.

Example 14.7. — For each $k \geqslant 2$, this is an example of a finitely presented group G such that $\Gamma(\mathcal{F}_{1,2,\dots,k+1} : G)$ has a V_0–vertex v which is of VAk–by–Fuchsian type, but none of the edge groups incident to v is $VA(k+1)$.

Let F denote the compact surface obtained from a disc by removing the interiors of two disjoint discs. Let a and b denote the generators of $\pi_1(F)$ carried by two boundary components of F, oriented so that the third component carries ab. Thus ab^{-1} is carried by a figure eight loop γ in F. Note that F is a (topological) regular neighbourhood of γ.

We will construct G as the fundamental group of a graph Γ of groups whose underlying graph is a tree with three edges all incident to one vertex v.

Fix an integer $k \geqslant 2$, and let A denote a free abelian group of rank k. Choose an automorphism φ of A of infinite order, and let $G(v)$ denote the extension of A by $\pi_1(F)$ in which a and b act on A by φ. It follows that ab^{-1} centralises A, and so determines a subgroup H of $G(v)$ isomorphic to $\mathbb{Z}^{k+1}$. Thus $G(v)$ is the fundamental group of a bundle W over F with fibre the k–torus, and the pre-image in W of the figure eight loop γ in F is a $(k+1)$–torus T which carries H. Further W is a topological regular neighbourhood of T. The three peripheral subgroups of $G(v)$ are extensions of A by $\mathbb{Z}$ for which the defining automorphism of A is φ or φ^2, so that none of these groups is VA. These subgroups of $G(v)$ are the three edge groups of Γ. Thus v is of VAk–by–Fuchsian type. To complete the construction of G, we attach to each remaining vertex the direct product of $\mathbb{Z}^2$ with the incident edge group.

We prove below in Lemma 14.9 that if G admits a splitting over a VPC subgroup of length $i \leqslant k+1$, then $i = k+1$ and the splitting is conjugate to one of the edge splittings of Γ. As we have noted before, the results of [**13**] and [**15**] imply that G has a splitting over some VPC subgroup of length $\leqslant k$ if and only if it has a nontrivial almost invariant subset over such a subgroup. It follows that G possesses no nontrivial almost invariant subset over a VPC subgroup of length $i \leqslant k$. Thus a H–almost invariant subset X of G associated to the torus T is automatically k–canonical for abelian groups and so lies in $\mathcal{F}_{1,2,\dots,k+1}$. Now X crosses some translate of itself strongly, so that the algebraic regular neighbourhood of $\mathcal{F}_{1,2,\dots,k+1}$ must have a V_0–vertex of VAk–by–Fuchsian type which encloses X. The fact that the only splittings of G over $VPC(k+1)$ subgroups are the edge splittings of Γ shows that this V_0–vertex must have the same edge and vertex groups as v. It follows that $\Gamma(\mathcal{F}_{1,2,\ldots,k+1} : G)$ has a V_0–vertex v which is of VAk–by–Fuchsian type, but none of the edge groups incident to v is $VA(k+1)$. One can show further that Γ is equal to $\Gamma(\mathcal{F}_{1,2,\ldots,k+1} : G)$, but we omit the argument.

In order to prove Lemma 14.9, we start by observing that A is normal in G. This convenient fact allows us to work with the quotient $G' = G/A$. Clearly G' is the fundamental group of a graph Γ' of groups with the same underlying graph as Γ, with the group attached to v being $\pi_1(F)$, the other vertex groups being isomorphic to $\mathbb{Z}^3$, and the edge groups all infinite cyclic. Further the edge groups are the peripheral subgroups of $\pi_1(F)$ so that v is of Fuchsian type in Γ'.

Lemma 14.8. — *If G' is the fundamental group of the graph of groups Γ' just described, then the three edge splittings of Γ' are the only splittings of G' over two-ended subgroups up to conjugacy.*

Proof. — Recall that the peripheral subgroups of $\pi_1(F)$ are generated by the elements a, b and ab. We denote the $\mathbb{Z}^3$ vertex groups which contain these elements by G_a, G_b and G_{ab} respectively, and denote the edge splittings corresponding to these edges by σ_a, σ_b and σ_{ab}.

Suppose that G' has a splitting σ over a two-ended subgroup H', so that G' acts on a tree T' with quotient a single edge and with edge stabilisers conjugate to H'. As every subgroup of H' is two-ended or finite, $\mathbb{Z}^3$ cannot split over any subgroup of H'. It follows that each of G_a, G_b and G_{ab} must fix a vertex of T'. Hence the three peripheral subgroups of $\pi_1(F)$ each fix vertices of T'. It follows that $\pi_1(F)$ also fixes a vertex of T'. For otherwise, the action of $\pi_1(F)$ on T' would yield a splitting of $\pi_1(F)$ in which each peripheral subgroup lies in a vertex group, and standard topological methods show that this is impossible. In particular, it follows that G' is generated by vertex stabilisers. If σ were a HNN extension, the subgroup of G' generated by all the vertex stabilisers would be the kernel of a surjection from G' to $\mathbb{Z}$. We conclude that σ cannot be a HNN extension, so that the quotient $G'\backslash T'$ has a single edge and two vertices.

Note that none of G_a, G_b, G_{ab} and $\pi_1(F)$ can be isomorphic to a subgroup of a two-ended group. Thus none of these groups can fix an edge of T', so that each must fix a unique vertex of T'. Let w denote the vertex of T' fixed by $\pi_1(F)$. As G' cannot fix a vertex of T', one of G_a, G_b and G_{ab}, say G_a, must fix a vertex u of T' where $u \neq w$. It follows that a fixes each edge of the path λ in T' between u and w, so that H' contains some conjugate of a. As H' is virtually infinite cyclic, it cannot also contain any conjugate of b or of ab, as a conjugate of a together with one of b or of ab would generate a free subgroup of H' of rank 2. It follows that G_b and G_{ab} must fix w. By conjugating σ, we can suppose that H' contains a. As H' contains the subgroup generated by a with finite index, our expression of G' as the fundamental group of Γ' shows that H' must be conjugate into a vertex group of Γ', and hence must be generated by a.

If u is adjacent to w in T', it will follow that σ must be conjugate to the splitting σ_a. If u is not adjacent to w, let z denote a vertex of T' in the interior of λ. As a fixes every edge of λ, it must fix two distinct edges l and m incident to z. As each vertex of $G'\backslash T'$ has valence 1, there is g in G' such that $gz = z$ and $gl = m$. As the stabilisers of l and of m are each generated by a, it follows that g must normalise a. But the normaliser of a in G' is G_a which fixes u. Thus g must also fix u, and hence every edge between u and z, which contradicts the fact that $gl = m$. This contradiction shows that if G_a does not fix w, then σ is conjugate to σ_a. If G_b or G_{ab} do not fix w, the same argument shows that σ is conjugate to σ_b or to σ_{ab}, so that σ must be conjugate to one of the three edge splittings of Γ' as required. □

Lemma 14.9. — *If G is constructed as the fundamental group of a graph Γ of groups as in Example 14.7, and if G admits a splitting over a VPC subgroup of length $i \leqslant k+1$, then $i = k+1$ and the splitting is conjugate to one of the edge splittings of Γ.*

Proof. — Suppose that G admits a splitting over a VPC subgroup K of length $i \leqslant k+1$. Thus G acts on a tree T with quotient a single edge, and with edge stabilisers conjugate to K. As A is abelian, either it fixes a vertex of T, or it has an axis l. This means that there is a line l in T on which elements of A act trivially or by translations. Suppose that A fixes some vertex v of T. As A is normal in G, it follows that A fixes all the vertices in the G–orbit of v, and hence fixes some edge of T. It follows that A lies in K, in this case. Now suppose that A has axis l. As A is normal in G, it follows that $T = l$. As the symmetry group of l is $\mathbb{Z}_2 * \mathbb{Z}_2$, there is a natural homomorphism from G to $\mathbb{Z}_2 * \mathbb{Z}_2$, with kernel K, which is clearly a contradiction. We conclude that A must lie in K. It follows that G/A splits over K/A. As K/A is VPC of length at most 1 and it is easy to see that $e(G/A) = 1$, it follows that K/A must have length 1. Now the preceding lemma shows that the splitting of G/A over K/A is conjugate to one of the edge splittings of Γ'. It follows that the given splitting of G over K must be conjugate to one of the edge splittings of Γ, as required. □

Recall from Example 11.7 that if there is a regular neighbourhood of all the nontrivial almost invariant subsets of G which are over VA subgroups of length 1 or 2, then it cannot be a refinement of Γ_1. Thus there may be nontrivial almost invariant subsets of G which are over $VA2$ subgroups but are not enclosed by any V_0–vertex of $\Gamma_{1,2}$. There are similar examples for higher rank groups. However, we prove below that, as in Lemma 13.14, the stabiliser of any nontrivial almost invariant subset of G which is over a VA subgroup of rank at most n must be 'almost' conjugate into a V_0–vertex of $\Gamma_{1,2,\ldots,n}$. As for that lemma, it will be helpful to have an idea of the level of a V_0–vertex. A V_0–vertex of $\Gamma_{1,2,\ldots,n}$ has level n if it is sent to a V_1–vertex of $\Gamma_{1,2,\ldots,n-1}$ by the refinement projection, and has level $< n$ otherwise. This allows an inductive definition of the level of any V_0–vertex of $\Gamma_{1,2,\ldots,n}$.

Proposition 14.10. — *Let G be a one-ended, finitely presented group and let $\Gamma_{1,2,\ldots,n}$ denote the regular neighbourhood of the previous theorem. Let X be any nontrivial almost invariant subset of G over a VA subgroup H of rank $l+1 \leqslant n$. Then either X represents an element of $\mathcal{F}_{1,2,\ldots,n}$, and so is enclosed by some V_0–vertex of $\Gamma_{1,2,\ldots,n}$, or some subgroup of finite index in H is conjugate into the vertex group of a V_0–vertex of $\Gamma_{1,2,\ldots,n}$ which is of large commensuriser type and level $< n$. In the second case, there is an abelian subgroup A of G of rank $< n$, such that $e(G,A) \geqslant 2$, and a V_0–vertex v of $\Gamma_{1,2,\ldots,n}$ which is of large commensuriser type such that $G(v) = \mathrm{Comm}_G(A)$, and some subgroup of finite index in H is conjugate to a subgroup of $G(v)$ which contains A.*

Proof. — The proof is based on the proof of Proposition 13.14. We will argue by induction on l. The induction starts when $l = 1$, and this is just the case $n = 1$ of Proposition 13.14. It suffices to consider only the case when H is abelian, so we will assume this during our proof.

Now we will assume that the proposition holds for nontrivial almost invariant subsets of G which are over a VA subgroup of rank at most l. If X does not cross any nontrivial almost invariant subset of G which is over a VA subgroup of rank at most l, then X is l–canonical and so lies in $\mathcal{F}_{1,2,\ldots,n}$. Otherwise X crosses some nontrivial K–almost invariant subset Y of G, where K is VA of rank $\leqslant l$, and we will choose Y so as to minimise the rank of K.

If X crosses Y strongly, then $H \cap K$ must have rank l, and hence K has rank l and $H \cap K$ has finite index in K. Now we apply our induction hypothesis. If Y lies in $\mathcal{F}_{1,2,\ldots,n}$, we use the fact that, as H is abelian, it centralises $H \cap K$. Thus $H \cap K$ has large commensuriser, and there is a V_0–vertex w of $\Gamma_{1,2,\ldots,n}$ which is of large commensuriser type such that $G(w) = \mathrm{Comm}_G(H \cap K)$. As $H \cap K$ has rank $l < n$, this proves the required result about H in this case. If Y does not lie in $\mathcal{F}_{1,2,\ldots,n}$, we apply our induction hypothesis. This implies that, after simplifying by a conjugation, there is an abelian subgroup A of G of rank $< l$, such that $e(G,A) \geqslant 2$, and a V_0–vertex v of $\Gamma_{1,2,\ldots,n}$ which is of large commensuriser type such that $G(v) = \mathrm{Comm}_G(A)$, and some subgroup of finite index in K is a subgroup of $G(v)$ which contains A. Now we simply note that H centralises $H \cap K$ which has a subgroup of finite index which contains a subgroup of A of finite index. Thus H commensurises A, and so lies in $G(v)$.

If X and Y cross each other weakly, then the first paragraph of the proof of Proposition 13.5 shows that one of $X^{(*)} \cap Y^{(*)}$, call it W, is almost invariant over $H \cap K$. As X and Y cross, W must be a nontrivial $(H \cap K)$–almost invariant subset of G. As G is one-ended, $H \cap K$ must have rank at least 1. Now we apply the argument of the preceding paragraph with W in place of Y, to prove the required result. Note that X does not cross W. We simply need the fact that H contains $H \cap K$.

Finally suppose that X crosses Y weakly but Y crosses X strongly. As in the proof of Proposition 13.14, we let $L = H \cap K$, and replace K by a subgroup of finite index to arrange that L is normal in K and that $L\backslash K$ is infinite cyclic. Since X crosses Y weakly, we can assume that $\delta X \cap Y$ is finite. By again replacing K by a subgroup of finite index, we may assume that for a generator k of $L\backslash K$, the translates of $\delta X \cap Y$ by the powers of k do not intersect. Consider the set $Z = X \cap Y \cap kX^*$. This set is almost invariant over L. If Z is L–finite, then it lies within a finite distance of δY. As $\cup_{i \geqslant 1} k^i Z = X \cap Y$, it follows that $X \cap Y$ lies within a finite distance of δY, contradicting the hypothesis that X and Y cross. If Z is not L–finite, then X crosses $Z \cup k^{-1}Z$, and as L has lesser rank than K, this contradicts our choice of Y so as to minimise the rank of K. □

Next we discuss the behaviour of the sequence $\Gamma_{1,2,\ldots,n}$ of graphs of groups structures for a fixed group G as n increases. For brevity, we will denote $\Gamma_{1,2,\ldots,n}$ by Γ^n in this paragraph only. Each Γ^n is a refinement of Γ^{n-1}. We would like to consider whether this sequence stabilises by applying Theorem 7.13, but this theorem does not apply because only those edge groups incident to V_0–vertices of types a) or b) need to be VPC. Instead we argue as follows. Each V_0–vertex of Γ^n which is not a vertex of Γ^{n-1} encloses at least one splitting over a VAn subgroup. By picking one such splitting each time Γ^n and Γ^{n-1} are distinct, we obtain a sequence σ_k of compatible splittings of G over VA subgroups of rank at least k. This yields a sequence of graphs of groups structures Δ_k for G whose edge splittings are precisely $\sigma_1, \ldots, \sigma_k$. If the sequence Γ^n does not stabilise, the sequence Δ_k will be infinite. This will not contradict Theorem 7.13, but this can only occur if there is a subsequence of the σ_i's, say τ_j, such that τ_j is a splitting of G over a VA subgroup C_j of rank at least j such that $C_j \subset C_{j+1}$, for all j. Further the sequence Γ^n stabilises apart from such subsequences of edge splittings, so there can only be finitely many such subsequences. In the case when there are no such sequences, then the sequence $\Gamma_{1,2,\ldots,n}$ of graphs of groups decompositions of G eventually stabilises yielding a decomposition which we denote by Γ_∞, whose V_0–vertices enclose all nontrivial almost invariant subsets of G over any finitely generated subgroup which is virtually abelian.

Finally, as in chapters 10, 12 and 13, it follows that one can also form a regular neighbourhood of only those almost invariant subsets which are associated to splittings. This is the result we obtain.

Theorem 14.11. — *Let G be a one-ended, finitely presented group. Let $\mathcal{S}_{1,2,\ldots,n}$ denote the collection of equivalence classes of all almost invariant subsets of G which are associated to splittings over a virtually abelian subgroup of rank i, for $1 \leqslant i \leqslant n$, and are $(i-1)$–canonical with respect to abelian groups.*

Then the regular neighbourhood construction of chapter 3 works and yields a regular neighbourhood $\Gamma(\mathcal{S}_{1,2,\ldots,n} : G)$.

Each V_0–vertex v of $\Gamma(\mathcal{S}_{1,2,\ldots,n} : G)$ satisfies one of the following conditions:

(1) *v is isolated, so that $G(v)$ is VA of rank $\leqslant n$.*

(2) *$G(v)$ is of VAk-by-Fuchsian type, for some k such that $1 \leqslant k \leqslant n-1$.*

(3) *$G(v)$ contains a VA subgroup H of rank at most n, which it commensurises, such that $e(G,H) \geqslant 2$.*

If $\Gamma(\mathcal{S}_{1,2,\dots,n} : G)$ consists of a single vertex, then either $\mathcal{S}_{1,2,\dots,n}$ is empty, or G itself satisfies one of the above three conditions.

As in examples 10.10, 12.8 and 13.17, the following example shows that if G itself satisfies the third condition of the above theorem, then $\Gamma(\mathcal{S}_{1,2,\dots,n} : G)$ need not consist of a single vertex.

Example 14.12. — Consider the group $G_{p,q}$ of Example 10.10. We claim that $\Gamma(\mathcal{S}_{1,2,\dots,n} : G_{p,q}) = \Gamma(\mathcal{S}_1 : G_{p,q})$, which we showed in Example 10.10 does not consist of a single vertex. Suppose that $G_{p,q}$ has a splitting over a VAi subgroup H, for some $i \geqslant 2$, which is $(i-1)$–canonical with respect to abelian groups. Such a splitting must be compatible with the initial splitting of $G_{p,q}$ over C, so that H must be isomorphic to a subgroup of A or B, which is impossible. It follows that $\mathcal{S}_{1,2,\dots,n}$ equals the set $\mathcal{S}_1$ of equivalence classes of all nontrivial almost invariant subsets of G which are associated to splittings over a two-ended subgroup. Thus $\Gamma(\mathcal{S}_{1,2,\dots,n} : G_{p,q}) = \Gamma(\mathcal{S}_1 : G_{p,q})$ as required.

Note that this example is very special as $G_{p,q}$ contains no VAi subgroups at all for $i \geqslant 3$, although it does contain subgroups isomorphic to $\mathbb{Z} \times \mathbb{Z}$.

CHAPTER 15

PREVIOUS DECOMPOSITIONS OVER TWO-ENDED SUBGROUPS

In this chapter, we will discuss the relationship between the JSJ-decompositions of previous authors and the canonical decomposition which we described in chapter 10. Let G be a one-ended, finitely presented group. Let Γ denote the decomposition of G which we obtained in Theorem 10.1, so that Γ is a regular neighbourhood of all equivalence classes of nontrivial almost invariant subsets of G which are over two-ended subgroups. It is clear that this is not usually the same as any of the JSJ-decompositions of previous authors because it may have edge groups which are not two-ended. Such edge groups can only occur for edges which are incident to V_0–vertices which are of large commensuriser type, and we will now describe how to alter Γ so as to obtain one of these other decompositions. Recall that each V_0–vertex v of Γ which is of large commensuriser type encloses at least one splitting of G over a two-ended subgroup. For each such V_0–vertex v of Γ, we pick a maximal family of compatible splittings of G each of which is over a two-ended group, is enclosed by v, and is not an edge splitting for an edge which is incident to v. This is possible by the accessibility result in Theorem 7.11. We refine Γ by splitting at each such vertex using all these splittings. The resulting graph of groups structure is no longer canonical, as the splittings enclosed by v are not usually unique. Next we simply collapse each edge of Γ which carries a group which is not two-ended. The result is a graph of groups structure Γ' for G in which every edge group is two-ended. In particular, it follows that every vertex group of Γ' is finitely generated. Of course, Γ' is no longer bipartite. Further, it is not true that any nontrivial almost invariant subset of G over a two-ended subgroup is enclosed by a vertex of Γ'. However it follows from our construction of Γ' from Γ that if G possesses a nontrivial almost invariant subset over a two-ended subgroup H, then H has a subgroup of finite index which is conjugate into some vertex group of Γ'. The known JSJ-decompositions along two-ended subgroups can all be refined to such a decomposition, but in [**36**] there are some assumptions on unfoldedness which may somewhat restrict the choice of splittings used to refine Γ.

However these decompositions are not canonical. We call any such decomposition of G along two-ended subgroups a non-canonical JSJ-decomposition of G.

Sela in [**49**] initiated the study of uniqueness of such decompositions up to some moves which he called sliding, conjugation and modifying the boundary homomorphism by a conjugation. In [**16**], Forester gave a complete description of the uniqueness properties of these decompositions. In [**17**], Forester considered two moves on graphs of groups or equivalently G–trees called a collapse move and an expansion move. The first move involves selecting an edge s in the graph which is not a loop and such that the induced map from the edge group $G(s)$ to the initial vertex group $G(v)$ is an isomorphism. One then collapses s down to v. An expansion move is the reverse of a collapse move. A move which factors as a composition of expansions and collapses is called an elementary deformation. He showed that two cocompact G–trees are related by an elementary deformation if and only if they have the same elliptic subgroups. More recently, Guirardel [**22**] gave a simplified proof of Forester's result.

Consider one of the non-canonical JSJ-decompositions derived as above from our canonical decomposition Γ. By construction none of the edge splittings is crossed strongly by any nontrivial almost invariant set over a two-ended subgroup. Thus they are elliptic with respect to any splitting of G over a two-ended group. Next consider any two G–trees T_1 and T_2 corresponding to two such decompositions. The Fuchsian vertex groups are the same in both and are thus elliptic with respect to both the trees. The other vertex groups of T_1 and T_2 do not admit any splittings over two-ended groups relative to the edge groups. This is because in our refinement of Γ, we used a maximal family of compatible splittings of G. It follows that the actions of G on the trees T_1 and T_2 have the same elliptic subgroups. Thus, by Forester's theorem, we have:

Theorem 15.1. — *Let G be a one-ended group and let T_1 and T_2 be two G–trees corresponding to non-canonical JSJ-decompositions of G. Then T_1 and T_2 are related by an elementary deformation.*

Corollary 15.2. — *Let G be a one-ended group and σ a splitting of G over a two-ended subgroup. Let Γ be any non-canonical JSJ-decomposition of G. Then either σ is enclosed by a V_0–vertex of Γ of Fuchsian type, or σ can be obtained by collapses and expansions from the edge splittings of Γ.*

CHAPTER 16

EXTENSIONS

One might wonder how far the techniques above can be used to obtain canonical decompositions enclosing almost invariant sets over other classes of groups. In our arguments, the first step was to show that in any cross-connected component, the crossings are either all weak or all strong. Then we used different kinds of arguments in these two cases. For CCC's in which all crossing is weak, we proved that such a CCC must enclose a splitting, then proved a finiteness result for the CCC and finally we needed an accessibility result. Here are two general results of this type. The first is a reformulation of Theorem B.3.13, and asserts the existence of a splitting very generally.

Theorem 16.1. — *Suppose that $H \subset G$ are finitely generated groups and that G does not contain any nontrivial almost invariant subsets over subgroups of infinite index in H. Let X be a nontrivial almost invariant subset of G over H and suppose that the translates of X do not cross each other strongly. Then G splits over a subgroup commensurable with H.*

Our arguments in chapter 7 of this paper extend to show the following accessibility result.

Theorem 16.2. — *Suppose $\mathcal{K}$ is a class of small groups closed under commensurability. Suppose G is a finitely presented group which does not split over a subgroup of infinite index in an element of $\mathcal{K}$. Let Γ_k be a graph of groups decomposition of G without redundant vertices and with all edge groups in $\mathcal{K}$, and suppose that for each k, Γ_{k+1} is a refinement of Γ_k. Then, the sequence Γ_k stabilises.*

These two results can be used to handle more cases of CCC's in which all crossing is weak. However, to handle CCC's in which all crossing is strong, we used the results and arguments of Bowditch [**8**], and of Dunwoody and Swenson [**15**] which in turn depend on the special structure of VPC groups. Some of these results can be summarised in the following theorems.

Theorem 16.3. — *Let G be a finitely generated group and let X be a nontrivial almost invariant subset of G over a virtually abelian group H of rank $n+1$. Suppose that X is n–canonical with respect to abelian groups, i.e. it does not cross any nontrivial almost invariant subset over a virtually abelian group of rank $\leqslant n$. Then G splits over a virtually abelian group of rank $n+1$. If X does not cross any translate of X strongly, then G splits over a subgroup commensurable with H. Moreover the almost invariant sets associated with the splittings obtained are n–canonical with respect to abelian groups.*

Theorem 16.4. — *Let G be a finitely generated group and let X be a nontrivial almost invariant subset of G over a $VPC(n+1)$ group H. Suppose G does not have any nontrivial almost invariant sets over VPC groups of length $< n$ and that X is n–canonical. Then G splits over a $VPC(n+1)$ group. If X does not cross any translate of X strongly, then G splits over a subgroup commensurable with H. Moreover the almost invariant sets associated with the splittings obtained are n–canonical.*

Example 14.1 suggests that it may not be possible to strengthen Theorem 16.3. It is possible that similar results may be provable for the special class of slender groups considered in [**14**], but these problems are still open. The techniques of [**14**] and [**20**] enclose splittings rather than almost invariant sets. The technique in [**20**] is particularly appealing. It is their construction that suggested to us regular neighbourhoods and Bowditch's use of pretrees provided us with a crucial technique. The crossing hypotheses used in their technique are weaker than ours, provided of course that one starts with splittings. Thus there may be further refinements of the decompositions that we obtained if one combines their techniques with ours. We recall that even in the case of 3–manifolds the canonical decompositions obtained by enclosing splittings only are different from the standard topological JSJ-decompositions (see [**34**]). So, our work seems to suggest that there are several possible generalisations of JSJ-decompositions to groups. Moreover, our theories of regular neighbourhoods and canonical splittings are very general and these may apply to almost invariant sets over groups more general than VPC groups.

APPENDIX A

THE SYMMETRY OF INTERSECTION NUMBERS IN GROUP THEORY

This appendix consists of the complete text of Scott's paper [**42**], and its later correction. We are grateful to the editors of Geometry and Topology for agreeing to this. The numbering of results in this appendix is the same as the numbering in the original paper with the addition of a prefix A.

If one considers two simple closed curves L and S on a closed orientable surface F, one can define their intersection number to be the least number of intersection points obtainable by isotoping L and S transverse to each other. (Note that the count is to be made without any signs attached to the intersection points.) By definition, this number is symmetric, i.e. the roles of L and S are interchangeable. This can be regarded as a definition of the intersection number of the two infinite cyclic subgroups Λ and Σ of the fundamental group of F which are carried by L and S. In this paper, we show that an analogous definition of intersection number of subgroups of a group can be given in much greater generality and proved to be symmetric. We also give an interpretation of these intersection numbers.

In [**36**], Rips and Sela considered a torsion free finitely presented group G and infinite cyclic subgroups Λ and Σ such that G splits over each. (A group G splits over a subgroup C if either G has a HNN decomposition $G = A*_C$, or G has an amalgamated free product structure $G = A *_C B$, where $A \neq C \neq B$.) They effectively considered the intersection number $i(\Lambda, \Sigma)$ of Λ with Σ, and they proved that $i(\Lambda, \Sigma) = 0$ if and only if $i(\Sigma, \Lambda) = 0$. Using this, they proved that G has what they call a JSJ decomposition. If $i(\Lambda, \Sigma)$ was not zero, it follows from their work that G can be expressed as the fundamental group of a graph of groups with some vertex group being a surface group H which contains Λ and Σ. Now it is intuitively clear (and we discuss it further at the end of section A.2 of this paper) that the intersection number of Λ with Σ is the same whether it is measured in G or in H. Also the intersection numbers of Λ and Σ in H are symmetric because of their topological interpretation. So it follows at the end of all their work that the intersection numbers of Λ and Σ in G

are also symmetric. In 1994, Rips asked if there was a simpler proof of this symmetry which does not depend on their proof of the JSJ splitting. The answer is positive, and the ideas needed for the proof are all essentially contained in earlier papers of the author. This paper is a belated response to Rips' question. The main idea is to reduce the natural, but not clearly symmetric, definition of intersection number to counting the intersections of suitably chosen sets. The most general possible algebraic situation in which to define intersection numbers seems to be that of a finitely generated group G and two finitely generated subgroups Λ and Σ, not necessarily cyclic, such that the number of ends of each of the pairs (G, Λ) and (G, Σ) is more than one. Note that any infinite cyclic subgroup Λ of $\pi_1(F)$ satisfies $e(\pi_1(F), \Lambda) = 2$. This is because F is closed and orientable so that the cover of F with fundamental group Λ is an open annulus which has two ends. In order to handle the general situation, we will need the concept of an almost invariant set, which is closely related to the theory of ends. We should note that Kropholler and Roller [**29**] introduced an intersection cohomology class in the special case of $PD(n-1)$–subgroups of PDn–groups. Their ideas are closely related to ours, and we will discuss the connections at the start of section A.3 of this paper. Finally, we should point out that since Rips asked the above question about symmetry of intersection numbers, Dunwoody and Sageev [**14**] have given a proof of the existence of a JSJ decomposition for any finitely presented group which is very much simpler and more elementary than that of Rips and Sela.

The preceding discussion is a little misleading, as the intersection numbers which we define are not determined simply by a choice of subgroups. In fact, we define intersection numbers for almost invariant sets. A special case occurs when one has a group G and subgroups Λ and Σ such that G splits over each, as a splitting of G has a well defined almost invariant set associated. This is discussed in section A.2. Thus we can define the intersection number of two splittings of G. In the case of cyclic subgroups of surface groups corresponding to simple closed curves, these curves determine splittings of the surface group over each cyclic subgroup, and the intersection number we define for these splittings is the same as the topological intersection number of the curves.

In the first section of this paper, we discuss in more detail intersection numbers of closed curves on surfaces. In the second section we introduce the concept of an almost invariant set and prove the symmetry results advertised in the title. In the third section, we discuss the interpretation of intersection numbers when they are defined, and how our ideas are connected with those of Kropholler and Roller.

Acknowledgments: This paper was written while the author was visiting the Mathematical Sciences Research Institute in Berkeley in 1996/7. Research at MSRI is supported in part by NSF grant DMS-9022140. He is also grateful for the partial support provided by NSF grants DMS-9306240 and DMS-9626537.

A.1. The symmetry for surface groups

In this section, we will discuss further the special case of two essential closed curves L and S on a compact surface F. This will serve to motivate the definitions in the following section, and also show that the results of that section do indeed answer the question of Rips. It is not necessary to assume that F is closed or orientable, but we do need to assume that L and S are two-sided on F. As described in the introduction in the case of simple curves, one defines their intersection number to be the least number of intersection points obtainable by homotoping L and S transverse to each other, where the count is to be made without any signs attached to the intersection points. (One should also insist that L and S be in general position, in order to make the count correctly.) Of course, this number is symmetric, i.e. the roles of L and S are interchangeable. We will show in section A.2 that one can define these intersection numbers in an algebraically natural way. There is also an idea of self-intersection number for a curve on a surface and we will discuss a corresponding algebraic idea.

For the next discussion, we will restrict our attention to the case when L and S are simple and introduce the algebraic approach to defining intersection numbers taken by Rips and Sela in [**36**]. Let G denote $\pi_1(F)$. Suppose that L and S cannot be made disjoint and choose a basepoint on $L \cap S$. Suppose that L represents the element λ of G. This element λ cannot be trivial, nor can L be parallel to a boundary component of F, because of our assumption that L and S cannot be made disjoint. Thus L induces a splitting of G over the infinite cyclic subgroup Λ of G which is generated by λ. Let σ denote the element of G represented by S. Define $d(\sigma, \lambda)$ to be the length of σ when written as a word in cyclically reduced form in the splitting of G determined by L. Similarly, define $d(\lambda, \sigma)$ to be the length of λ when written as a word in cyclically reduced form in the splitting of G determined by S. For convenience, suppose also that L and S are separating. Then each of these numbers is equal to the intersection number of L and S described above and therefore $d(\lambda, \sigma) = d(\sigma, \lambda)$. What is interesting is that this symmetry is not obvious from the purely algebraic point of view, but it is obvious topologically because the intersection of two sets is symmetric.

In the above discussion, we restricted attention to simple closed curves on a surface F, because the algebraic analogue is clear. If F is closed, then not only does a simple closed curve on F determine a splitting of $\pi_1(F)$ over the infinite cyclic subgroup carried by the curve, but any splitting of $\pi_1(F)$ over an infinite cyclic subgroup is induced in this way by some simple closed curve on F. Hence the algebraic situation described above exactly corresponds to the topological situation when F is closed.

Now we continue with further discussion of the intersection number of two closed curves L and S which need not be simple. As in [**18**], it will be convenient to assume that L and S are shortest closed geodesics in some Riemannian metric on F so that they automatically intersect minimally. Instead of defining the intersection number of L and S in the "obvious" way, we will interpret our intersection numbers in suitable covers of F, exactly as in [**18**] and [**19**]. Let F_Λ denote the cover of F with fundamental group equal to Λ. Then L lifts to F_Λ and we denote its lift by L again. Let l denote the pre-image of this lift in the universal cover $\widetilde{F}$ of F. The full pre-image of L in $\widetilde{F}$ consists of disjoint lines which we call L–lines, which are all translates of l by the action of G. Similarly, we define F_Σ, the line s and S–lines in $\widetilde{F}$. Now we consider the images of the L–lines in F_Σ. Each L–line has image in F_Σ which is a possibly singular line or circle. Then we define $d(L, S)$ to be the number of images of L–lines in F_Σ which meet S. Similarly, we define $d(S, L)$ to be the number of images of S–lines in F_Λ which meet L. It is shown in [**18**], using the assumption that L and S are shortest closed geodesics, that each L–line in F_Σ crosses S at most once, and similarly for S–lines in F_Λ. It follows that $d(L, S)$ and $d(S, L)$ are each equal to the number of points of $L \cap S$, and so they are equal to each other. (This assumes that L and S are in general position.)

Here is an argument which shows that $d(L, S)$ and $d(S, L)$ are equal without reference to the situation in the surface F. Recall that the L–lines are translates of l by elements of G. Of course, there is not a unique element of G which sends l to a given L–line. In fact, the L–lines are in natural bijective correspondence with the cosets $g\Lambda$ of Λ in G. (Our groups act on the left on covering spaces.) The images of the L–lines in F_Σ are in natural bijective correspondence with the double cosets $\Sigma g\Lambda$, and $d(L, S)$ counts the number of these double cosets such that the line gl crosses s. Similarly, $d(S, L)$ counts the number of the double cosets $\Lambda h\Sigma$ such that the line hs crosses l. Note that it is trivial that gl crosses s if and only if l crosses $g^{-1}s$. Now we use the bijection from G to itself given by sending each element to its inverse. This induces a bijection between the set of all double cosets $\Sigma g\Lambda$ and the set of all double cosets $\Lambda h\Sigma$ by sending $\Sigma g\Lambda$ to $\Lambda g^{-1}\Sigma$. It follows that it also induces a bijection between those double cosets $\Sigma g\Lambda$ such that gl crosses s and those double cosets $\Lambda h\Sigma$ such that hs crosses l, which shows that $d(L, S)$ equals $d(S, L)$ as required.

This argument has more point when one applies it to a more complicated situation than that of curves on surfaces. In [**19**], we considered least area maps of surfaces into a 3-manifold. The intersection number which we used there was defined in essentially the same way but it had no obvious topological interpretation such as the number of double curves of intersection. We proved that our intersection numbers were symmetric by the above double coset argument, in [**19**] just before Theorem 6.3.

A.2. Intersection numbers in general

In order to handle the general case, we will need the idea of an almost invariant set. This idea was introduced by Cohen in [**9**] and was first used in the relative context by Houghton in [**24**]. We will introduce this idea and explain its connection with the foregoing.

Let E and F be sets. We say that E and F are almost equal, and write $E \overset{a}{=} F$, if the symmetric difference $(E - F) \cup (F - E)$ is finite. If E is contained in some set W on which a group G acts on the right, we say that E is almost invariant if $Eg \overset{a}{=} E$, for all g in G. An almost invariant subset E of W will be called non-trivial if it is infinite and has infinite complement. The connection of this idea with the theory of ends of groups is via the Cayley graph Γ of G with respect to some finite generating set of G. (Note that in this paper groups act on the left on covering spaces and, in particular, G acts on its Cayley graph on the left.) Using $\mathbf{Z}_2$ as coefficients, we can identify 0-cochains and 1-cochains on Γ with sets of vertices or edges. A subset E of G represents a set of vertices of Γ which we also denote by E, and it is a beautiful fact, due to Cohen [**9**], that E is an almost invariant subset of G if and only if δE is finite, where δ is the coboundary operator. If H is a subgroup of G, we let $H\backslash G$ denote the set of cosets Hg of H in G, i.e. the quotient of G by the left action of H. Of course, G will no longer act on the left on this quotient, but it will still act on the right. Thus we have the idea of an almost invariant subset of $H\backslash G$.

Now we again consider the situation of simple closed curves L and S on a compact surface F and let $\widetilde{F}$ denote the universal cover of F. Pick a generating set for G which can be represented by a bouquet of circles embedded in F. We will assume that the wedge point of the bouquet does not lie on L or S. The pre-image of this bouquet in $\widetilde{F}$ will be a copy of the Cayley graph Γ of G with respect to the chosen generating set. The pre-image in F_Λ of the bouquet will be a copy of the graph $\Lambda\backslash\Gamma$, the quotient of Γ by the action of Λ on the left. Consider the closed curve L on F_Λ. Let D denote the set of all vertices of $\Lambda\backslash\Gamma$ which lie on one side of L. Then D has finite coboundary, as δD equals exactly the edges of $\Lambda\backslash\Gamma$ which cross L. Hence D is an almost invariant subset of $\Lambda\backslash G$. Let X denote the pre-image of D in Γ, so that X equals the set of vertices of Γ which lie on one side of the line l. There is an algebraic description of X in terms of canonical forms for elements of G as follows. Suppose that L separates F, so that $G = A *_\Lambda B$. Also suppose that L and D are chosen so that all the vertices of Γ labelled with an element of Λ do not lie in X. Pick right transversals T and T' for Λ in A and B respectively, both of which contain the identity e of G. (A right transversal of Λ in A consists of a choice of coset representative for each coset $a\Lambda$.) Each element of G can be expressed uniquely in the form $a_1b_1 \dots a_nb_n\lambda$, where $n \geqslant 1$, λ lies in Λ, each a_i lies in $T - \{e\}$ except that a_1 may be trivial, and each b_i lies in $T' - \{e\}$ except that b_n may be trivial. Then X consists of those elements for which a_1 is non-trivial. If Λ is non-separating in F, there is a similar description for X. See Theorem 1.7 of [**47**] for details. Similarly, we can define a set E in F_Σ and its pre-image Y in $\widetilde{F}$ which equals the set of vertices of Γ which lie on one side of the line s. Now finally the connection between the earlier arguments and almost invariant sets can be given. For we can decide whether the lines l and s cross by considering instead the sets X and Y. The lines l and s together divide G into the four sets

$X \cap Y, X^* \cap Y, X \cap Y^*$ and $X^* \cap Y^*$, where X * denotes $G - X$, and l crosses s if and only if each of these four sets projects to an infinite subset of $\Sigma \backslash G$. Equally, s crosses l if and only if each of these four sets projects to an infinite subset of $\Lambda \backslash G$. As we know that l crosses s if and only if s crosses l, it follows that these conditions are equivalent. We will show that this symmetry holds in a far more general context.

Note that in the preceding example the subset X of G is Λ–invariant under the left action of Λ on G, i.e. $\lambda X = X$, for all λ in Λ.

For the most general version of this symmetry result, we can consider any finitely generated group G. Note that the subgroups of G which we consider need not be finitely generated.

Definition A.2.1. — If G is a finitely generated group and H is a subgroup, then a subset X of G is H–almost invariant if X is invariant under the left action of H, and simultaneously the quotient set $H \backslash X$ is almost invariant under the right action of G. In addition, X is a non-trivial H–almost invariant subset of G if $H \backslash X$ and $H \backslash X^*$ are both infinite.

Note that if X is a non-trivial H–almost invariant subset of G, then $e(G, H)$ is at least 2, as $H \backslash X$ is a non-trivial almost invariant subset of $H \backslash G$.

Definition A.2.2. — Let X be a Λ–almost invariant subset of G and let Y be a Σ–almost invariant subset of G. We will say that X crosses Y if each of the four sets $X \cap Y, X^* \cap Y, X \cap Y^*$ and $X^* \cap Y^*$ projects to an infinite subset of $\Sigma \backslash G$.

Note that it is obvious that if Y is trivial, then X cannot cross Y. Our first and most basic symmetry result is the following. This is essentially proved in Lemma 2.3 of [**41**], but the context there is less general.

Lemma A.2.3. — *If G is a finitely generated group with subgroups Λ and Σ, and X is a non-trivial Λ–almost invariant subset of G and Y is a non-trivial Σ–almost invariant subset of G, then X crosses Y if and only if Y crosses X.*

Remark A.2.4. — If X and Y are both trivial, then neither can cross the other, so the above symmetry result is clear. However, this symmetry result fails if only one of X or Y is trivial. Here is a simple example. Let Λ and Σ denote infinite cyclic groups with generators λ and σ respectively, and let G denote the group $\Lambda \times \Sigma$. We identify G with the set of integer points in the plane. Let $X = (m, n) \in G : n > 0$, and let $Y = (m, n) \in G : m = 0$. Then X is a non-trivial Λ–almost invariant subset of G and Y is a trivial Σ–almost invariant subset of G. One can easily check that Y crosses X, although X cannot cross Y as Y is trivial.

Proof. — Suppose that X does not cross Y. By replacing one or both of X and Y by its complement if needed, we can assume that $X \cap Y$ projects to a finite subset of $\Sigma \backslash G$. The fact that Y is non-trivial implies that $\Sigma \backslash Y$ is an infinite subset of $\Sigma \backslash G$, so there is a point z in $\Sigma\backslash Y$ which is not in the image of $X \cap Y$. Now we need to use some choice of generators for G and consider the corresponding Cayley graph Γ of G. The vertices of Γ are identified with G and the action of G on itself on the left extends to an action on Γ. We consider z and the image of $X \cap Y$ in the quotient graph $\Sigma\backslash\Gamma$. As $X \cap Y$ has finite image, there is a number d such that each point of its image can be joined to z by a path of length at most d. As the projection of Γ to $\Sigma\backslash\Gamma$ is a covering map, it follows that each point of $X \cap Y$ can be joined to some point lying above z by a path of length at most d. As any point above z lies in X^*, it follows that each point of $X \cap Y$ can be joined to some point of X^* by a path of length at most d. Hence each point of $X \cap Y$ lies at most distance d from δX. Thus the image of $X \cap Y$ in $\Lambda\backslash\Gamma$ lies within the d–neighbourhood of the compact set $\delta(\Lambda\backslash X)$, and so must itself be finite. It follows that Y does not cross X, which completes the proof of the symmetry result. □

At the start of this section, we explained how to connect the topological intersection of simple closed curves on a surface with crossing of sets. One can construct many other interesting examples in much the same way.

Example A.2.5. — As before, let F denote a closed surface with fundamental group G, and let $\widetilde{F}$ denote the universal cover of F. Pick a generating set of G which can be represented by a bouquet of circles embedded in F, so that $\widetilde{F}$ contains a copy of the Cayley graph Γ of G with respect to the chosen generators. Let F_1 denote a cover of F which is homeomorphic to a four punctured torus and let Λ denote its fundamental group. For example, if F is the closed orientable surface of genus three, we can consider a compact subsurface F' of F which is homeomorphic to a torus with four open discs removed, and take the cover F_1 of F such that $\pi_1(F_1) = \pi_1(F')$. For notational convenience, we identify F_1 with $S^1 \times S^1$ with the four points $(1,1), (1,i), (1,-1)$ and $(1,-i)$ removed. Now we choose 1-dimensional submanifolds of F_1 each consisting of two circles and each separating F_1 into two pieces. Let L denote $S^1 \times \{e^{\pi i/4}, e^{5\pi i/4}\}$ and let S denote $S^1 \times \{e^{3\pi i/4}, e^{7\pi i/4}\}$. As before, we let D denote all the vertices of the graph $\Lambda\backslash\Gamma$ in F_1 which lie on one side of L, and let E denote all the vertices of the graph $\Lambda\backslash\Gamma$ in F_1 which lie on one side of S. Let X and Y denote the pre-images of D and E in G. Now D is an almost invariant subset of $\Lambda\backslash G$, as δD equals exactly the edges of $\Lambda\backslash\Gamma$ which cross L, and E is almost invariant for similar reasons. Hence X and Y are each Λ–almost invariant subsets of G. Clearly X and Y cross. An important feature of this example is that although X and Y cross, the boundaries L and S of the corresponding surfaces in F_1 are disjoint. This is quite different from the example with which we introduced almost invariant sets, but this is a much more typical situation.

Definition A.2.6. — Let Λ and Σ be subgroups of a finitely generated group G. Let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$ and let X and Y denote the pre-images in G of D and E respectively. We define $i(D, E)$ to equal the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y.

For this definition to be interesting, we need to show that $i(D, E)$ is finite, which is not obvious from the definition in this general situation. In fact, it may well be false if one does not assume that the groups Λ and Σ are finitely generated, although we have no examples[(1)]. From now on, we will assume that Λ and Σ are finitely generated.

Lemma A.2.7. — *Let Λ and Σ be finitely generated subgroups of a finitely generated group G. Let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. Then $i(D, E)$ is finite.*

Proof. — This is again proved by using the Cayley graph, so it appears to depend on the fact that G is finitely generated. However, we have no examples where $i(D, E)$ is not finite when G is not finitely generated. The proof we give is essentially contained in that of Lemmas 4.3 and 4.4 of [**40**]. Start by considering the finite graph δD in $\Lambda\backslash\Gamma$. As Λ is finitely generated, we can add edges and vertices to δD to obtain a finite connected subgraph $\delta_1 D$ of $\Lambda\backslash\Gamma$ which contains δD and has the property that its inclusion in $\Lambda\backslash\Gamma$ induces a surjection of its fundamental group to Λ. Thus the pre-image of $\delta_1 D$ in Γ is a connected graph which we denote by $\delta_1 X$. Similarly, we obtain a finite connected graph $\delta_1 E$ of $\Sigma\backslash\Gamma$ which contains δE and has connected pre-image $\delta_1 Y$ in Γ. As usual, we will denote the pre-images of D and E in G by X and Y respectively.

Next we claim that if gX crosses Y then $g(\delta_1 X)$ intersects $\delta_1 Y$. (The converse need not be true.) Suppose that $g(\delta_1 X)$ and $\delta_1 Y$ are disjoint. Then $g(\delta_1 X)$ cannot meet δY. As $g(\delta_1 X)$ is connected, it must lie in Y or Y^*. It follows that $g(\delta X)$ lies in Y or Y^*, so that one of the four sets $X \cap Y, X^* \cap Y, X \cap Y^*$ and $X^* \cap Y^*$ must be empty, which implies that gX does not cross Y.

Now we can show that $i(D, E)$ must be finite. Recall that $i(D, E)$ is defined to be the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y. The preceding paragraph implies that $i(D, E)$ is bounded above by the number of double cosets $\Sigma g\Lambda$ such that $g(\delta_1 X)$ meets $\delta_1 Y$. Let P and Q be finite subgraphs of $\delta_1 X$ and $\delta_1 Y$ which project onto $\delta_1 D$ and $\delta_1 E$ respectively. If $g(\delta_1 X)$ meets $\delta_1 Y$, then there exist elements λ of Λ and σ of Σ such that $g(\lambda P)$ meets σQ. Thus $\sigma^{-1} g\lambda P$ meets Q. Now there are only finitely many elements of G which can translate P to meet Q, and it follows that $i(D, E)$ is bounded above by this number. □

We have just shown that, as in the preceding section, the intersection numbers we have defined are symmetric, but we will need a little more information.

(1) Since this was written, Guirardel has given such examples. See Example 2.34.

Lemma A.2.8. — *Let G be a finitely generated group with subgroups Λ and Σ, let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. Then the following statements hold:*

(1) $i(D, E) = i(E, D)$,

(2) $i(D, E) = i(D^*, E) = i(D, E^*) = i(D^*, E^*)$,

(3) *if D' is almost equal to D and E' is almost equal to E, and X, X' and Y, Y' denote their pre-images in G, then X crosses Y if and only if X' crosses Y', so that $i(D, E) = i(D', E')$.*

Proof. — The first part is proved by using the bijection from G to itself given by sending each element to its inverse. This induces a bijection between all double cosets $\Sigma g\Lambda$ and $\Lambda h\Sigma$ by sending $\Sigma g\Lambda$ to $\Lambda g^{-1}\Sigma$, and it further induces a bijection between those double cosets $\Sigma g\Lambda$ such that gX crosses Y and those double cosets $\Lambda h\Sigma$ such that hY crosses X.

The second part is clear from the definitions.

For the third part, we note that, as E and E' are almost equal, so are their complements in $\Sigma\backslash G$, and it follows that X crosses Y if and only if it crosses Y'. Hence the symmetry proved in Lemma A.2.3, shows that Y crosses X if and only Y' crosses X. Now the same argument reversing the roles of D and E yields the required result. □

At this point, we have defined in a natural way a number which can reasonably be called the intersection number of D and E, but have not yet defined an intersection number for subgroups of G. First note that if $e(G, \Lambda)$ is equal to 2, then all choices of non-trivial almost invariant sets in $\Lambda\backslash G$ are almost equal or almost complementary. Let D denote some choice here. Suppose that $e(G, \Sigma)$ is also equal to 2, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. The third part of the preceding lemma implies that $i(D, E)$ is independent of the choices of D and E and so depends only on the subgroups Λ and Σ. This is then the definition of the intersection number $i(\Lambda, \Sigma)$. In the special case when G is the fundamental group of a closed orientable surface and Λ and Σ are cyclic subgroups of G, it is automatic that $e(G, \Lambda)$ and $e(G, \Sigma)$ are each equal to 2. The discussion of the previous section clearly shows that this definition coincides with the topological definition of intersection number of loops representing generators of these subgroups, whether or not those loops are simple. Note that one can also define the self-intersection number of an almost invariant subset D of $\Lambda\backslash G$ to be $i(D, D)$, and hence can define the self-intersection number of a subgroup Λ of G such that $e(G, \Lambda) = 2$. Again this idea generalises the topological idea of self-intersection number of a loop on a surface.

If one considers subgroups Λ and Σ such that $e(G, \Lambda)$ or $e(G, \Sigma)$ is greater than 2, there are possibly different ideas for their intersection number depending on which almost invariant sets we pick. (It is tempting to simply define $i(\Lambda, \Sigma)$ to be the minimum possible value for $i(D, E)$, where D is a non-trivial Λ–almost invariant subset of G and E is a non-trivial Σ–almost invariant subset of G. But this does not seem to be the "right" definition.) However, there is a natural way to choose these almost invariant sets if we are given splittings of G over Λ and Σ. As discussed in the previous section in the case of surface groups, the standard way to do this when $G = A *_\Lambda B$ is in terms of canonical forms for elements of G as follows. Pick right transversals T and T' for Λ in A and B respectively, both of which contain the identity e of G. Then each element can be expressed uniquely in the form $a_1 b_1 \ldots a_n b_n \lambda$, where $n \geqslant 1$, λ lies in Λ, each a_i lies in $T - \{e\}$ except that a_1 may be trivial, and each b_i lies in $T' - \{e\}$ except that b_n may be trivial. Let X denote the subset of G consisting of elements for which a_1 is non-trivial, and let D denote $\Lambda \backslash X$. It is easy to check directly that X is Λ–almost invariant. One must check that $\lambda X = X$, for all λ in Λ and that $Dg \overset{a}{=} D$, for all g in G. The first equation is trivial, and the second is easily checked when g lies in A or B, which implies that it holds for all g in G. Note also that the definition of X is independent of the choices of transversals of Λ in A and B. Then D is the almost invariant set determined by the given splitting of G. This definition seems asymmetric, but if instead we consider the Λ–almost invariant subset of G consisting of elements whose canonical form begins with a non-trivial element of B, we will obtain an almost invariant subset of $\Lambda \backslash G$ which is almost equal to the complement of D. There is a similar description of D when $G = A *_\Lambda$. For details see Theorem 1.7 of [**47**]. The connection between D and the given splitting of G can be seen in several ways. From the topologists' point of view, one sees this as described earlier for surface groups. From the point of view of groups acting on trees, there is also a very natural description. One identifies a splitting of G with an action of G on a tree T without inversions, such that the quotient $G \backslash T$ has a single edge. Let e denote the edge of T with stabiliser Λ, let v denote the vertex of e with stabiliser A, and let E denote the component of $T - e$ which contains v. Then we can define $X = g \in G : ge \subset E$. It is easy to check directly that this set is the same as the set X defined above using canonical forms.

In the preceding paragraph, we showed how to obtain a well defined intersection number of given splittings over Λ and Σ. An important point to notice is that this intersection number is not determined by the subgroups Λ and Σ of G only. It depends on the given splittings. In the case when G is a surface group, this is irrelevant as there can be at most one splitting of a surface group over a given infinite cyclic subgroup. But in general, a group G with subgroup Λ can have many different splittings over Λ.

Example A.2.9. — Here is a simple example to show that intersection numbers depend on splittings, not just on subgroups. First we note that the self-intersection number of any splitting is zero. Now construct a group G by amalgamating four groups G_1, G_2, G_3 and G_4 along a common subgroup Λ. Thus G can be expressed as $G_{12} *_\Lambda G_{34}$, where G_{ij} is the subgroup of G generated by G_i and G_j, but it can also be expressed as $G_{13} *_\Lambda G_{24}$ or $G_{14} *_\Lambda G_{23}$. The intersection number of any distinct pair of these splittings of G is non-zero, but all the splittings being considered are splittings over the same group Λ.

A question which arose in our introduction in connection with the work of Rips and Sela was how the intersection number of two subgroups of a group G alters if one replaces G by a subgroup. In general, nothing can be said, but in interesting cases one can understand the answer to this question. The particular case considered by Rips and Sela was of a finitely presented group G which is expressed as the fundamental group of a graph of groups with some vertex group being a group H which contains infinite cyclic subgroups Λ and Σ. Further H is the fundamental group of a surface F and Λ and Σ are carried by simple closed curves L and S on F. A point deliberately left unclear in our earlier discussion of their work was that F is not a closed surface. It is a compact surface with non-empty boundary. The curves L and S are not homotopic to boundary components and so define splittings of H. The edges in the graph of groups which are attached to H all carry some subgroup of the fundamental group of a boundary component of F. This implies that L and S also define splittings of G. It is clear from this picture that the intersection number of Λ and Σ should be the same whether measured in G or in H, as it should equal the intersection number of the curves L and S, but this needs a little more thought to make precise. As usual, the first point to make is that we are really talking about the intersection numbers of the splittings defined by L and S, rather than intersection numbers of Λ and Σ. For the number of ends $e(H, \Lambda)$ and $e(H, \Sigma)$ are infinite when F is a surface with boundary. As G is finitely presented, we can attach cells to the boundary of F to construct a finite complex K with fundamental group G. Now the identification of the intersection number of the given splittings of G with the intersection number of L and S proceeds exactly as at the start of this section, where we showed how to identify the intersection number of the given splittings of H with the intersection number of L and S.

A.3. Interpreting intersection numbers

It is natural to ask what is the meaning of the intersection numbers defined in the previous section. The answer is already clear in the case of a surface group with cyclic subgroups. In this section[(2)], we will give an interpretation of the intersection number of two splittings of a finitely generated group G over finitely generated subgroups. We start by discussing the connection with the work of Kropholler and Roller.

In [**29**], Kropholler and Roller introduced an intersection cohomology class for $PD(n-1)$–subgroups of a PDn–group. The pairs involved always have two ends,

[(2)] This section contains a serious error which was corrected in a later paper. This correction has been added to this appendix as the next section.

so the work of the previous section defines an intersection number in this situation. The connection between our intersection number and their intersection cohomology class is the following. Recall that if one has subgroups Λ and Σ of a finitely generated group G, such that $e(G,\Lambda)$ and $e(G,\Sigma)$ are each equal to 2, then one chooses a non-trivial Λ–almost invariant subset X of G and a non-trivial Σ–almost invariant subset Y of G and defines our intersection number $i(\Lambda,\Sigma)$ to equal the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y. Their cohomology class encodes the information about which double cosets have this crossing property. Thus their invariant is much finer than the intersection number and it is trivial to deduce the intersection number from their cohomology class.

To interpret the intersection number of two splittings of a group G, we need to discuss the Subgroup Theorem for amalgamated free products. Let G be a finitely generated group, which splits over finitely generated subgroups Λ and Σ. We will write $G = A_1 *_\Lambda (B_1)$ to denote that either G has the HNN structure $A_1 *_\Lambda$ or G has the structure $A_1 *_\Lambda B_1$. Similarly, we will write $G = A_2 *_\Sigma (B_2)$. The Subgroup Theorem, see [**47**] and [**50**] (or [**51**]) for discussions from the topological and algebraic points of view, yields a graph of groups structure $\Phi_1(\Sigma)$ for Σ, with vertex groups lying in conjugates of A_1 or B_1 and edge groups lying in conjugates of Λ. Typically this graph will not be finite or even locally finite. However, as Σ is finitely generated, there is a finite subgraph Ψ_1 which still carries Σ. If we reverse the roles of Λ and Σ, we will obtain a graph of groups structure $\Phi_2(\Lambda)$ for Λ, with vertex groups lying in conjugates of A_2 or B_2 and edge groups lying in conjugates of Σ, and there is a finite subgraph Ψ_2 which still carries Λ. We show below that, in most cases, the intersection number of Λ and Σ measures the minimal possible number of edges of these finite subgraphs. Notice that if we consider the special case when G is the fundamental group of a closed surface and Λ and Σ are infinite cyclic subgroups, this statement is clear. Now the symmetry of intersection numbers implies the surprising fact that the minimal number of edges for Ψ_1 and Ψ_2 are the same.

There is an alternative point of view which we will use for our proof. The splitting $A_2 *_\Sigma (B_2)$ of G corresponds to an action of G on a tree T such that the quotient $G\backslash T$ has one edge. The edge stabilisers in this action on T are all conjugate to Σ and the vertex stabilisers are conjugate to A_2 or B_2 as appropriate. If one has a subgroup Λ of G, the quotient $\Lambda\backslash T$ will be the graph underlying $\Phi_2(\Lambda)$. There is a Λ–invariant subtree T' of T, such that the graph $\Lambda\backslash T'$ is the graph underlying Ψ_2. Whichever point of view you take, it is necessary to connect it with the ideas about almost invariant sets which we have already discussed. Here is our interpretation of intersection numbers.

Theorem A.3.1. — *Let G be a finitely generated group, which splits over finitely generated subgroups Λ and Σ, such that if U and V are any conjugates of Λ and Σ respectively, then $U \cap V$ has infinite index in both U and V. Then the intersection* *number of the two splittings equals the minimal number of edges in each of the graphs* Ψ_1 *and* Ψ_2.

Remark A.3.2. — This result is clearly false if the condition on conjugates is omitted. For example, if $\Lambda = \Sigma$, then $\Psi_1(\Sigma)$ and $\Psi_2(\Lambda)$ will each consist of a single vertex, but the intersection number of the two splittings need not be zero.

The proof will use the following sequence of lemmas.

We start with a general result about minimal G-invariant subtrees of a tree T on which a group G acts. If every element of G fixes each point of a non-trivial subtree T' of T, then any vertex of T' is a minimal G-invariant subtree of T. Otherwise, there is a unique minimal G-invariant subtree of T. An orientation of an edge e of T consists of a choice of one vertex as the initial vertex $i(e)$ of e and the other as the terminal vertex $t(e)$. An oriented path in T consists of a finite sequence of oriented edges $e_1, e_2, \dots, e_k$ of T, such that $t(e_j) = i(e_{j+1})$, for $1 \leqslant j \leqslant k-1$. If we consider two oriented edges e and e' of T we say that they are coherently oriented if there is an oriented path which begins with one and ends with the other. Finally, given an edge e of T and an element g of G, we will say that e and ge are coherently oriented if for some (and hence either) orientation on e and the induced orientation on ge, the edges e and ge are coherently oriented.

Lemma A.3.3. — *Suppose that a group G acts on a tree T without inversions and without fixing a point. Let T' denote the minimal G-invariant subtree. Then an edge e of T lies in T' if and only if there exists an element g of G such that e and ge are distinct and coherently oriented.*

Proof. — First consider an edge e not lying in T'. Orient e so that it is the first edge of an oriented path λ in T which starts with e, has no edge in T', and ends at a vertex of T'. Thus ge, with the induced orientation, is the first edge of an oriented path $g\lambda$ in T which starts with ge, has no edge in T', and ends at a vertex of T'. Now the unique path in T which joins e and ge must consist either of λ and $g\lambda$ together with a path in T' or of an initial segment of e together with an initial segment of ge. In either case, it follows that e and ge are not coherently oriented.

Now we consider an edge e of T' and its image $\overline{e}$ in $G\backslash T'$.

If $\overline{e}$ is non-separating in $G\backslash T'$, let μ denote an oriented path in $G\backslash T'$ which joins the ends of $\overline{e}$ and meets $\overline{e}$ only in its endpoints. Then the loop formed by $\mu \cup \overline{e}$ lifts to an oriented path in T', which shows that there is g in G such that e and ge are distinct and coherently oriented.

If $\overline{e}$ separates $G\backslash T'$, we can write the graph $G\backslash T'$ as $\Gamma_1 \cup \overline{e} \cup \Gamma_2$, where each Γ_i is connected and meets $\overline{e}$ in one endpoint only. Now consider the graph of groups structure given by $G\backslash T'$. By contracting each Γ_i to a point, we obtain an amalgamated free product structure of G as $G_1 *_C G_2$, where $C = \mathit{stab}(e)$ and each G_i is the fundamental group of the graph of groups Γ_i. Let T_i denote the tree on which G_i acts with quotient Γ_i. Then the complement in T' of the edge e and its translates consists of disjoint copies of T_1 and T_2. We identify T_i with the copy of T_i which meets e. Note that T_1 and T_2 are disjoint. Now it is clear that $G_1 \neq C \neq G_2$. For if $G_1 = C$, then $G = G_2$, which implies that T_2 is a G-invariant subtree of T', contradicting the minimality of T'. As $G_1 \neq C$, there is an element g_1 of G_1 such that $g_1 e \neq e$, and similarly there is an element g_2 of G_2 such that $g_2 e \neq e$. For each i, there is a path λ_i in T_i which begins at e and ends at $g_i e$. As T_1 and T_2 are disjoint, so are λ_1 and λ_2. It follows that of the three edges $e, g_1 e, g_2 e$, at least one pair is coherently oriented, which completes the proof of the lemma. □

The following result is clear.

Lemma A.3.4. — *Suppose that a group G acts on a tree T without inversions and without fixing a point. Let e denote an edge of T, let E denote a component of $T - \{e\}$ and let g denote an element of G. Then e and ge are distinct and coherently oriented if and only if either $gE \subsetneqq E$ or $gE^* \subsetneqq E^*$.*

Next we need to connect this with almost invariant sets, although the following result does not use the almost invariance property.

Lemma A.3.5. — *Suppose that a group G acts on a tree T without inversions and without fixing a point and suppose that the quotient graph $G\backslash T$ has only one edge. Let e denote an edge of T, let E denote a component of $T - \{e\}$ and let $Y = k \in G : ke \subset E$. Then the following statements hold for all elements g of G:*

(1) *$gY \subset Y$ if and only if $gE \subset E$, and $gY^* \subset Y^*$ if and only if $gE^* \subset E^*$.*
(2) *$gY = Y$ if and only if $gE = E$, and $gY^* = Y^*$ if and only if $gE^* = E^*$.*
(3) *$gY \subsetneqq Y$ if and only if $gE \subsetneqq E$, and $gY^* \subsetneqq Y^*$ if and only if $gE^* \subsetneqq E^*$.*

Proof. — Suppose that $gE \subset E$. If k lies in Y, then $ke \subset E$, so that $gke \subset gE \subset E$. Thus gk also lies in Y. It follows that $gY \subset Y$.

Conversely, suppose that $gY \subset Y$ and consider an edge f of E. As $G\backslash T$ has only one edge, $f = ke$ for some k in G. As f lies in E, k lies in Y, and hence gk also lies in Y by our assumption that $gY \subset Y$. Thus $gke \subset E$, so that $gf \subset E$. Thus implies that $gE \subset E$ as required.

The proof for the second equivalence in part 1 is essentially the same.

The equivalences in part 2 follow by applying part 1 for g and g^{-1}. Now the equivalences in part 3 are clear. □

Next we connect the above inclusions with crossing of sets.

Lemma A.3.6. — *Suppose that a finitely generated group G splits over a finitely generated subgroup Λ with corresponding Λ–almost invariant set X and also splits over a finitely generated subgroup Σ with corresponding Σ–almost invariant set Y. Suppose further that if U and V are any conjugates of Λ and Σ respectively, then $U \cap V$ has* *infinite index in U. Then X crosses Y if and only if there is an element λ in Λ such that either $\lambda Y \subsetneqq Y$ or $\lambda Y^* \subsetneqq Y^*$.*

Proof. — We claim that there exists $\lambda_1 \in \Lambda$ such that either $\lambda_1 Y \subsetneqq Y$ or $\lambda_1 Y^* \subsetneqq Y$, and there exists $\lambda_2 \in \Lambda$ such that either $\lambda_2 Y \subsetneqq Y^*$ or $\lambda_2 Y^* \subsetneqq Y^*$. Assuming this, either $\lambda_1 Y \subsetneqq Y$ or $\lambda_2 Y^* \subsetneqq Y^*$, and our proof is complete, or we have $\lambda_1 Y^* \subsetneqq Y$ and $\lambda_2 Y \subsetneqq Y^*$. The last possibility implies that $\lambda_2\lambda_1 Y^* \subsetneqq \lambda_2 Y \subsetneqq Y^*$, again completing the proof.

To prove our claim, we pick a finite generating set for G, and consider the Cayley graph Γ of G with respect to this generating set. As Y is a Σ–almost invariant set associated to a splitting $A_2 *_\Sigma (B_2)$ of G over Σ, we can choose Γ and Y so that, for every g in G, $g\delta Y$ is disjoint from or coincides with δY. A simple way to arrange this is to take as generators of G the union of a set of generators of Σ and of A_2 and B_2, so that $\Gamma(G)$ contains a copy of the Cayley graph $\Gamma(\Sigma)$ of Σ and $\Sigma\backslash\Gamma$ contains $\Sigma\backslash\Gamma(\Sigma)$ which is a wedge of circles. (Note that this uses the hypothesis that Σ is finitely generated.) Let v denote the wedge point, and let E denote the collection of vertices of $\Sigma\backslash\Gamma$ which can be joined to v by a path whose interior is disjoint from v such that the last edge is labelled by an element of A. Then clearly δE consists of exactly those edges of $\Sigma\backslash\Gamma$ which have one end at v and are labelled by an element of A. Further, if we let Y denote the pre-image of E in G, then, for every g in G, $g\delta Y$ is disjoint from or coincides with δY.

In order to prove that λ_1 exists, we argue as follows. As $\Lambda \cap \Sigma$ has infinite index in Λ, and as δX is Λ-invariant, it follows that δX must contain points which are arbitrarily far from δY on each side of δY. Recall that $\Lambda\backslash X$ is an almost invariant subset of $\Lambda\backslash G$, so that it has finite coboundary which equals $\Lambda\backslash\delta X$. Hence there is a number d such that any point of $\Lambda\backslash\delta X$ can be joined to the image of δY in $\Lambda\backslash\Gamma$ by a path of length at most d. It follows that any point of δX can be joined to $\lambda\delta Y$, for some λ in Λ, by a path in Γ of length at most d. Hence there is a translate of δY which contains points on one side of δY and another translate which contains points on the other side of δY. Hence there are elements λ_1 and λ_2 of Λ such that $\lambda_1\delta Y$ lies on one side of δY and $\lambda_2\delta Y$ lies on the other. Without loss of generality, we can suppose that $\lambda_1\delta Y$ lies on the side containing Y so that either $\lambda_1 Y \subsetneqq Y$ or $\lambda_1 Y^* \subsetneqq Y$. As $\lambda_2\delta Y$ lies on the side of δY containing Y^*, either $\lambda_2 Y \subsetneqq Y^*$ or $\lambda_2 Y^* \subsetneqq Y^*$. This completes the proof of the claim made at the start of the proof. □

Now we can give the proof of Theorem A.3.1.

Proof. — Recall that G splits over finitely generated subgroups Λ and Σ such that if U and V are any conjugates of Λ and Σ, then $U \cap V$ has infinite index in both U and V. Also G acts on a tree T so as to induce the given splitting over Σ. Let e denote an edge of T with stabiliser Σ and consider the action of Λ on T. Our hypothesis on conjugates of Λ and Σ implies, in particular, that Λ is not contained in any conjugate of Σ so that Λ cannot fix an edge of T. Thus there is a unique minimal Λ-invariant subtree T' of T. Lemma A.3.3 shows that an edge he of T lies in T' if and only if there is λ in Λ such that he and λhe are distinct and coherently oriented. Lemma A.3.4 shows that this occurs if and only if either $\lambda hE \subsetneqq hE$ or $\lambda hE^* \subsetneqq hE^*$, and Lemma A.3.5 shows that this occurs if and only if $\lambda hY \subsetneqq hY$ or $\lambda hY^* \subsetneqq hY^*$. Finally Lemma A.3.6 shows that this occurs if and only if X crosses hY. We conclude that an edge he of T lies in T' if and only if X crosses hY. Thus the edges of T which lie in the minimal Λ-invariant subtree T' naturally correspond to the cosets $h\Sigma$ such that X crosses hY. Hence the number of edges in $\Psi_2(\Lambda)$ equals the number of double cosets $\Lambda h\Sigma$ such that X crosses hY, which was defined to be the intersection number of the given splittings. Similarly, one can show that the intersection number of the given splittings equals the minimal possible number of edges in the graph $\Psi_1(\Sigma)$. This completes the proof of Theorem A.3.1. □

A.4. Correction

This section consists of the complete text of the correction to Scott's paper [**42**].

Theorem A.3.1 is false as stated. The error in the argument occurs in the proof of Lemma A.3.6. See below for a counterexample.

Lemma A.3.6 asserts that, under suitable hypotheses, X crosses Y if and only if there is an element λ in Λ such that either $\lambda Y \subsetneqq Y$ or $\lambda Y^* \subsetneqq Y^*$. One of these implications is correct. If such a λ exists, then it is true that X must cross Y. However, I failed to give any argument for this, and I provide one below. The other implication is false. The mistake is contained in the second sentence of the last paragraph of the proof of Lemma A.3.6. A simple fix is to amend the statements of Theorem A.3.1 and Lemma A.3.6 to take this into account. Thus we need the additional hypothesis for Lemma A.3.6 that if X crosses Y, then δX must contain points which are arbitrarily far from δY on each side of δY. We also need the additional hypothesis for Theorem A.3.1 that if X crosses gY then δX must contain points which are arbitrarily far from δgY on each side of δgY. This technical assumption is often but not always satisfied.

Here is the half of the proof of Lemma A.3.6 which was omitted. This asserts that if there is an element λ in Λ such that either $\lambda Y \subsetneqq Y$ or $\lambda Y^* \subsetneqq Y^*$, then X must cross Y. We will assume that $\lambda Y \subsetneqq Y$, as the argument in the other case is essentially identical. As Y is associated to a splitting of G, it is easy to see that the distance of $\lambda^n \delta Y$ from δY must tend to infinity as $n \to \infty$. (For example, if $G = A *_C B$, and Y is the set of words in G which begin in $A - C$, then λ must begin in $A - C$ and end in $B - C$.) Now consider an element $g \in G$, and let d denote the distance of g from δY. Then d is also the distance of $\lambda^n g$ from $\lambda^n \delta Y$. Hence, for any element g of G, all translates $\lambda^n g$ lie in Y, for suitably large n. If we apply these statements to an edge of δX, and recall that δX is preserved by λ, we see that δX must contain points which are arbitrarily far from δY and lie in Y. By applying the same discussion to λ^{-1}, we see that δX must also contain points which are arbitrarily far from δY and lie in Y^*. Hence δX must contain points which are arbitrarily far from δY on each side of δY as required.

Now we come to the promised counterexample. Let G denote the fundamental group of the closed orientable surface M of genus two. Let D denote a simple closed curve on M which separates M into two once-punctured tori S and T and let D' denote a non-separating simple closed curve in the interior of S. Let W denote the surface obtained from S by removing a regular neighbourhood of D'. Let C denote a non-separating simple closed curve on M whose intersection number with D is two, and which is disjoint from D'. We will describe two splittings of G. The first will be the HNN splitting over an infinite cyclic group determined by C. The second will be the amalgamated free product splitting of G over $\pi_1(W)$ with vertex groups $\pi_1(S)$ and $\pi_1(W \cup T)$. These two splittings satisfy the hypotheses of Theorem A.3.1. If one considers $\pi_1(C)$ as a subgroup of the splitting over $\pi_1(W)$, the minimal graph of groups obtained has no edges, because $\pi_1(C)$ is contained in $\pi_1(W \cup T)$ which is a vertex group. If one considers $\pi_1(W)$ as a subgroup of the HNN splitting determined by C, the minimal graph of groups obtained has at least one edge because $\pi_1(W)$ does not lie in a conjugate of any vertex group. (The graph in question has exactly one edge, but this fact is not needed here.) This shows that Theorem A.3.1 must fail for this example, because the numbers of edges in these two graphs are not equal. It is also true that Lemma A.3.6 fails for this example. Let X and Y be the usual subsets of G associated to the two splittings. I claim that X crosses Y but δX does not contain points which are arbitrarily far from δY on each side of δY. To see this, consider the picture in the cover M_C of M whose fundamental group equals $\pi_1(C)$. This cover is an open annulus which contains a lift of C which we will continue to denote by C. As in section A.2, we pick a generating set for G which can be represented by a bouquet of circles embedded in M, so that the pre-image in the universal cover $\widetilde{M}$ of M of this bouquet is a copy of the Cayley graph Γ of G, and we identify the vertices of this graph with G. Now let E denote the set of all vertices of $\pi_1(C)\backslash\Gamma$ in M_C which lie on one side of C. Then E represents an almost invariant subset of $\pi_1(C)\backslash G$ and the pre-image of E in Γ can be taken to be X. Now consider the picture in the cover M_W of M whose fundamental group equals $\pi_1(W)$. This cover consists of a lift of W, which we will continue to denote by W and open collars attached to the boundary components of W. Let F denote the set of all vertices of $\pi_1(W)\backslash\Gamma$ which lie in the union of W together with the collar attached to the component D of ∂W. Then F represents an almost invariant subset of $\pi_1(W)\backslash G$ and the pre-image of F in G can be taken to be Y. The pre-image in $\widetilde{M}$ of C is a line whose image in M_W is a properly embedded line meeting W in a compact arc which projects homeomorphically to $C \cap W$. Now inspection shows that each of the four sets $X^{(*)} \cap Y^{(*)}$ has infinite image in M_W so that X crosses Y but δX does not contain points which are arbitrarily far from δY on each side of δY.

The new version of Theorem A.3.1 described here is, of course, rather unsatisfactory as the extra hypothesis is technical and it is not clear when it holds. However, there is a little more which can be said without any extra work. For it follows from the preceding discussion that the number of edges in each of the minimal graphs of groups described above is always less than or equal to the intersection number of the two splittings being considered.

APPENDIX B

SPLITTINGS OF GROUPS AND INTERSECTION NUMBERS

This appendix consists of the complete text of our paper [**44**]. We are grateful to the editors of Geometry and Topology for agreeing to this. The numbering of results in this appendix is the same as the numbering in the original paper with the addition of a prefix B.

In this paper, we will discuss an algebraic version of intersection numbers which was introduced by Scott in [**42**]. First we need to discuss intersection numbers in the topological setting. Let F denote a surface and let L and S each be a properly immersed two-sided circle or compact arc in F. Here 'properly' means that the boundary of the 1–manifold lies in the boundary of F. One can define the intersection number of L and S to be the least number of intersection points obtainable by homotoping L and S transverse to each other. (The count is to be made without any signs attached to the intersection points.) It is obvious that this number is symmetric in the sense that it is independent of the order of L and S. It is also obvious that L and S have intersection number zero if and only if they can be properly homotoped to be disjoint. It seems natural to define the self-intersection number of an immersed two-sided circle or arc L in F to be the least number of transverse intersection points obtainable by homotoping L into general position. With this definition, L has self-intersection number zero if and only if it is homotopic to an embedding. However, in light of later generalisations, it turns out that this definition should be modified a little in order to ensure that the self-intersection number of any cover of a simple closed curve is also zero. No modification is needed unless L is a circle which can be homotoped to cover another immersion with degree greater than 1. In this case, suppose that the maximal degree of covering which can occur is k and that L covers L' with degree k. Then we define the self-intersection number of L to be k^2 times the self-intersection number of L'. With this modified definition, L has self-intersection number zero if and only if it can be homotoped to cover an embedding.

In [**19**], Freedman, Hass and Scott introduced a notion of intersection number and self-intersection number for two-sided π_1–injective immersions of compact surfaces into 3–manifolds which generalises the preceding ideas. Their intersection number cannot be described as simply as for curves on a surface, but it does share some important properties. In particular, it is a non-negative integer and it is symmetric, although this symmetry is not obvious from the definition. Further, two surfaces have intersection number zero if and only if they can be homotoped to be disjoint, and a single surface has self-intersection number zero if and only if it can be homotoped to cover an embedding. These two facts are no longer obvious consequences of the definition, but are non-trivial applications of the theory of least area surfaces.

In [**42**], Scott extended the ideas of [**19**] to define intersection numbers in a purely group theoretic setting. The details will be discussed in the first section of this paper, but we give an introduction to the ideas here. It seems clear that everything discussed in the preceding two paragraphs should have a purely algebraic interpretation in terms of fundamental groups of surfaces and 3–manifolds, and the aim is to find an interpretation which makes sense for any group. It seems natural to attempt to define the intersection number of two subgroups $H$and K of a given group G. This is exactly what the topological intersection number of simple closed curves on a surface does when G is the fundamental group of a closed orientable surface and we restrict attention to infinite cyclic subgroups H and K. However, if one considers two simple arcs on a surface F with boundary, they each carry the trivial subgroup of $G = \pi_1(F)$, whereas we know that some arcs have intersection number zero and others do not. Thus intersection numbers are not determined simply by the groups involved. We need to look a little deeper in order to formulate the algebraic analogue. First we need to think a bit more about curves on surfaces. Let L be a simple arc or closed curve on an orientable surface F, let G denote $\pi_1(F)$ and let H denote the image of $\pi_1(L)$ in G. If L separates F then, in most cases, it gives G the structure of an amalgamated free product $A *_H B$, and if L is non-separating, it gives G the structure of a HNN extension $A*_H$. In order to avoid discussing which of these two structures G has, it is convenient to say that a group G *splits over a subgroup* H if G is isomorphic to $A*_H$ or to $A *_H B$, with $A \neq H \neq B$. (Note that the condition that $A \neq H \neq B$ is needed as otherwise any group G would split over any subgroup H. For one can always write $G = G *_H H$.) Thus, in most cases, L determines a splitting of $G = \pi_1(F)$. Usually one ignores base points, so that the splitting of G is only determined up to conjugacy. In [**42**], Scott defined the intersection number of two splittings of any group G over any subgroups H and K. In the special case when G is the fundamental group of a compact surface F and these splittings arise from embedded arcs or circles on F, the algebraic intersection number of the splittings equals the topological intersection number of the corresponding 1–manifolds. The analogous statement holds when G is the fundamental group of a compact 3-manifold and these splittings arise from π_1–injective embedded surfaces. In general, the algebraic intersection number shares some properties of the topological intersection number. Algebraic intersection numbers are symmetric, and if G, H and K are finitely generated, the intersection number of splittings of G over H and over K is a non-negative integer.

The first main result of this paper is a generalisation to the algebraic setting of the fact that two simple arcs or closed curves on a surface have intersection number zero if and only if they can be isotoped apart. Of course, the idea of isotopy makes no sense in the algebraic setting, so we need some algebraic language to describe multiple disjoint curves on a surface. Let $L_1, \ldots, L_n$ be disjoint simple arcs or closed curves on a compact orientable surface F with fundamental group G, such that each L_i determines a splitting of G. Together they determine a graph of groups structure on G with n edges. We say that a collection of n splittings of a group G is *compatible* if G can be expressed as the fundamental group of a graph of groups with n edges, such that, for each i, collapsing all edges but the i–th yields the i–th splitting of G. We will say that the splittings are *compatible up to conjugacy* if collapsing all edges but the i–th yields a splitting of G which is conjugate to the i–th given splitting. Clearly disjoint essential simple arcs or closed curves on F define splittings of G which are compatible up to conjugacy. The precise statement we obtain is the following.

Theorem B.2.5. — *Let G be a finitely generated group with n splittings over finitely generated subgroups. This collection of splittings is compatible up to conjugacy if and only if each pair of splittings has intersection number zero. Further, in this situation, the graph of groups structure on G obtained from these splittings has a unique underlying graph, and the edge and vertex groups are unique up to conjugacy.*

So far, we have not discussed any algebraic analogue of non-embedded arcs or circles on surfaces. There is such an analogue which is the idea of an almost invariant subset of the quotient $H\backslash G$, where H is a subgroup of G. This generalises the idea of an immersed curve in a surface or of an immersed π_1–injective surface in a 3–manifold which carries the subgroup H of G. We give the definitions in section B.1. There is also an idea of intersection number of such things, which we give in Definition B.1.3. This too was introduced by Scott in [**42**]. Our second main result, Theorem B.2.8, is an algebraic analogue of the fact that a singular curve on a surface or a singular surface in a 3–manifold which has self-intersection number zero can be homotoped to cover an embedding. It asserts that if $H\backslash G$ has an almost invariant subset with self-intersection number zero, then G has a splitting over a subgroup H' commensurable with H. We leave the precise statement until section B.2.

In a separate paper [**45**], we use the ideas about intersection numbers of splittings developed in [**42**] and in this paper to study JSJ decompositions of Haken 3–manifolds. The problem there is to recognize which splittings of the fundamental group of such a manifold arise from the JSJ decomposition (see [**32**] and [**34**]). It turns out that a class of splittings which we call canonical can be defined using intersection numbers and we use this to show that the JSJ decomposition for Haken 3–manifolds depends only on the fundamental group. This leads to an algebraic proof of Johannson' Deformation Theorem. It seems very likely that similar ideas apply to Sela's JSJ decompositions [**49**] of hyperbolic groups and thus provide a common thread to the two types of JSJ decomposition. Thus, the use of intersection numbers seems to provide a tool in the study of diverse topics in group theory and this paper together with [**42**] provides some of the foundational material.

This paper is organised as follows. In section B.1, we recall from [**42**] the basic definitions of intersection numbers in the algebraic context. We also prove a technical result which was essentially proved by Scott [**41**] in 1980. However, Scott's results were all formulated in the context of surfaces in 3–manifolds, so we give a complete proof of the generalisation to the purely group theoretic context. Section B.2 is devoted to the proofs of our two main results discussed above.

There is a second natural idea of intersection number, which we discuss in section B.3. We call it the strong intersection number. It is not symmetric in general, but this is not a problem when one is considering self-intersection numbers. We also discuss when the two kinds of intersection number are equal, which then forces the strong intersection number to be symmetric. We use these ideas to give a new approach to a result of Kropholler and Roller [**29**] on splittings of Poincaré duality groups. We also discuss applications of our ideas to prove a special case of a conjecture of Kropholler and Roller [**30**] on splittings of groups in general. We point out that these ideas lead to an alternative approach to the algebraic Torus Theorem [**15**]. We end the section with a brief discussion of an error in [**42**]. In section 3 of that paper, Scott gave an incorrect interpretation of the intersection number of two splittings. His error was caused by confusing the ideas of strong and ordinary intersection. However, the arguments in [**42**] work to give a nice interpretation of the intersection number in the case when it is equal to the strong intersection number. Without this condition, finding nice interpretations of the two intersection numbers is an open problem.

B.1. Preliminaries and statements of main results

We will start by recalling from [**42**] how to define intersection numbers in the algebraic setting. We will connect this with the natural topological idea of intersection number already discussed in the introduction. Consider two simple closed curves L and S on a closed orientable surface F. As in [**18**], it will be convenient to assume that L and S are shortest geodesics in some Riemannian metric on F so that they automatically intersect minimally. We will interpret the intersection number of L and S in suitable covers of F, exactly as in [**18**] and [**19**]. Let G denote $\pi_1(F)$, let H denote the infinite cyclic subgroup of G carried by L, and let F_H denote the cover of F with fundamental group equal to H. Then L lifts to F_H and we denote its lift by L again. Let l denote the pre-image of this lift in the universal cover $\widetilde{F}$ of F. The full pre-image of L in $\widetilde{F}$ consists of disjoint lines which we call L–lines, which are all translates of l by the action of G. (Note that in this paper groups act on the left on covering spaces.) Similarly, we define K, F_K, the line s and S–lines in $\widetilde{F}$. Now we consider the images of the L–lines in F_K. Each L–line has image in F_K which is a line or circle. Then we define $d(L,S)$ to be the number of images of L–lines in F_K which meet S. Similarly, we define $d(S,L)$ to be the number of images of S–lines in F_H which meet L. It is shown in [**18**], using the assumption that L and S are shortest closed geodesics, that each L–line in F_K crosses S at most once, and similarly for S–lines in F_H. It follows that $d(L,S)$ and $d(S,L)$ are each equal to the number of points of $L \cap S$, and so they are equal to each other.

We need to take one further step in abstracting the idea of intersection number. As the stabiliser of l is H, the L–lines naturally correspond to the cosets gH of H in G. Hence the images of the L–lines in F_K naturally correspond to the double cosets KgH. Thus we can think of $d(L,S)$ as the number of double cosets KgH such that gl crosses s. This is the idea which we generalise to define intersection numbers in a purely algebraic setting.

First we need some terminology.

Two sets P and Q are *almost equal* if their symmetric difference $P - Q \cup Q - P$ is finite. We write $P \overset{a}{=} Q$.

If a group G acts on the right on a set Z, a subset P of Z is *almost invariant* if $Pg \overset{a}{=} P$ for all g in G. An almost invariant subset P of Z is *non-trivial* if P and its complement $Z - P$ are both infinite. The complement $Z - P$ will be denoted simply by P^*, when Z is clear from the context

For finitely generated groups, these ideas are closely connected with the theory of ends of groups via the Cayley graph Γ of G with respect to some finite generating set of G. (Note that G acts on its Cayley graph on the left.) Using $\mathbb{Z}_2$ as coefficients, we can identify 0–cochains and 1–cochains on Γ with sets of vertices or edges. A subset P of G represents a set of vertices of Γ which we also denote by P, and it is a beautiful fact, due to Cohen [**9**], that P is an almost invariant subset of G if and only if δP is finite, where δ is the coboundary operator. Now Γ has more than one end if and only if there is an infinite subset P of G such that δP is finite and P^* is also infinite. Thus Γ has more than one end if and only if G contains a non-trivial almost invariant subset. If H is a subgroup of G, we let $H \backslash G$ denote the set of cosets Hg of H in G, i.e. the quotient of G by the left action of H. Of course, G will no longer act on the left on this quotient, but it will still act on the right. Thus we also have the idea of an almost invariant subset of $H \backslash G$, and the graph $H \backslash \Gamma$ has more than one end if and only if $H \backslash G$ contains a non-trivial almost invariant subset. Now the number of ends $e(G)$ of G is equal to the number of ends of Γ, so it follows that $e(G) > 1$ if and only if G contains a non-trivial almost invariant subset. Similarly, the number of ends

$e(G, H)$ of the pair (G, H) equals the number of ends of $H\backslash\Gamma$, so that $e(G, H) > 1$ if and only if $H\backslash G$ contains a non-trivial almost invariant subset.

Now we return to the simple closed curves L and S on the surface F. Pick a generating set for G which can be represented by a bouquet of circles embedded in F. We will assume that the wedge point of the bouquet does not lie on L or S. The pre-image of this bouquet in $\widetilde{F}$ will be a copy of the Cayley graph Γ of G with respect to the chosen generating set. The pre-image in F_H of the bouquet will be a copy of the graph $H\backslash\Gamma$, the quotient of Γ by the action of H on the left. Consider the closed curve L on F_H. Let P denote the set of all vertices of $H\backslash\Gamma$ which lie on one side of L. Then P has finite coboundary, as δP equals exactly the edges of $H\backslash\Gamma$ which cross L. Hence P is an almost invariant subset of $H\backslash G$. Let X denote the pre-image of P in Γ, so that X equals the set of vertices of Γ which lie on one side of the line l. Now finally the connection between the earlier arguments and almost invariant sets can be given. For we can decide whether the lines l and s cross by considering instead the sets X and Y. The lines l and s together divide G into the four sets $X \cap Y$, $X^* \cap Y$, $X \cap Y^*$ and $X^* \cap Y^*$, where X * denotes $G - X$, and l crosses s if and only if each of these four sets projects to an infinite subset of $K\backslash G$.

Now let G be a group with subgroups H and K, let P be a non-trivial almost invariant subset of $H\backslash G$ and let Q be a non-trivial almost invariant subset of $K\backslash G$. We will define the intersection number $i(P, Q)$ of P and Q. First we need to consider the analogues of the sets X and Y in the preceding paragraph, and to say what it means for them to cross.

Definition B.1.1. — If G is a group and H is a subgroup, then a subset X of G is *H–almost invariant* if X is invariant under the left action of H, and simultaneously $H\backslash X$ is an almost invariant subset of $H\backslash G$. In addition, X is a *non-trivial* H–almost invariant subset of G, if the quotient sets $H\backslash X$ and $H\backslash X^*$ are both infinite.

Note that if H is trivial, then a H–almost invariant subset of G is the same as an almost invariant subset of G.

Definition B.1.2. — Let X be a H–almost invariant subset of G and let Y be a K–almost invariant subset of G. We will say that X *crosses* Y if each of the four sets $X \cap Y$, $X^* \cap Y$, $X \cap Y^*$ and $X^* \cap Y^*$ projects to an infinite subset of $K\backslash G$.

We will often write $X^{(*)} \cap Y^{(*)}$ instead of listing the four sets $X \cap Y$, $X^* \cap Y$, $X \cap Y^*$ and $X^* \cap Y^*$.

If G is a group and H is a subgroup, then we will say that a subset W of G is *H–finite* if it is contained in the union of finitely many left cosets Hg of H in G, and we will say that two subsets V and W of G are *H–almost equal* if their symmetric difference is H-finite.

In this language, X crosses Y if each of the four sets $X^{(*)} \cap Y^{(*)}$ is not K–finite.

This definition of crossing is not symmetric, but it is shown in [**42**] that if G is a finitely generated group with subgroups H and K, and X is a non-trivial H–almost invariant subset of G and Y is a non-trivial K–almost invariant subset of G, then X crosses Y if and only if Y crosses X. If X and Y are both trivial, then neither can cross the other, so the above symmetry result is clear. However, this symmetry result fails if only one of X or Y is trivial. This lack of symmetry will not concern us as we will only be interested in non-trivial almost invariant sets.

Now we come to the definition of the intersection number of two almost invariant sets.

Definition B.1.3. — Let H and K be subgroups of a finitely generated group G. Let P denote a non-trivial almost invariant subset of $H\backslash G$, let Q denote a non-trivial almost invariant subset of $K\backslash G$ and let X and Y denote the pre-images of P and Q respectively in G. Then the intersection number $i(P,Q)$ of P and Q equals the number of double cosets KgH such that gX crosses Y.

Remark B.1.4. — The following facts about the intersection number are proved in Lemmas A.2.7 and A.2.8 of Appendix A.

(1) Intersection numbers are symmetric, i.e. $i(P,Q) = i(Q,P)$.

(2) $i(P,Q)$ is finite when G, H, and K are all finitely generated.

(3) If P' is an almost invariant subset of $H\backslash G$ which is almost equal to P or to P^* and if Q' is an almost invariant subset of $K\backslash G$ which is almost equal to Q or to Q^*, then $i(P',Q') = i(P,Q)$.

We will often be interested in situations where X and Y do not cross each other and neither do many of their translates. This means that one of the four sets $X^{(*)} \cap Y^{(*)}$ is K-finite, and similar statements hold for many translates of $X$and Y. If $U = uX$ and $V = vY$ do not cross, then one of the four sets $U^{(*)} \cap V^{(*)}$ is K^v-finite, but probably not K-finite. Thus one needs to keep track of which translates of X and Y are being considered in order to have the correct conjugate of K, when formulating the condition that U and V do not cross. The following definition will be extremely convenient because it avoids this problem, thus greatly simplifying the discussion at certain points.

Definition B.1.5. — Let U be a H–almost invariant subset of G and let V be a K–almost invariant subset of G. We will say that $U \cap V$ is *small* if it is H–finite.

Remark B.1.6. — As the terminology is not symmetric in U and V and makes no reference to H or K, some justification is required. If U is also H'–almost invariant for a subgroup H' of G, then H' must be commensurable with H. Thus $U \cap V$ is H-finite if and only if it is H'–finite. In addition, the fact that crossing is symmetric tells us that $U \cap V$ is H-finite if and only if it is K–finite. This provides the needed justification of our terminology.

Finally, the reader should be warned that this use of the word small has nothing to do with the term small group which means a group with no subgroups which are free of rank 2.

At this point we have the machinery needed to define the intersection number of two splittings. This definition depends on the fact, which we recall from Appendix A, that if a group G has a splitting over a subgroup H, there is a H–almost invariant subset X of G associated to the splitting in a natural way. This is entirely clear from the topological point of view as follows. If $G = A *_H B$, let N denote a space with fundamental group G constructed in the usual way as the union of N_A, N_B and $N_H \times I$. If $G = A*_H$, then N is constructed from N_A and $N_H \times I$ only. Now let M denote the based cover of N with fundamental group H, and denote the based lift of $N_H \times I$ into M by $N_H \times I$. Then X corresponds to choosing one side of $N_H \times I$ in M. We now give a purely algebraic description of this choice of X (see [**47**] for example). If $G = A *_H B$, choose right transversals T_A, T_B of H in A, B, both of which contain the identity element. (A right transversal for a subgroup H of a group G consists of one representative element for each right coset gH of H in G.) Each element of G can be expressed uniquely in the form $a_1 b_1 a_2 \cdots a_n b_n h$ with $h \in H$, $a_i \in T_A$, $b_i \in T_B$, where only h, a_1 and b_n are allowed to be trivial. Then X consists of elements for which a_1 is non-trivial. In the case of a HNN-extension $A*_H$, let α_i, $i = 1, 2$, denote the two inclusions of H in A so that $t^{-1}\alpha_1(h)t = \alpha_2(h)$, and choose right transversals T_i of $\alpha_i(H)$ in A, both of which contain the identity element. Each element of G can be expressed uniquely in the form $a_1 t^{\epsilon_1} a_2 t^{\epsilon_2} \cdots a_n t^{\epsilon_n} a_{n+1}$ where a_{n+1} lies in A and, for $1 \leqslant i \leqslant n$, $\epsilon_i = 1$ or -1, $a_i \in T_1$ if $\epsilon_i = 1$, $a_i \in T_2$ if $\epsilon_i = -1$ and moreover $a_i \neq 1$ if $\epsilon_{i-1} \neq \epsilon_i$. In this case, X consists of elements for which a_1 is trivial and $\epsilon_1 = 1$. In both cases, the stabiliser of X under the left action of G is exactly H and, for every $g \in G$, at least one of the four sets $X^{(*)} \cap gX^{(*)}$ is empty. Note that this is equivalent to asserting that one of the four inclusions $X \subset gX$, $X \subset gX^*$, $X^* \subset gX$, $X^* \subset gX^*$ holds.

The following terminology will be useful.

Definition B.1.7. — A collection E of subsets of G which are closed under complementation is called *nested* if for any pair U and V of sets in the collection, one of the four sets $U^{(*)} \cap V^{(*)}$ is empty. If each element U of E is a H_U–almost invariant subset of G for some subgroup H_U of G, we will say that E is *almost nested* if for any pair U and V of sets in the collection, one of the four sets $U^{(*)} \cap V^{(*)}$ is small.

The above discussion shows that the translates of X and X^* under the left action of G are nested.

Note that X is not uniquely determined by the splitting. In both cases, we made choices of transversals, but it is easy to see that X is independent of the choice of transversal. However, in the case when $G = A *_H B$, we chose X to consist of elements for which a_1 is non-trivial whereas we could equally well have reversed the roles of A and B. This would simply replace X by $X^* - H$. Also either of these sets could be replaced by its complement. We will use the term *standard almost invariant set* for the images in $H\backslash G$ of any one of X, $X \cup H$, X^*, $X^* - H$. In the case when $G = A*_H$, reversing the roles of the two inclusion maps of H into A also replaces X by $X^* - H$. Again we have four standard almost invariant sets which are the images in $H\backslash G$ of any one of X, $X \cup H$, X^*, $X^* - H$. There is a subtle point here. In the amalgamated free product case, we use the obvious isomorphism between $A *_H B$ and $B *_H A$. In the HNN case, let us write $A*_{H,i,j}$ to denote the group $< A, t : t^{-1}i(h)t = j(h) >$. Then the correct isomorphism to use between $A*_{H,i,j}$ and $A*_{H,j,i}$ is not the identity on A. Instead it sends t to t^{-1} and A to $t^{-1}At$. In all cases, we have four standard almost invariant subsets of $H\backslash G$.

Definition B.1.8. — If a group G has splittings over subgroups H and K, and if P and Q are standard almost invariant subsets of $H\backslash G$ and $K\backslash G$ respectively associated to these splittings, then the *intersection number* of this pair of splittings of G is the intersection number of P and Q.

Remark B.1.9. — As any two of the four standard almost invariant subsets of $H\backslash G$ associated to a splitting of G over H are almost equal or almost complementary, Remark B.1.4 tells us that this definition does not depend on the choice of standard almost invariant subsets P and Q.

If X and Y denote the pre-images in G of $P$and Q respectively, and if we conjugate the first splitting by a and the second by b, then X is replaced by aXa^{-1} and $Y$is replaced by bYb^{-1}. Now Xg is H–almost equal to X and Yg is K–almost equal to Y, because of the general fact that for any subset W of G and any element g of G, the set Wg lies in a l-neighbourhood of W, where l equals the length of g. This follows from the equations $d(wg, w) = d(g, e) = l$. It follows that the intersection number of a pair of splittings is unchanged if we replace them by conjugate splittings.

Now we can state two easy results about the case of zero intersection number. Recall that if X is one of the standard H–almost invariant subsets of G determined by a splitting of G over H, then the set of translates of X and X^* is nested. It follows at once that the self-intersection number of $H\backslash X$ is zero. Also if two splittings of G over subgroups H and K are compatible, and if X and Y denote corresponding standard H–almost and K–almost invariant subsets of G, then the set of all translates of X, X^*, Y, Y^* is also nested, so that the intersection number of the two splittings is zero. The next section is devoted to proving converses to each of these statements.

Before going further, we need to say a little more about splittings. Recall from the introduction that a group G is said to split over a subgroup H if G is isomorphic to $A*_H$ or to $A*_H B$, with $A \neq H \neq B$. We will need a precise definition of a splitting. We will say that a *splitting* of G consists either of proper subgroups A and B of G and a subgroup H of $A \cap B$ such that the natural map $A *_H B \to G$ is an isomorphism, or it consists of a subgroup A of G and subgroups H_0 and H_1 of A such that there is an element t of G which conjugates H_0 to H_1 and the natural map $A*_H \to G$ is an isomorphism.

Recall also that a collection of n splittings of a group G is *compatible* if G can be expressed as the fundamental group of a graph of groups with n edges, such that, for each i, collapsing all edges but the i–th yields the i–th splitting of G. We note that if a splitting of a group G over a subgroup H is compatible with a conjugate of itself by some element g of G, then g must lie in H. This follows from a simple analysis of the possibilities. For example, if the splitting $G = A *_H B$ is compatible with its conjugate by some $g \in G$, then G is the fundamental group of a graph of groups with two edges, which must be a tree, such that collapsing one edge yields the first splitting and collapsing the other yields its conjugate by g. This means that each of the two extreme vertex groups of the tree must be one of A, A^g, B or B^g, and the same holds for the subgroup of G generated by the two vertex groups of an edge. Now it is easy to see that $A \subset A^g$ and $B^g \subset B$, or the same inclusions hold with the roles of A and B reversed. In either case it follows that g lies in H as claimed. The case when $G = A*_H$ is slightly different, but the conclusion is the same. This leads us to the following idea of equivalence of two splittings. We will say that two amalgamated free product splittings of G are equivalent, if they are obtained from the same choice of subgroups A, B and H of G. This means that the splittings $A *_H B$ and $B *_H A$ of G are equivalent. Similarly, a splitting $A*_H$ of G is equivalent to the splitting obtained by interchanging the two subgroups H_0 and H_1 of A. Also we will say that any splitting of a group G over a subgroup H is equivalent to any conjugate by some element of H. Then the equivalence relation on all splittings of G which this generates is the idea of equivalence which we will need. Stated in this language, we see that if two splittings are compatible and conjugate, then they must be equivalent.

Note that two splittings of a group G are equivalent if and only if they are over the same subgroup H, and they have exactly the same four standard almost invariant sets.

Next we need to recall the connection between splittings of groups and actions on trees. Bass-Serre theory, [**50**] or [**51**], tells us that if a group G splits over a subgroup H, then G acts without inversions on a tree T, so that the quotient is a graph with a single edge and the vertex stabilisers are conjugate to A or B and the edge stabilisers are conjugate to H. In his important paper [**11**], Dunwoody gave a method for constructing such a G-tree starting from the subset X of G defined above. The crucial property of X which is needed for the construction is the nestedness of the set of translates of X under the left action of G. We recall Dunwoody's result:

Theorem B.1.10. — *Let E be a partially ordered set equipped with an involution $e \to \overline{e}$, where $e \neq \overline{e}$, such that the following conditions hold:*

(1) *If* $e,\ f \in E$ *and* $e \leqslant f$, *then* $\overline{f} \leqslant \overline{e}$,
(2) *If* $e,\ f \in E$, *there are only finitely many* $g \in E$ *such that* $e \leqslant g \leqslant f$,
(3) *If* $e,\ f \in E$, *at least one of the four relations* $e \leqslant f$, $e \leqslant \overline{f}$, $\overline{e} \leqslant f$, $\overline{e} \leqslant \overline{f}$ *holds,*
(4) *If* $e,\ f \in E$, *one cannot have* $e \leqslant f$ *and* $e \leqslant \overline{f}$.

Then there is an abstract tree T *with edge set equal to* E *such that the order relation which* E *induces on the edge set of* T *is equal to the order relation in which* $e \leqslant f$ *if and only if there is an oriented path in* T *which begins with* e *and ends with* f.

One applies this result to the set $E = gX, gX^* : g \in G$ with the partial order given by inclusion and the involution by complementation. There is a natural action of G on E and hence on the tree T. In most cases, G acts on T without inversions and we can recover the original decomposition from this action as follows. Let e denote the edge of T determined by X. Then X can be described as the set $\{g : g \in G,\ ge < e \text{ or } g\overline{e} < e\}$. If the action of G on T has inversions, then the original splitting must have been an amalgamated free product decomposition $G = A *_H B$, with H of index 2 in A. In this case, subdividing the edges of T yields a tree T_1 on which G acts without inversions. If e_1 denotes the edge of T_1 contained in e and containing the terminal vertex of e, then X can be described as the set $\{g : g \in G, ge_1 < e_1 \text{ or } g\overline{e_1} < e_1\}$.

Now we will prove the following result. This implies part 2) of Remark B.1.4. We give the proof here because the proof in Lemma A.2.7 is not complete, and we will need to apply the methods of proof later in this paper.

Lemma B.1.11. — *Let* G *be a finitely generated group with finitely generated subgroups* H *and* K, *a non-trivial* H*–almost invariant subset* X *and a non-trivial* K*–almost invariant subset* Y. *Then* $\{g \in G : gX \text{ and } Y \text{ are not nested}\}$ *consists of a finite number of double cosets* KgH.

Proof. — Let Γ denote the Cayley graph of G with respect to some finite generating set for G. Let P denote the almost invariant subset $H\backslash X$ of $H\backslash G$ and let Q denote the almost invariant subset $K\backslash Y$ of $K\backslash G$. Recall from the start of this section, that if we identify P with the 0-cochain on $H\backslash\Gamma$ whose support is P, then P is an almost invariant subset of $H\backslash G$ if and only if δP is finite. Thus δP is a finite collection of edges in $H\backslash\Gamma$ and similarly δQ is a finite collection of edges in $K\backslash\Gamma$. Now let C denote a finite connected subgraph of $H\backslash\Gamma$ such that C contains δP and the natural map $\pi_1(C) \to H$ is onto, and let E denote a finite connected subgraph of $K\backslash\Gamma$ such that E contains δQ and the natural map $\pi_1(E) \to K$ is onto. Thus the pre-image D of C in Γ is connected and contains δX, and the pre-image F of E in Γ is connected and contains δY. Let Δ denote a finite subgraph of D which projects onto C, and let Φ denote a finite subgraph of F which projects onto E. If gD meets F, there must be elements h and k in H and K such that $gh\Delta$ meets $k\Phi$. Now $\{\gamma \in G : \gamma\Delta \text{ meets } \Phi\}$ is finite, as G acts freely on Γ. It follows that $\{g \in G : gD \text{ meets } F\}$ consists of a finite number of double cosets KgH.

The result would now be trivial if X and Y were each the vertex set of a connected subgraph of Γ. As this need not be the case, we need to make a careful argument as in the proof of Lemma 5.10 of [**47**]. Consider g in G such that gD and F are disjoint. We will show that gX and Y are nested. As D is connected, the vertex set of gD must lie entirely in Y or entirely in Y^*. Suppose that the vertex set of gD lies in Y. For a set S of vertices of Γ, let $\overline{S}$ denote the maximal subgraph of Γ with vertex set equal to S. Each component W of $\overline{X}$ and $\overline{X^*}$ contains a vertex of D. Hence gW contains a vertex of gD and so must meet Y. If gW also meets Y^*, then it must meet F. But as F is connected and disjoint from gD, it lies in a single component gW. It follows that there is exactly one component gW of $\overline{gX}$ and $\overline{gX^*}$ which meets Y^*, so that we must have $gX \subset Y$ or $gX^* \subset Y$. Similarly, if gD lies in Y^*, we will find that $gX \subset Y^*$ or $gX^* \subset Y^*$. It follows that in either case gX and Y are nested as required. □

In Theorem 2.2 of [**41**], Scott used Dunwoody's theorem to prove a general splitting result in the context of surfaces in 3–manifolds. We will use the ideas in his proof a great deal. The following theorem is the natural generalisation of his result to our more general context and will be needed in the proofs of Theorems B.2.5 and B.2.8. The first part of the theorem directly corresponds to the result proved in [**41**], and the second part is a simple generalisation which will be needed later.

Theorem B.1.12

(1) *Let H be a finitely generated subgroup of a finitely generated group G. Let X be a non-trivial H–almost invariant set in G such that $E = gX, gX^* : g \in G$ is almost nested and if two of the four sets $X^{(*)} \cap gX^{(*)}$ are small, then at least one of them is empty. Then G splits over the stabilizer H' of X and H' contains H as a subgroup of finite index. Further, one of the H'–almost invariant sets Y determined by the splitting is H–almost equal to X.*

(2) *Let $H_1, \dots, H_k$ be finitely generated subgroups of a finitely generated group G. Let X_i, $1 \leqslant i \leqslant k$, be a non-trivial H_i–almost invariant set in G such that $E = gX_i, gX_i^* : 1 \leqslant i \leqslant k, g \in G$ is almost nested. Suppose further that, for any pair of elements U and V of E, if two of the four sets $U^{(*)} \cap V^{(*)}$ are small, then at least one of them is empty. Then G can be expressed as the fundamental group of a graph of groups whose i–th edge corresponds to a conjugate of a splitting of G over the stabilizer H_i' of X_i, and H_i' contains H_i as a subgroup of finite index. Further, for each i, one of the H_i'–almost invariant sets determined by the i–th splitting is H_i–almost equal to X_i.*

Most of the arguments needed to prove this theorem are contained in the proof of Theorem 2.2 of [**41**], but in the context of 3–manifolds. We will present the proof of the first part of this theorem, and then briefly discuss the proof of the second part. The idea in the first part is to define a partial order on $E = gX, gX^* : g \in G$, which coincides with inclusion whenever possible. Let U and V denote elements of E. If $U \cap V^*$ is small, we want to define $U \leqslant V$. There is a difficulty, which is what to do if U and V are distinct but $U \cap V^*$ and $V \cap U^*$ are both small. However, the assumption in the statement of Theorem B.1.12 is that if two of the four sets $U^{(*)} \cap V^{(*)}$ are small, then one of them is empty. Thus, as in [**41**], we define $U \leqslant V$ if and only if $U \cap V^*$ is empty or the only small set of the four. Note that if $U \subset V$ then $U \leqslant V$. We will show that this definition yields a partial order on E.

As usual, we let Γ denote the Cayley graph of G with respect to some finite generating set. The distance between two points of G is the usual one of minimal edge path length. Our first step is the analogue of Lemma 2.3 of [**41**].

Lemma B.1.13. — *$U \cap V^*$ is small if and only if it lies in a bounded neighbourhood of each of U, U^*, V, V^*.*

Proof. — As U and V are translates of X or X^*, it suffices to prove that $gX \cap X^*$ is small if and only if it lies in a bounded neighbourhood of each of X, X^*, gX, gX^*. If $gX \cap X^*$ is small, it projects to a finite subset of $H\backslash G$ which therefore lies within a bounded neighbourhood of the image of δX. By lifting paths, we see that each point of $gX \cap X^*$ lies in a bounded neighbourhood of δX, and hence lies in a bounded neighbourhood of X and X^*. By reversing the roles of gX and X^*, we also see that $gX \cap X^*$ lies in a bounded neighbourhood of each of gX and gX^*.

For the converse, suppose that $gX \cap X^*$ lies in a bounded neighbourhood of each of X and X^*. Then it must lie in a bounded neighbourhood of δX, so that its image in $H\backslash G$ must lie in a bounded neighbourhood of the image of δX. As this image is finite, it follows that $gX \cap X^*$ must be small, as required. □

Now we can prove that our definition of $\leqslant$ yields a partial order on E. Our proof is essentially the same as in Lemma 2.4 of [**41**].

Lemma B.1.14. — *If a relation $\leqslant$ is defined on E by the condition that $U \leqslant V$ if and only if $U \cap V^*$ is empty or the only small set of the four sets $U^{(*)} \cap V^{(*)}$, then $\leqslant$ is a partial order.*

Proof. — We need to show that $\leqslant$ is transitive and that if $U \leqslant V$ and $V \leqslant U$ then $U = V$.

Suppose first that $U \leqslant V$ and $V \leqslant U$. The first inequality implies that $U \cap V^*$ is small and the second implies that $V \cap U^*$ is small, so that two of the four sets $U^{(*)} \cap V^{(*)}$ are small. The assumption of Theorem B.1.12 implies that one of these two sets must be empty. As $U \leqslant V$, our definition of $\leqslant$ implies that $U \cap V^*$ is empty. Similarly, the fact that $V \leqslant U$ tells us that $V \cap U^*$ is empty. This implies that $U = V$ as required.

To prove transitivity, let U, V and W be elements of E such that $U \leqslant V \leqslant W$. We must show that $U \leqslant W$.

Our first step is to show that $U \cap W^*$ is small. As $U \cap V^*$ and $V \cap W^*$ are small, we let d_1 be an upper bound for the distance of points of $U \cap V^*$ from V and let d_2 be an upper bound for the distance of points of $V \cap W^*$ from W. Let x be a point of $U \cap W^*$. If x lies in V, then it lies in $V \cap W^*$ and so has distance at most d_2 from W. Otherwise, it must lie in $U \cap V^*$ and so have distance at most d_1 from some point x' of V. If x' lies in W, then x has distance at most d_1 from W. Otherwise, x' lies in $V \cap W^*$ and so has distance at most d_2 from W. In this case, x has distance at most $d_1 + d_2$ from W. It follows that in all cases, x has distance at most $d_1 + d_2$ from W, so that $U \cap W^*$ lies in a bounded neighbourhood of W as required. As $U \cap W^*$ is contained in W^*, it follows that it lies in bounded neighbourhoods of W and W^*, so that $U \cap W^*$ is small as required.

The definition of $\leqslant$ now shows that $U \leqslant W$, except possibly when two of the four sets $U^{(*)} \cap W^{(*)}$ are small. The only possibility is that $U^* \cap W$ and $U \cap W^*$ are both small. As one must be empty, either $U \subset W$ or $W \subset U$. We conclude that if $U \leqslant V \leqslant W$, then either $U \leqslant W$ or $W \subset U$. Now we consider two cases.

First suppose that $U \subset V \leqslant W$, so that either $U \leqslant W$ or $W \subset U$. If $W \subset U$, then $W \subset V$, so that $W \leqslant V$. As $V \leqslant W$ and $W \leqslant V$, it follows from the first paragraph of the proof of this lemma that $V = W$. Hence, in either case, $U \leqslant W$.

Now consider the general situation when $U \leqslant V \leqslant W$. Again either $U \leqslant W$ or $W \subset U$. If $W \subset U$, then we have $W \subset U \leqslant V$. Now the preceding paragraph implies that $W \leqslant V$. Hence we again have $V \leqslant W$ and $W \leqslant V$ so that $V = W$. Hence $U \leqslant W$ still holds. This completes the proof of the lemma. $\square$

Next we need to verify that the set E with the partial order which we have defined satisfies all the hypotheses of Dunwoody's Theorem B.1.10.

Lemma B.1.15. — *E together with $\leqslant$ satisfies the following conditions.*

(1) *If U, $V \in E$ and $U \leqslant V$, then $V^* \leqslant U^*$,*

(2) *If U, $V \in E$, there are only finitely many $Z \in E$ such that $U \leqslant Z \leqslant V$,*

(3) *If U, $V \in E$, at least one of the four relations $U \leqslant V$, $U \leqslant V^*$, $U^* \leqslant V$, $U^* \leqslant V^*$ holds,*

(4) *If U, $V \in E$, one cannot have $U \leqslant V$ and $U \leqslant V^*$.*

Proof. — Conditions (1) and (3) are obvious from the definition of $\leqslant$ and the hypotheses of Theorem B.1.12.

To prove (4), we observe that if $U \leqslant V$ and $U \leqslant V^*$, then $U \cap V^*$ and $U \cap V$ must both be small. This implies that U itself is small, so that X or X^* must be small. But this contradicts the hypothesis that X is a non-trivial H–almost invariant subset of G.

Finally we prove condition (2). Let $Z = gX$ be an element of E such that $Z \leqslant X$. Recall that, as $Z \cap X^*$ projects to a finite subset of $H\backslash G$, we know that $Z \cap X^*$ lies in a d-neighbourhood of X, for some $d > 0$. If $Z \leqslant X$ but Z is not contained in X, then Z and X are not nested. Now Lemma B.1.11 tells us that if Z is such a set, then g belongs to one of only finitely many double cosets HkH. It follows that if we consider all elements Z of E such that $Z \leqslant X$, we will find either $Z \subset X$, or $Z \cap X^*$ lies in a d-neighbourhood of X, for finitely many different values of d. Hence there is $d_1 > 0$ such that if $Z \leqslant X$ then Z lies in the d_1-neighbourhood of X. Similarly, there is $d_2 > 0$ such that if $Z \leqslant X^*$, then Z lies in the d_2-neighbourhood of X^*. Let d denote the larger of d_1 and d_2. Then for any elements U and V of E with $U \leqslant V$, the set $U \cap V^*$ lies in the d-neighbourhood of each of U, U^*, V and V^*.

Now suppose we are given $U \leqslant V$ and wish to prove condition (2). Choose a point u in U whose distance from U^* is greater than d, choose a point v in V^* whose distance from V is greater than d and choose a path L in Γ joining u to v. If $U \leqslant Z \leqslant V$, then u must lie in Z and v must lie in Z^* so that L must meet δZ. As L is compact, the proof of Lemma B.1.11 shows that the number of such Z is finite. This completes the proof of part 2) of the lemma. □

We are now in a position to prove Theorem B.1.12.

Proof. — To prove the first part, we let E denote the set of all translates of X and X^* by elements of G, let $U \to U^*$ be the involution on E and let the relation $\leqslant$ be defined on E by the condition that $U \leqslant V$ if $U \cap V^*$ is empty or the only small set of the four sets $U^{(*)} \cap V^{(*)}$. Lemmas B.1.14 and B.1.15 show that $\leqslant$ is a partial order on E and satisfies all of Dunwoody's conditions (1)-(4). Hence we can construct a tree T from E. As G acts on E, we have a natural action of G on T. Clearly, G acts transitively on the edges of T. If G acts without inversions, then $G \backslash T$ has a single edge and gives G the structure of an amalgamated free product or HNN decomposition. The stabiliser of the edge of T which corresponds to X is the stabiliser H' of X, so we obtain a splitting of G over H' unless G fixes a vertex of T. Note that as $H \backslash \delta X$ is finite, and H' preserves δX, it follows that H' contains H with finite index as claimed in the theorem. If G acts on T with inversions, we simply subdivide each edge to obtain a new tree T' on which G acts without inversions. In this case, the quotient $G \backslash T'$ again has one edge, but it has distinct vertices. The edge group is H' and one of the vertex groups contains H' with index two. As H has infinite index in G, it follows that in this case also we obtain a splitting of G unless G fixes a vertex of T.

Suppose that G fixes a vertex v of T. As G acts transitively on the edges of T, every edge of T must have one vertex at v, so that all edges of T are adjacent to each other. We will show that this cannot occur. The key hypothesis here is that X is non-trivial.

Let W denote $\{g : gX \leqslant X \text{ or } gX^* \leqslant X\}$, and note that condition 3) of Lemma B.1.15 shows that $W^* = \{g : gX \leqslant X^* \text{ or } gX^* \leqslant X^*\}$. Recall that there is $d_1 > 0$ such that if $Z \leqslant X$ then Z lies in the d_1–neighbourhood of X. If d denotes $d_1 + 1$, and $g \in W$, it follows that $g\delta X$ lies in the d–neighbourhood of X. Let c denote the distance of the identity of G from δX. Then g must lie within the $(c+d)$–neighbourhood of X, for all $g \in W$, so that W itself lies in the $(c+d)$–neighbourhood of X. Similarly, W^* lies in the $(c+d)$–neighbourhood of X^*. Now both X and X^* project to infinite subsets of $H\backslash G$, so G cannot equal W or W^*. It follows that there are elements $U$and V of E such that $U < X < V$, so that $U$and V represent non-adjacent edges of T. This completes the proof that G cannot fix a vertex of T.

To prove the last statement of the first part of Theorem B.1.12, we will simplify notation by supposing that the stabiliser H' of X is equal to H. One of the standard H–almost invariant sets associated to the splitting we have obtained from the action of G on the tree T is the set W in the preceding paragraph. We will show that W is H–almost equal to X. The preceding paragraph shows that W lies in the $(c+d)$–neighbourhood of X, and that W^* lies in the $(c+d)$-neighbourhood of X^*. It follows that W is H–almost contained in X and W^* is H–almost contained in X^*, so that W and X are H–almost equal as claimed. This completes the proof of the first part of Theorem B.1.12.

For the second part, we will simply comment on the modifications needed to the preceding proof. The statement of Lemma B.1.13 remains true though the proof needs a little modification. The statement and proof of Lemma B.1.14 apply unchanged. The statement of Lemma B.1.15 remains true, though the proof needs some minor modifications. Finally the proof of the first part of Theorem B.1.12 applies with minor modifications to show that G acts on a tree T with quotient consisting of k edges in the required way. This completes the proof of Theorem B.1.12. □

B.2. Zero intersection numbers

In this section, we prove our two main results about the case of zero intersection number. First we will need the following little result.

Lemma B.2.1. — *Let G be a finitely generated group which splits over a subgroup H. If the normaliser N of H in G has finite index in G, then H is normal in G.*

Proof. — The given splitting of G over H corresponds to an action of G on a tree T such that $G\backslash T$ has a single edge, and some edge of Thas stabiliser H. Let T' denote the fixed set of H, i.e. the set of all points fixed by H. Then T' is a (non-empty) subtree of T. As N normalises H, it must preserve T', i.e. $NT' = T'$. Suppose that $N \neq G$. As N has finite index in G, we let $e, g_1, \ldots, g_n$ denote a set of coset representatives for N in G, where $n \geqslant 1$. As G acts transitively on T, we have $T = T' \cup g_1T' \cup \cdots \cup g_nT'$. Edges of T' all have stabiliser H, and so edges of g_iT' all have stabiliser $g_iHg_i^{-1}$. As g_i does not lie in N, these stabilisers are distinct so the intersection $T' \cap g_iT'$ contains no edges. The intersection of two subtrees of a tree must be empty or a tree, so it follows that $T' \cap g_iT'$ is empty or a single vertex v_i, for each i. Now N preserves T' and permutes the translates g_iT', so N preserves the collection of all the v_i's. As this collection is finite, N has a subgroup N_1 of finite index such that N_1 fixes a vertex v of T'. As N_1 has finite index in G, it follows that G itself fixes some vertex of T, which contradicts our assumption that our action of G on T corresponds to a splitting of G. This contradiction shows that N must equal G, so that H is normal in G as claimed. □

Recall that if X is a H–almost invariant subset of G associated to a splitting of G, then the set of translates of X and X^* is nested. Equivalently, for every $g \in G$, one of the four sets $X^{(*)} \cap gX^{(*)}$ is empty. We need to consider carefully how it is possible for two of the four sets to be small, and a similar question arises when one considers two splittings of G.

Lemma B.2.2. — *Let G be a finitely generated group with two splittings over finitely generated subgroups H and K with associated H–almost invariant subset X of G and associated K–almost invariant subset Y of G.*

(1) *If two of the four sets $X^{(*)} \cap Y^{(*)}$ are small, then $H = K$.*
(2) *If two of the four sets $X^{(*)} \cap gX^{(*)}$ are small, then g normalises H.*

Proof. — Our first step will be to show that H and K must be commensurable. Without loss of generality, we can suppose that $X \cap Y$ is small. The other small set can only be $X^* \cap Y^*$, as otherwise X or Y would be small which is impossible. It follows that for each edge of δY, either it is also an edge of δX or it has (at least) one end in one of the two small sets. As the images in $H\backslash\Gamma$ of δX and of each small set is finite, and as the graph Γ is locally finite, it follows that the image of δY in $H\backslash\Gamma$ must be finite. This implies that $H \cap K$ has finite index in the stabiliser K of δY. By reversing the roles of H and K, it follows that $H \cap K$ has finite index in H, so that H and K must be commensurable, as claimed.

Now let L denote $H \cap K$, so that L stabilises both X and Y, and consider the images P and Q of X and Y in $L\backslash\Gamma$. As L has finite index in H and K, it follows that δP and δQ are each finite, so that P and Q are almost invariant subsets of $L\backslash G$. Further, two of the four sets $X^{(*)} \cap Y^{(*)}$ have finite image in $L\backslash\Gamma$, so we can assume that P and Q are almost equal, by replacing one of X or Y by its complement in G, if needed. Let L' denote the intersection of the conjugates of L in H, so that L' is normal in H, though it need not be normal in K. We do not have $L' = H \cap K$, but because L has finite index in H, we know that L' has finite index in H and hence also in K, which is all we need. Let P' and Q' denote the images of X and Y respectively in $L'\backslash\Gamma$, and consider the action of an element h of H on $L'\backslash\Gamma$. Trivially $hP' = P'$. As P' and Q' are almost equal, hQ' must be almost equal to Q'. Now we use the key fact that Y is associated to a splitting of G so that its translates by G are nested. Thus for any element g of G, one of the following four inclusions holds: $gY \subset Y$, $gY \subset Y^*$, $gY^* \subset Y$, $gY^* \subset Y^*$. As hQ' is almost equal to Q', we must have $hY \subset Y$ or $hY^* \subset Y^*$. But h has a power which lies in L and hence stabilises Y. It follows that $hY = Y$, so that h lies in K. Thus H is a subgroup of K. Similarly, K must be a subgroup of H, so that $H = K$. This completes the proof of part 1 of the lemma. Note that it follows that $L = H = K$, that $H\backslash X = P$ and $K\backslash Y = Q$ and that P and Q are almost equal or almost complementary.

In order to prove part 2 of the lemma, we apply the preceding work to the case when the second splitting is obtained from the first by conjugating by some element g of G. Thus $K = gHg^{-1}$ and $Y = gXg^{-1}$ which is K–almost equal to gX by Remark B.1.9. Hence if two of the four sets $X^{(*)} \cap gX^{(*)}$ are small, then so are two of the four sets $X^{(*)} \cap Y^{(*)}$ small. Now the above shows that $H = K = gHg^{-1}$, so that g normalises H. This completes the proof of the lemma. □

Lemma B.2.3. — *Let G be a finitely generated group with two splittings over finitely generated subgroups H and K with associated H-almost invariant subset X of G and associated K–almost invariant subset Y of G. If two of the four sets $X^{(*)} \cap Y^{(*)}$ are small, then the two splittings of G are conjugate. Further one of the following holds:*

(1) *the two splittings are equivalent, or*

(2) *the two splittings are of the form $G = L *_H C$, where H has index 2 in L, and the splittings are conjugate by an element of L, or*

(3) *H is normal in G and $H\backslash G$ is isomorphic to $\mathbb{Z}$ or to $\mathbb{Z}_2 * \mathbb{Z}_2$.*

Proof. — The preceding lemma showed that the hypotheses imply that H equals K and also that the images P and Q of X and Y in $H\backslash G$ are almost equal or almost complementary. By replacing one of X or Y by its complement if needed, we can arrange that P and Q are almost equal. We will show that in most cases, the two given splittings over H and K must be equivalent, and that the exceptional cases can be analysed separately to show that the splittings are conjugate.

Recall that by applying Theorem B.1.10, we can use information about X and its translates to construct a G–tree T_X and hence the original splitting of G over H. Similarly, we can use information about Y and its translates to construct a G–tree T_Y and hence the original splitting of G over K. We will compare these two constructions in order to prove our result.

As P and Q are almost equal subsets of $H\backslash G$, it follows that there is $\delta \geqslant 0$ such that, in the Cayley graph Γ of G, we have X lies in a δ-neighbourhood of Y and Y lies in a δ-neighbourhood of X. Now let U_X denote one of X or X^*, let V_X denote one of gX or gX^* and let U_Y and V_Y denote the corresponding sets obtained by replacing X with Y. Recall that $U_X \cap V_X$ is small if and only if its image in $H\backslash G$ is finite. Clearly this occurs if and only if V_X lies in a δ-neighbourhood of U_X^*, for some $\delta \geqslant 0$. It follows that $U_X \cap V_X$ is small if and only if $U_Y \cap V_Y$ is small.

As X and Y are associated to splittings, we know that for each $g \in G$, at least one of the four sets $X^{(*)} \cap gX^{(*)}$ is empty and at least one of the four sets $Y^{(*)} \cap gY^{(*)}$ is empty. Further the information about which of the four sets is empty completely determines the trees T_X and T_Y. Thus we would like to show that when we compare the four sets $X^{(*)} \cap gX^{(*)}$ with the four sets $Y^{(*)} \cap gY^{(*)}$, then corresponding sets are empty. Note that when g lies in H, we have $gX = X$, so that two of the four sets $X^{(*)} \cap gX^{(*)}$ are empty.

First we consider the case when, for each $g \in G-H$, only one of the sets $X^{(*)} \cap gX^{(*)}$ is small and hence empty. Then only the corresponding one of the four sets $Y^{(*)} \cap gY^{(*)}$ is small and hence empty. Now the correspondence $gX \to gY$ gives a G–isomorphism of T_X with T_Y and thus the splittings are equivalent.

Next we consider the case when two of the sets $X^{(*)} \cap gX^{(*)}$ are small, for some $g \in G-H$. Part 2 of Lemma B.2.2 implies that g normalises H. Further if $R = H\backslash gX$, then $P$is almost equal to R or R^*. Let $N(H)$ denote the normaliser of H in G, so that $N(H)$ acts on the left on the graph $H\backslash\Gamma$ and we have $R = gP$. Let L denote the subgroup of $N(H)$ consisting of elements k such that kP is almost equal to P or P^*. Now we apply Theorem 5.8 from [**47**] to the action of $H\backslash L$ on the left on the graph $H\backslash\Gamma$. This result tells us that if $H\backslash L$ is infinite, then it has an infinite cyclic subgroup of finite index. Further the proof of this result in [**47**] shows that the quotient of $H\backslash\Gamma$ by $H\backslash L$ must be finite. This implies that $H\backslash\Gamma$ has two ends and that L has finite index in G. To summarise, either $H\backslash L$ is finite, or it has two ends and L has finite index in G. Let k be an element of L whose image in $H\backslash L$ has finite order such that $kP \overset{a}{=} P$. As X is associated to a splitting of G, we must have $kX \subset X$ or $X \subset kX$. As k has finite order in $H\backslash L$, we have $k^n X = X$, for some positive integer n, which implies that $kX = X$ so that k itself lies in H. It follows that the group $H\backslash L$ must be trivial, $\mathbb{Z}_2$, $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. In the first case, the two trees T_X and T_Y will be G-isomorphic, showing that the given splittings are equivalent. In the other three cases, $L - H$ is non-empty and we know that, for any $g \in L - H$, two of the four sets $X^{(*)} \cap gX^{(*)}$ are small. Thus in these cases, it seems possible that T_X and T_Y will not be G–isomorphic, so we need some special arguments.

We start with the case when $H\backslash L$ is $\mathbb{Z}_2$. In this case, the given splitting must be an amalgamated free product of the form $L *_H C$, for some group C. If k denotes an element of $L - H$, then $kP \overset{a}{=} P^*$. Thus G acts on T_X and T_Y with inversions. Recall that either the two partial orders on the translates of X and Y are the same under the bijection $gX \to gY$, or they differ only in that $kX^* \subset X$ but $Y \subset kY^*$, for all $k \in L - H$. If they differ, we replace the second splitting by its conjugate by some element $k \in L - H$, so that Y is replaced by $Y' = kY$ and we replace X by $X' = X^*$. As Y' is H–almost equal to X', the partial orders on the translates of X' and Y' respectively are the same under the bijection $gX' \to gY'$ except possibly when one compares X', kX' and Y', kY', where $k \in L - H$. In this case, the inclusion $kX^* \subset X$ tells us that $kX' \subset (X')^*$, and the inclusion $Y \subset kY^*$ tells us that $kY' = k^2 Y = Y \subset kY^* = (Y')^*$. We conclude that the partial orders on the translates of X' and Y' respectively are exactly the same, so that T_X and T_Y are G–isomorphic, and the two given splittings are conjugate by an element of L.

Now we turn to the two cases where $H\backslash L$ is infinite, so that L has finite index in G and $H\backslash\Gamma$ has two ends. As L normalises H, Lemma B.2.1 shows that H is normal in G. As $H\backslash\Gamma$ has two ends, it follows that $L = G$, so that $H\backslash G$ is $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. It is easy to check that there is only one splitting of $\mathbb{Z}$ over the trivial group and that all splittings of $\mathbb{Z}_2 * \mathbb{Z}_2$ over the trivial group are conjugate. It follows that, in either case, all splittings of G over H are conjugate. This completes the proof of Lemma B.2.3. □

Lemma B.2.4. — *Let G be a finitely generated group with two splittings over finitely generated subgroups H and K with associated H–almost invariant subset X of G and associated K–almost invariant subset Y of G. Let $E = gX, gX^*, gY, gY^* : g \in G$, and let U and V denote two elements of E such that two of the four sets $U^{(*)} \cap V^{(*)}$ are small. Then either one of the two sets is empty, or the two given splittings of G are conjugate.*

Proof. — Recall that X is associated to a splitting of G over H. It follows that gX is associated to the conjugate of this splitting by g. Thus U and V are associated to splittings of G which are each conjugate to one of the two given splittings. If U and V are each translates of X or X^*, the nestedness of the translates of X shows that one of the two small sets must be empty as claimed. Similarly if both are translates of Y or Y^*, then one of the two small sets must be empty. If U is a translate of X or X^* and V is a translate of Y or Y^*, we apply Lemma B.2.3 to show that the splittings to which U and V are associated are conjugate. It follows that the two original splittings were conjugate as required. □

Now we come to the proof of our first main result.

Theorem B.2.5. — *Let G be a finitely generated group with n splittings over finitely generated subgroups. This collection of splittings is compatible up to conjugacy if and only if each pair of splittings has intersection number zero. Further, in this situation, the graph of groups structure on G obtained from these splittings has a unique underlying graph, and the edge and vertex groups are unique up to conjugacy.*

Proof. — Let the n splittings s_i of G be over subgroups $H_1, \ldots, H_n$ with associated H_i–almost invariant subsets X_i of G, and let $E = gX_i, gX_i^* : g \in G, 1 \leqslant i \leqslant n$. We will start by supposing that no two of the s_i's are conjugate. We will handle the general case at the end of this proof.

We will apply the second part of Theorem B.1.12 to E. Recall that our assumption that the s_i's have intersection number zero implies that no translate of X_i can cross any translate of X_j, for $1 \leqslant i \neq j \leqslant n$. As each X_i is associated to a splitting, it is also true that no translate of X_i can cross any translate of X_i. This means that the set E is almost nested. In order to apply Theorem B.1.12, we will also need to show that for any pair of elements U and V of E, if two of the four sets $U^{(*)} \cap V^{(*)}$ are small then one is empty. Now Lemma B.2.4 shows that if two of these four sets are small, then either one is empty or there are distinct i and j such that s_i and s_j are conjugate. As we are assuming that no two of these splittings are conjugate, it follows that if two of the four sets $U^{(*)} \cap V^{(*)}$ are small then one is empty, as required.

Theorem B.1.12 now implies that G can be expressed as the fundamental group of a graph Γ of groups whose i–th edge corresponds to a conjugate of a splitting of G over the stabilizer H_i' of X_i. As X_i is associated to a splitting of G over H_i, its stabiliser H_i' must equal H_i. Further, it is clear from the construction that collapsing all but the i–th edge of Γ yields a conjugate of s_i, as the corresponding G-tree has edges which correspond precisely to the translates of X_i.

Now suppose that we have a graph of groups structure Γ' for G such that, for each i, $1 \leqslant i \leqslant n$, collapsing all edges but the i–th yields a conjugate of the splitting s_i of G. This determines an action of G on a tree T' without inversions. We want to show that T and T' are G-isomorphic. For this implies that Γ and Γ' have the same underlying graph, and that corresponding edge and vertex groups are conjugate, as required. Let e denote an edge of T', and let $Y(e)$ denote $\{g \in G : ge < e \text{ or } g\overline{e} < e\}$. There are edges e_i of T', $1 \leqslant i \leqslant n$, such that the set E' of all translates of $Y(e_i)$ and $Y(e_i)^*$ is nested and Dunwoody's construction applied to E' yields the G-tree T' again. We will denote $Y(e_i)$ by Y_i. The hypotheses imply that there is $k \in G$ such that the stabiliser K_i of e_i equals $k^{-1}H_ik$, and that Y_i is K_i–almost equal to $k^{-1}X_ik$, where X_i is one of the standard H_i–almost invariant subsets of G associated to the splitting s_i. Let Z_i denote kY_i so that Z_i is H_i–almost equal to X_ik. Now Remark B.1.9 shows that X_ik is H_i–almost equal to X_i, so that Z_i is H_i–almost equal to X_i. Now consider the G–equivariant bijection $E \to E'$ determined by sending X_i to Z_i. The above argument shows that if U is any element of E, and U' is the corresponding element of E', then U and U' are $stab(U)$–almost equal. We will show that in most cases, this bijection automatically preserves the partial orders on E and E', implying that T and T' are G–isomorphic, as required. We compare the partial orders on E and E' rather as in the proof of Lemma B.2.3.

For any elements U and V of E, let U' and V' denote the corresponding elements of E'. Thus $U \cap V$ is small if and only if $U' \cap V'$ is small. We would like to show that when we compare the four sets $U^{(*)} \cap V^{(*)}$ with the four sets $U'^{(*)} \cap V'^{(*)}$, then corresponding sets are empty, so that the partial orders are preserved by our bijection. Otherwise, there must be U and V in E such that two of the sets $U^{(*)} \cap V^{(*)}$ are small. If U and V are translates of X_i and X_j, then Lemma B.2.3 tells us that the splittings s_i and s_j are conjugate. As we are assuming that distinct splittings are not conjugate, it follows that $i = j$. Now the arguments in the proof of Lemma B.2.3 show that either the splitting s_i is an amalgamated free product of the form $L *_H C$, with $|L : H| = 2$, or H is normal in G and $H\backslash G$ is $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. If the second case occurs, then there can be only one splitting in the given family, so it is immediate that

Γ and Γ' have the same underlying graph, and that corresponding edge and vertex groups are conjugate. If the first case occurs and the partial orders on translates of X_i and Z_i do not match, we must have $lX_i^* \subset X_i$ but $Z_i \subset lZ_i^*$, for all $l \in L - H$. We now pick $l \in L - H$ and alter our bijection from E to E' so that X_i maps to $W_i = lZ_i^*$ and extend G–equivariantly to the translates of X_i and X_i^*. This ensures that the partial orders on E and E' match for translates of X_i. By repeating this for other values of i as necessary, we can arrange that the partial orders match completely, and can then conclude that T and T' are G–isomorphic as required.

We end by discussing the case when some of the given n splittings are conjugate. We divide the splittings into conjugacy classes and discard all except one splitting from each conjugacy class, to obtain k splittings. Now we apply the preceding argument to express G uniquely as the fundamental group of a graph Γ of groups with k edges. If an edge of Γ corresponds to a splitting over a subgroup H which is conjugate to $r-1$ other splittings, we simply subdivide this edge into r sub-edges, and label all the sub-edges and the $r-1$ new vertices by H. This shows the existence of the required graph of groups structure Γ' corresponding to the original n splittings. The uniqueness of Γ' follows from the uniqueness of Γ, and the fact that the collection of all the edges of Γ' which correspond to a given splitting of G must form an interval in Γ' in which all the interior vertices have valence 2. This completes the proof of Theorem B.2.5. $\square$

Now we turn to the proof of Theorem B.2.8 that splittings exist. It will be convenient to make the following definitions. We will use H^g to denote gHg^{-1}.

Definition B.2.6. — If X is a H–almost invariant subset of G and Y is a K–almost invariant subset of G, and if X and Y are H–almost equal, then we will say that X and Y are equivalent and write $X \sim Y$. (Note that H and K must be commensurable.)

Definition B.2.7. — If H is a subgroup of a group G, the commensuriser in G of H consists of those elements g in G such that H and H^g are commensurable subgroups of G. The commensuriser is clearly a subgroup of G and is denoted by $\mathrm{Comm}_G(H)$ or just $Comm(H)$, when the group G is clear from the context.

Now we come to the proof of our second main result.

Theorem B.2.8. — *Let G be a finitely generated group with a finitely generated subgroup H, such that $e(G,H) \geqslant 2$. If there is a non-trivial H–almost invariant subset X of G such that $i(H\backslash X, H\backslash X) = 0$, then G has a splitting over some subgroup H' commensurable with H. Further, one of the H'–almost invariant sets Y determined by the splitting is equivalent to X.*

Remark B.2.9. — This is the best possible result of this type, as it is clear that one cannot expect to obtain a splitting over H itself. For example, suppose that H is carried by a proper power of a two-sided simple closed curve on a closed surface whose fundamental group is G, so that $e(G,H) = 2$. There are essentially only two *non-trivial almost invariant subsets of* $H\backslash G$, each with vanishing self-intersection number, but there is no splitting of G over H.

Proof. — The idea of the proof is much as before. We let P denote the almost invariant subset $H\backslash X$ of $H\backslash G$, and let E denote $gX, gX^* : g \in G$. We want to apply the first part of Theorem B.1.12. As before, the assumption that $i(P,P) = 0$ implies that E is almost nested. However, in order to apply Theorem B.1.12, we also need to know that for any pair of elements U and V of E, if two of the four sets $U^{(*)} \cap V^{(*)}$ are small then one is empty. In the proof of Theorem B.2.5, we simply applied Lemma B.2.4. However, here the situation is somewhat more complicated. Lemma B.2.10 below shows that if $X \cap gX^*$ and $gX \cap X^*$ are both small, then g must lie in a certain subgroup $\mathcal{K}$ of $\mathrm{Comm}_G(H)$. Thus it would suffice to arrange that E is nested with respect to $\mathcal{K}$, i.e. that gX and X are nested so long as g lies in $\mathcal{K}$. Now Proposition B.2.14 below tells us that there is a subgroup H' commensurable with H and a H'–almost invariant set Y equivalent to X such that $E' = gY, gY^* : g \in G$ is nested with respect to $\mathcal{K}$. It follows that if U and V are any elements of E' and if $U \cap V^*$ and $V \cap U^*$ are both small, then one of them is empty. We also claim that, like E, the set E' is almost nested. This means that if we let P' denote $H'\backslash Y$, we are claiming that $i(P',P') = 0$. Let H'' denote $H \cap H'$. The fact that Y is equivalent to X means that the pre-images in $H''\backslash G$ of P and of P' are almost equal almost invariant sets which we denote by Q and Q'. If d denotes the index of H' in H, then $i(Q,Q) = d^2 i(P,P) = 0$ and similarly $i(Q',Q')$ is an integral multiple of $i(P',P')$. As Q and Q' are almost equal, it follows that $i(Q',Q') = i(Q,Q)$, and hence that $i(P',P') = 0$ as claimed. This now allows us to apply Theorem B.1.12 to the set E'. We conclude that G splits over the stabiliser H'' of Y, that H'' contains H' with finite index and that one of the H''–almost invariant sets associated to the splitting is equivalent to X'. It follows that H'' is commensurable with H and that one of the H''–almost invariant sets determined by the splitting is equivalent to X. This completes the proof of Theorem B.2.8 apart from the proofs of Lemma B.2.10 and Proposition B.2.14. □

It remains to prove the two results we just used. The proofs do not use the hypothesis that the set of all translates of X and X^* are almost nested. Thus for the rest of this section, we will consider the following general situation.

Let G be a finitely generated group with a finitely generated subgroup H such that $e(G,H) \geqslant 2$, and let X denote a non-trivial H–almost invariant subset of G.

Recall that our problem in the proof of Theorem B.2.8 is the possibility that two of the four sets $X^{(*)} \cap gX^{(*)}$ are small. As this would imply that $gX \sim X$ or X^*, it is clear that the subgroup $\mathcal{K}$ of G defined by $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$ is very relevant to our problem. We will consider this subgroup carefully. Here is the first result we quoted in the proof of Theorem B.2.8.

Lemma B.2.10. — *If $\mathcal{K} = \{g \in G : gX \sim X$ or $X^*\}$, then $H \subset \mathcal{K} \subset \mathrm{Comm}_G(H)$.*

Proof. — The first inclusion is clear. The second is proved in essentially the same way as the proof of the first part of Lemma B.2.3. Let g be an element of $\mathcal{K}$, and consider the case when $gX \sim X$ (the other case is similar). Recall that this means that the sets $X \cap gX^*$ and $X^* \cap gX$ are both small. Now for each edge of δgX, either it is also an edge of δX or it has (at least) one end in one of the two small sets. As the images in $H\backslash\Gamma$ of δX and of each small set is finite, and as the graph Γ is locally finite, it follows that the image of δgX in $H\backslash\Gamma$ must be finite. This implies that $H \cap H^g$ has finite index in the stabiliser H^g of δgX. By reversing the roles of X and gX, it follows that $H \cap H^g$ has finite index in H, so that H and H^g must be commensurable, as claimed. It follows that $\mathcal{K} \subset \mathrm{Comm}_G(H)$, as required. □

Another way of describing our difficulty in applying Theorem B.1.12 is to say that it is caused by the fact that the translates of X and X^* may not be nested. However, Lemma B.1.11 assures us that "most" of the translates are nested. The following result gives us a much stronger finiteness result.

Lemma B.2.11. — *Let G, H, X, $\mathcal{K}$ be as above. Then $\{g \in \mathcal{K} : gX$ and X are not nested$\}$ consists of a finite number of right cosets gH of H in G.*

Proof. — Lemma B.1.11 tells us that the given set is contained in the union of a finite number of double cosets HgH. If $k \in \mathcal{K}$, we claim that the double coset HkH is itself the union of only finitely many cosets gH, which proves the required result. To prove our claim, recall that $k^{-1}Hk$ is commensurable with H. Thus $k^{-1}Hk$ can be expressed as the union of cosets $g_i(k^{-1}Hk \cap H)$, for $1 \leqslant i \leqslant n$. Hence

$$HkH = k(k^{-1}Hk)H = k\left(\cup_{i=1}^n g_i(k^{-1}Hk \cap H)\right)H = k\left(\cup_{i=1}^n g_iH\right) = \cup_{i=1}^n kg_iH,$$

so that HkH is the union of finitely many cosets gH as claimed. □

Now we come to the key result.

Lemma B.2.12. — *Let G, H, X, $\mathcal{K}$ be as above. Then there are a finite number of finite index subgroups $H_1, \dots, H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1 \leqslant i \leqslant m$, where $N(H_i)$ denotes the normaliser of H_i in G.*

Proof. — Consider an element g in $\mathcal{K}$. Lemma B.2.10 tells us that H and H^g are commensurable subgroups of G. Let L denote their intersection and let L' denote the intersection of the conjugates of L in H. Thus L' is of finite index in H and H^g and is normal in H. Now consider the quotient $L'\backslash G$. Let P and Q denote the images of X and gX respectively in $L'\backslash G$. As before, P and Q are almost invariant subsets of $L'\backslash G$ which are almost equal or almost complementary. Now consider the action of $L'\backslash H$ on the left on $L'\backslash G$. If h is in H, then $hP = P$, so that $hQ \overset{a}{=} Q$. If $h(gX)$ and gX are nested, there are four possible inclusions, but the fact that $hQ \overset{a}{=} Q$ excludes two of them. Thus we must have $hQ \subset Q$ or $Q \subset hQ$. This implies that $hQ = Q$ as some power of h lies in L' and so acts trivially on $L'\backslash G$. We conclude that if h is an element of $H-L'$ such that $h(gX)$ and gX are nested, then h stabilises gX and so lies in H^g. Hence h lies in L. It follows that for each element h of $H-L$, the sets $h(gX)$ and gX are not nested. Recall from Lemma B.2.11 that $\{g \in \mathcal{K} : gX$ and X are not nested$\}$ consists of a finite number of cosets gH of H in G. It will be convenient to denote this number by $d-1$. Thus, for $g \in \mathcal{K}$, the set $\{h \in \mathcal{K} : h(gX)$ and gX are not nested$\}$ consists of $d-1$ cosets hH^g of H^g in G. It follows that $H-L$ lies in the union of $d-1$ cosets hH^g of H^g in G. As $L = H \cap H^g$, it follows that $H-L$ lies in the union of $d-1$ cosets hL of L in G and hence that L has index at most d in H.

A similar argument shows also that L has index at most d in H^g. Of course, the same bound applies to the index of $H \cap H^{g^i}$ in H, for each i. Now we define $H' = \cap_{i\in\mathbb{Z}} H^{g^i}$. Clearly H' is a subgroup of H which is normalised by g. Now each intersection $H \cap H^{g^i}$ has index at most d in H, and so $H' = \cap_{i\in\mathbb{Z}} \left(H \cap H^{g^i}\right)$ is an intersection of subgroups of H of index at most d. If H has n subgroups of index at most d, it follows that H' has index at most d^n in H. Hence each element of $\mathcal{K}$ normalises a subgroup of H of index at most d^n in H. As H has only finitely many such subgroups, we have proved that there are a finite number of finite index subgroups $H_1, \dots, H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1 \leqslant i \leqslant m$, as required. □

Using this result, we can prove the following.

Lemma B.2.13. — *Let G, H, X, $\mathcal{K}$ be as above. Then there is a subgroup H' of finite index in H, such that $\mathcal{K}$ normalises H'.*

Proof. — We will consider how $\mathcal{K}$ can intersect the normaliser of a subgroup of finite index in H. Let H_1 denote a subgroup of H of finite index. We denote the image of X in $H_1\backslash G$ by P. Then P is an almost invariant subset of $H_1\backslash G$. We consider the group $\mathcal{K}\cap N(H_1)$, which we will denote by $\mathcal{K}_1$. Then $H_1\backslash\mathcal{K}_1$ acts on the left on $H_1\backslash G$, and we have $kP \overset{a}{=} P$ or P^*, for every element k of $H_1\backslash\mathcal{K}_1$, because every element of $\mathcal{K}$ satisfies $kX \sim X$ or X^*. Now we apply Theorem 5.8 from [**47**] to the action of $H_1\backslash\mathcal{K}_1$ on the left on the graph $H_1\backslash\Gamma$. This result tells us that if $H_1\backslash\mathcal{K}_1$ is infinite, then it has an infinite cyclic subgroup of finite index. Further the proof of this result in [**47**] shows that the quotient of $H_1\backslash\Gamma$ by $H_1\backslash\mathcal{K}_1$ must be finite. This implies that $H_1\backslash\Gamma$ has two ends and that $\mathcal{K}_1$ has finite index in G. Hence either $H_1\backslash\mathcal{K}_1$ is finite, or it has two ends and $\mathcal{K}_1$ has finite index in G.

Recall that there are a finite number of finite index subgroups $H_1, \dots, H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1 \leqslant i \leqslant m$. The above discussion shows that, for each i, if $\mathcal{K}_i$ denotes $\mathcal{K}\cap N(H_i)$, either $H_i\backslash\mathcal{K}_i$ is finite, or it has two ends and $\mathcal{K}_i$ has finite index in G. We consider two cases depending on whether or not every $H_i\backslash\mathcal{K}_i$ is finite.

Suppose first that each $H_i\backslash\mathcal{K}_i$ is finite. We claim that $\mathcal{K}$ contains H with finite index. To see this, let $H'' = \cap H_i$, so that H'' is a subgroup of H of finite index, and note that $\mathcal{K}$ is the union of a finite collection of groups $\mathcal{K}_i$ each of which contains H'' with finite index, so that $\mathcal{K}$ is the union of finitely many cosets of H''. It follows that $\mathcal{K}$ also contains H'' with finite index and hence contains H with finite index as claimed. If we let H' denote the intersection of the conjugates of H in $\mathcal{K}$, then H' is the required subgroup of H which is normalised by $\mathcal{K}$.

Now we turn to the case when $H_1\backslash\mathcal{K}_1$ is infinite and so $H_1\backslash\mathcal{K}_1$ has two ends and $\mathcal{K}_1$ has finite index in G. Define H' to be $\cap_{k\in\mathcal{K}}(H_1)^k$. As $\mathcal{K}$ contains $\mathcal{K}_1$ with finite index, H' is the intersection of only finitely many conjugates of H_1. As $\mathcal{K}$ is contained in $Comm(H)$, each of these conjugates of H_1 is commensurable with H_1. It follows that H' is a subgroup of H of finite index in H which is normalised by $\mathcal{K}$. This completes the proof of the lemma. □

The key point here is that $\mathcal{K}$ normalises H' rather than just commensurises it. Now we can prove the second result which we quoted in the proof of Theorem B.2.8.

Proposition B.2.14. — *Suppose that (G, H) is a pair of finitely generated groups and that X is a non-trivial H–almost invariant subset of G. Then, there is a subgroup H' of G which is commensurable with H, and a non-trivial H'–almost invariant set Y equivalent to X such that $gY, gY^* : g \in G$ is nested with respect to the subgroup $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$ of G.*

Proof. — The previous lemma tells us that there is a subgroup H' of finite index in H such that $\mathcal{K}$ normalises H'. Let P denote the almost invariant subset $H\backslash X$ of $H\backslash G$, and let P' denote the almost invariant subset $H'\backslash X$ of $H'\backslash G$.

Suppose that the index of H' in $\mathcal{K}$ is infinite. Recall from the proof of the preceding lemma that $H'\backslash\mathcal{K}$ has two ends and that $\mathcal{K}$ has finite index in G. We construct a new non-trivial H''–almost invariant set Y as follows. Since the quotient group $H'\backslash\mathcal{K}$ has two ends, $\mathcal{K}$ splits over a subgroup H'' which contains H' with finite index. Thus there is a H''-almost invariant set X'' in $\mathcal{K}$ which is nested with respect to $\mathcal{K}$. Further, H'' is normal in $\mathcal{K}$ and the quotient group must be isomorphic to $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. Let $g_1 = e, g_2, \ldots, g_n$ be coset representatives of $\mathcal{K}$ in G so that $G = \cup_i \mathcal{K} g_i$. We take $Y = \cup_i X'' g_i$. It is easy to check that Y is H''–almost invariant and that $gY, gY^* : g \in G$ is nested with respect to $\mathcal{K}$.

Now suppose that the index of H' in $\mathcal{K}$ is finite. We will define the subgroup $\mathcal{K}_0 = g \in G : gX \sim X$ of $\mathcal{K}$. The index of $\mathcal{K}_0$ in $\mathcal{K}$ is at most two.

First we consider the case when $\mathcal{K} = \mathcal{K}_0$. We define P'' to be the intersection of the translates of P' under the action of $H'\backslash\mathcal{K}$. Thus P'' is invariant under the action of $H'\backslash\mathcal{K}$. As all the translates of P' by elements of $H'\backslash\mathcal{K}$ are almost equal to P', it follows that $P'' \overset{a}{=} P'$ so that P'' is also an almost invariant subset of $H'\backslash G$. Let Y

denote the inverse image of P'' in G, so that Y is invariant under the action of $\mathcal{K}$. In particular, $gY, gY^* : g \in G$ is nested with respect to $\mathcal{K}$, as required.

Now we consider the general case when $\mathcal{K} \neq \mathcal{K}_0$. We can apply the above arguments using $\mathcal{K}_0$ in place of $\mathcal{K}$ to obtain a subgroup H'' of G and a H''–almost invariant subset Y of G which is equivalent to X, and whose translates are nested with respect to $\mathcal{K}_0$. We also know that Y is $\mathcal{K}_0$–invariant. Let Q denote the image of Y in $\mathcal{K}_0\backslash G$, let k denote an element of $\mathcal{K} - \mathcal{K}_0$ and consider the involution of $\mathcal{K}_0\backslash G$ induced by k. Then Q is a non-trivial almost invariant subset of $\mathcal{K}_0\backslash G$ and $kQ \overset{a}{=} Q^*$. Define $R = Q - kQ$, so that $R \overset{a}{=} Q$ and let Z denote the pre-image of R in G. We claim that the translates of Z and Z^* are nested with respect to $\mathcal{K}$. First we show that they are nested with respect to $\mathcal{K}_0$, by showing that $Z = Y - kY$ is $\mathcal{K}_0$-invariant. For $k_0 \in \mathcal{K}_0$, we have $k^{-1}k_0k \in \mathcal{K}_0$ as $\mathcal{K}_0$ must be normal in $\mathcal{K}$. It follows that $k_0kY = kY$. As $k_0Y = Y$, we see that Z is $\mathcal{K}_0$–invariant as required. In order to show that the translates of Z and Z^* are nested with respect to $\mathcal{K}$, we will also show that $Z \cap kZ$ is empty. This follows from the fact that $R \cap kR = (Q - kQ) \cap k(Q - kQ) = (Q - kQ) \cap (kQ - Q)$ which is clearly empty.

This completes the proof of Proposition B.2.14. □

B.3. Strong intersection numbers

Let G be a finitely generated group and let H and K be subgroups of G. Let X be a non-trivial H–almost invariant subset of G and let Y be a non-trivial K–almost invariant subset of G. In section B.1, we discussed what it means for X to cross Y and the fact that this is symmetric. As mentioned in the introduction, there is an alternative way to define crossing of almost invariant sets. Recall that, in section B.1, we introduced our definition of crossing by discussing curves on surfaces. Thus it seems natural to discuss the crossing of X and Y in terms of their boundaries. We call this strong crossing. However, this leads to an asymmetric intersection number. In this section, we define strong crossing and discuss its properties and some applications.

We consider the Cayley graph Γ of G with respect to a finite system of generators. We will usually assume that H and K are finitely generated though this does not seem necessary for most of the definitions below. We will also think of δX as a set of edges in Γ or as a set of points in G, where the set of points will simply be the collection of endpoints of all the edges of δX.

Definition B.3.1. — We say that Y *crosses* X *strongly* if both $\delta Y \cap X$ and $\delta Y \cap X^*$ project to infinite sets in $H\backslash G$.

Remark B.3.2. — This definition is independent of the choice of generators for G which is used to define Γ. Clearly, if Y crosses X strongly, then Y crosses X.

Strong crossing is not symmetric. For an example, one need only consider an essential two-sided simple closed curve S on a compact surface F which intersects a simple arc L transversely in a single point. Let G denote $\pi_1(F)$, and let H and K respectively denote the subgroups of G carried by S and L, so that H is infinite cyclic and K is trivial. Then S and L each define a splitting of G over H and K respectively. Let X and Y denote associated standard H–almost invariant and K–almost invariant subsets of G. These correspond to submanifolds of the universal cover of F bounded respectively by a line $\widetilde{S}$ lying above S and by a compact interval $\widetilde{L}$ lying above L, such that $\widetilde{S}$ meets $\widetilde{L}$ transversely in a single point. Clearly, X crosses Y strongly but Y does not cross X strongly.

However, a strong intersection number can be defined as before. It is usually asymmetric, but we will be particularly interested in the case of self-intersection numbers when this asymmetry will not arise.

Definition B.3.3. — The strong intersection number $si(H\backslash X, K\backslash Y)$ is defined to be the number of double cosets KgH such that gX crosses Y strongly. In particular, $si(H\backslash X, H\backslash X) = 0$ if and only if at least one of $\delta gX \cap X$ and $\delta gX \cap X^*$ is H–finite, for each $g \in G$.

Remark B.3.4. — If s and t are splittings of a group G over subgroups H and K, with associated almost invariant subsets X and Y of G, it is natural to say that s crosses t strongly if $si(H\backslash X, K\backslash Y) \neq 0$. It is easy to show that this is equivalent to the idea introduced by Sela [**49**] that s is hyperbolic with respect to t.

Remark B.3.2 shows that $si(H\backslash X, H\backslash X) \leqslant i(H\backslash X, H\backslash X)$. Recall that Theorem B.2.8 shows that if $i(H\backslash X, H\backslash X) = 0$, then G splits over a subgroup H' commensurable with H. Thus the vanishing of the strong self-intersection number may be considered as a first obstruction to splitting G over some subgroup related to H. We will show in Corollary B.3.11 that the vanishing of the strong self-intersection number has a nice algebraic formulation. This is that when $si(H\backslash X, H\backslash X)$ vanishes, we can find a subgroup K of G, commensurable with H, and a K–almost invariant subset Y of G which is nested with respect to $\mathrm{Comm}_G(H) = \mathrm{Comm}_G(K)$. However, Y may be very different from X. This leads to some splitting results when we place further restrictions on H.

Proposition B.3.5. — *Let G be a finitely generated group with finitely generated subgroup H, and let X be a non-trivial H–almost invariant subset of G. Then $si(H\backslash X, H\backslash X) = 0$ if and only if there is a subset Y of G which is H–almost equal to X (and hence H–almost invariant) such that $HYH = Y$.*

Proof. — Suppose that there exists a subset Y of G which is H–almost equal to X, such that $HYH = Y$. We have

$$si(H\backslash X, H\backslash X) = si(H\backslash Y, H\backslash Y),$$

as X and Y are H–almost equal. So, it is enough to show that for every $g \in G$, either $g\delta Y \cap X$ or $g\delta Y \cap X^*$ is H–finite. Suppose that $g \in Y$. Consider $\delta Y \cap Y$ which is a union of a finite number of right cosets Hg_i, $1 \leqslant i \leqslant n$. Since $g \in Y$, $gH \subset Y$. For any $h \in H$, $d(gh, ghg_i) = d(1, g_i)$. Thus $g\delta Y$ is at a bounded distance from Y and hence $g\delta Y \cap Y^*$ has finite image in $H\backslash G$. Similarly, if $g \in Y^*$, $g\delta Y \cap Y$ projects to a finite set in $H\backslash G$.

For the converse, suppose that $si(H\backslash X, H\backslash X) = 0$ and let π denote the projection from G to $H\backslash G$. By hypothesis, $\pi(g\delta X) \cap (H\backslash X)$ or $\pi(g\delta X) \cap (H\backslash X^*)$ is finite. The proof of Lemma B.1.15 tells us that there is a positive number d such that, for every $g \in G$, the set $g\delta X$ is contained in a d-neighbourhood of X or X^*. Let $V = N(X, d)$, the d–neighbourhood of X and let $Y = g|g(\delta X) \subset V$. If $g \in Y$ and $h \in H$, then $hg\delta X \subset hV = V$ and thus $HY = Y$. If $g \in Y$ and $h \in H$, then $gh(\delta X) = g(\delta X) \subset V$ and thus $YH = Y$. It only remains to show that Y is H–almost equal to X. This is essentially shown in the third and fourth paragraphs of the proof of Theorem B.1.12. □

Definition B.3.6. — We will say that a pair of finitely generated groups (G, H) is of *surface type* if $e(G, H') = 2$ for every subgroup H' of finite index in H and $e(G, H') = 1$ for every subgroup H' of infinite index in H.

This terminology is suggested by the dichotomy in [**43**]. Note that for such pairs any two non-trivial H–almost invariant sets in G are H–almost equal or H–almost complementary. We will see that for pairs of surface type, strong and ordinary intersection numbers are equal.

Proposition B.3.7. — *Let (G, H) be a pair of surface type, let X be a non-trivial H–almost invariant subset of G and let Y be a non-trivial K–almost invariant subset of G for some subgroup K of G. Then Y crosses X if and only if Y crosses X strongly.*

Proof. — Let Γ be the Cayley graph of G with respect to a finite system of generators and let $P = H\backslash X$. As in the proof of Lemma B.1.11, for a set S of vertices in a graph, we let $\overline{S}$ denote the maximal subgraph with vertex set equal to S. We will show that exactly one component of $\overline{X}$ has infinite image in $H\backslash\Gamma$. Note that $\overline{P}$ has exactly one infinite component as $H\backslash\Gamma$ has only two ends. Let Q denote the set of vertices of the infinite component of $\overline{P}$ and let W denote the inverse image of Q in G. If $\overline{W}$ has components with vertex set L_i, then we have $\cup\delta(L_i) = \delta W \subseteq \delta X$. Let L denote the vertex set of a component of $\overline{W}$, and let H_L be the stabilizer in H of L. Since δQ is finite, we see that $H_L\backslash\delta L$ is finite. Hence $H_L\backslash\Gamma$ has more than one end. Now our hypothesis that (G, H) is of surface type implies that H_L has finite index in H and thus $H_L\backslash\delta W$ is finite. If $H_L \neq H$, we see that $H_L\backslash\delta W$ divides $H_L\backslash\Gamma$ into at least three infinite components. Thus $H_L = H$ and so $\overline{W}$ is connected. The other components of $\overline{X}$ have finite image in $H\backslash\Gamma$. Similarly, exactly one component of $\overline{X^*}$ has infinite image in $H\backslash\Gamma$. The same argument shows that for any finite subset D of $H\backslash\Gamma$ containing δP, the two infinite components of $((H\backslash\Gamma) - D) \cap P$ and$((H\backslash\Gamma) - D) \cap P^*$ have connected inverse images in Γ.

Recall that if Y crosses X strongly, then Y crosses X. We will next show that if Y does not cross X strongly, then Y does not cross X. Suppose that $\delta Y \cap X$ projects to a finite set in $H\backslash\Gamma$. Take a compact set D in $H\backslash\Gamma$ large enough to contain $\delta Y \cap X$ and δP. By the argument above, if R is the infinite component of $((H\backslash\Gamma) - D) \cap P$, then its inverse image Z is connected and is contained in $\overline{X}$. Any two points in Z can be connected by a path in Z and thus the path does not intersect δY. Thus Z is contained in Y or Y^*. Hence $Z \cap Y$ or $Z \cap Y^*$ is empty. Suppose that $Z \cap Y$ is empty. Then $Z^* \supseteq Y$. Since $Z^* \cap X$ projects to a finite set, we see that $Y \cap X$ projects to a finite set. Similarly, if $Z \cap Y^*$ is empty, then $Y^* \cap X$ projects to a finite set in $H\backslash G$. Thus, we have shown that if $\delta Y \cap X$ projects to a finite set, then either $Y \cap X$ or $Y^* \cap X$ projects to finite set. Thus Y does not cross X. □

From the above proposition and the fact that ordinary crossing is symmetric, we deduce:

Corollary B.3.8. — *If (G, H) and (G, K) are both of surface type and X is a non-trivial H–almost invariant set in G, and Y is a non-trivial K–almost invariant set in G then $si(H\backslash X, K\backslash Y) = i(H\backslash X, K\backslash Y)$. In particular $i(H\backslash X, H\backslash X) = 0$ if and only if $si(H\backslash X, H\backslash X) = 0$.*

Let K be a Poincaré duality group of dimension $(n-1)$ which is a subgroup of a Poincaré duality group G of dimension n. Thus the pair (G, K) is of surface type. In [**29**], Kropholler and Roller defined an obstruction $sing(K)$ to splitting G over a subgroup commensurable with K. Their main result was that $sing(K)$ vanishes if and only if G splits over a subgroup commensurable with K. At an early stage in their proof, they showed that $sing(K)$ vanishes if and only if there is a K–almost invariant subset Y of G such that $KYK = Y$. Starting from this point, Proposition B.3.5, the above Corollary and then Theorem B.2.8 give an alternative proof of their splitting result. Thus Theorem B.2.8 may be considered as a generalization of their splitting theorem. We next reformulate in our language a conjecture of Kropholler and Roller [**30**] :

Conjecture B.3.9. — *If G is a finitely generated group with a finitely generated subgroup H, and if X is a non-trivial H–almost invariant subset of G such that $si(H\backslash X, H\backslash X) = 0$, then G splits over a subgroup commensurable with a subgroup of H.*

Note that Theorem B.2.8 has a stronger hypothesis than this conjecture, namely the vanishing of the self-intersection number $i(H\backslash X, H\backslash X)$, rather than the vanishing of the strong self-intersection number, and it has a correspondingly stronger conclusion, namely that G splits over a subgroup commensurable with H itself. A key difference between the two statements is that, in the above conjecture, one does not expect the almost invariant set associated to the splitting of G to be at all closely related to X. Dunwoody and Roller proved this conjecture when H is virtually polycyclic [**13**], and Sageev [**38**] proved it for quasiconvex subgroups of hyperbolic groups. The paper of Dunwoody and Roller [**13**] contains information useful in the general case. The second step in their proof, which uses a theorem of Bergman [**3**], proves the following result, stated in our language. (There is an exposition of Bergman's argument and parts of [**13**] in the later versions of [**15**].)

Theorem B.3.10. — *Let (G, H) be a pair of finitely generated groups, and let X be a H–almost invariant subset of G. If $si(H\backslash X, H\backslash X) = 0$, then there is a subgroup H' commensurable with H, and a non-trivial H'–almost invariant set Y with $si(H'\backslash Y, H'\backslash Y) = 0$ such that the set $gY, gY^* : g \in G$ is almost nested with respect to $\mathrm{Comm}_G(H) = \mathrm{Comm}_G(H')$.*

This combined with Proposition B.2.14 gives:

Corollary B.3.11. — *With the hypotheses of the above theorem we can choose H' and a non-trivial H'–almost invariant set Y with $si(H'\backslash Y, H'\backslash Y) = 0$ such that $gY, gY^* : g \in G$ is almost nested with respect to $\mathrm{Comm}_G(H)$ and is nested with respect to the subgroup $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$ of $\mathrm{Comm}_G(H)$.*

Now Theorem B.1.12 yields the following generalization of Stallings' Theorem [**53**] already noted by Dunwoody and Roller [**13**]:

Theorem B.3.12. — *If G, H are finitely generated groups with $e(G, H) > 1$ and if G commensurises H, then G splits over a subgroup commensurable with H.*

Corollary B.3.11 leads to the following partial solution of the above conjecture of Kropholler and Roller :

Theorem B.3.13. — *If G, H are finitely generated groups with $e(G, H) > 1$, if $e(G, K) = 1$ for every subgroup K commensurable with a subgroup of infinite index in H, and if X is a H–almost invariant subset of G such that $si(H\backslash X, H\backslash X) = 0$, then G splits over a subgroup commensurable with H.*

Proof. — Observe that Corollary B.3.11 shows that, by changing H up to commensurability, and changing X, we may assume that the translates of X are almost nested with respect to $\mathrm{Comm}_G(H)$ and nested with respect to $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$. If we do not have almost nesting for all translates of X, then there is g outside $\mathrm{Comm}_G(H)$ such that none of $X^{(*)} \cap gX^{(*)}$ is H–finite. In particular, none of these sets is $(H \cap H^g)$–finite. But these four sets are each invariant under $H \cap H^g$ and the fact that the strong intersection number vanishes shows that at least one of them has boundary which is $(H \cap H^g)$–finite. Since g is not in $\mathrm{Comm}_G(H)$, we have a contradiction to our hypothesis that $e(G, K) = 1$ with $K = H \cap H^g$. This completes the proof. $\square$

We note another application of groups of surface type which provides an approach to the Algebraic Torus Theorem [**15**] similar to ours in [**43**]. We will omit a complete discussion of this approach, but will prove the following proposition to illustrate the ideas.

Proposition B.3.14. — *If (G, H) is of surface type and if H has infinite index in $\mathrm{Comm}_G(H)$, then there is a subgroup H' of finite index in H such that the normalizer $N(H')$ of H' is of finite index in G and $H' \backslash N(H')$ is virtually infinite cyclic. In particular, if H is virtually polycyclic, then G is virtually polycyclic.*

Proof. — Let X be a non-trivial H–almost invariant subset of G, let g be an element of $\mathrm{Comm}_G(H)$ and let $Y = gX$, so that Y has stabiliser H^g. Let H' denote the intersection $H \cap H^g$ which has finite index in both H and in H^g because g lies in $\mathrm{Comm}_G(H)$. Thus $H' \backslash X$ and $H' \backslash Y$ are both almost invariant subsets of $H' \backslash G$. As (G, H) is of surface type, the pair (G, H') has two ends so that $H' \backslash X$ and $H' \backslash Y$ are almost equal or almost complementary. It follows that X is H–almost equal to Y or Y^*, i.e. $gX \sim X$ or $gX \sim X^*$. Recall from Lemma B.2.10, that if $\mathcal{K}$ denotes $\{g \in G : gX \sim X \text{ or } gX \sim X^*\}$, then $\mathcal{K} \subset \mathrm{Comm}_G(H)$. It follows that in our present situation $\mathcal{K}$ must equal $\mathrm{Comm}_G(H)$. By Lemma B.2.12, we see that there are a finite number of subgroups $H_1, \ldots, H_m$ of finite index in H such that $\mathcal{K}$ is contained in the union of the normalizers $N(H_i)$. As H has infinite index in $\mathcal{K} = \mathrm{Comm}_G(H)$, one of the H_i, say H_1, has infinite index in its normalizer $N(H_1)$. As (G, H) is of surface type, the pair (G, H_1) has two ends, so we can apply Theorem 5.8 from [**47**] to the action of $H_1 \backslash N(H_1)$ on the left on the graph $H_1 \backslash \Gamma$. This result tells us that $H_1 \backslash N(H_1)$ is virtually infinite cyclic. Further the proof of this result in [**47**] shows that the quotient of $H_1 \backslash \Gamma$ by $H_1 \backslash N(H_1)$ must be finite so that $N(H_1)$ has finite index in G. $\square$

The arguments of [**43**] can be extended to show:

Theorem B.3.15. — *Let (G, H) be a pair of finitely generated groups with H virtually polycyclic and suppose that G does not split over a subgroup commensurable with a subgroup of infinite index in H. If for some subgroup K of H, $e(G, K) \geqslant 3$, then G splits over a subgroup commensurable with H.*

We end this section with an interpretation of intersection numbers in the case when the strong and ordinary intersection numbers are equal. This corrects a mistake in [**42**]. Suppose that a group G splits over subgroups H and K and let the corresponding H–almost and K–almost invariant subsets of G be X and Y. Let T denote the Bass-Serre tree corresponding to the splitting of G over K and consider the action of H on T. Let T' denote the minimal H-invariant subtree of T, and let Ψ denote the quotient graph $H\backslash T'$. Similarly, we get a graph Φ by considering the action of K on the Bass-Serre tree corresponding to the splitting of G over H. We have:

Theorem B.3.16. — *With the above notation, suppose that*

$$i(H\backslash X, K\backslash Y) = si(H\backslash X, K\backslash Y).$$

Then the number of edges in Ψ is the same as the number of edges in Φ and both are equal to $si(H\backslash X, K\backslash Y)$.

Proof. — The proof of Theorem 3.1 of [**42**] goes through because of our assumption that $i(H\backslash X, K\backslash Y) = si(H\backslash X, K\backslash Y)$. The mistake in [**42**] occurs in the proof of Lemma 3.6 of [**42**] where it is implicitly assumed that if X crosses Y, then it crosses Y strongly. Since we have assumed that the two intersection numbers are equal, the argument is now valid. □

BIBLIOGRAPHY

[1] S.A. Adeleke & P.M. Neumann – *Relations related to betweenness: their structure and automorphisms*, Mem. Amer. Math. Soc., vol. 623, American Mathematical Society, Providence, RI, 1998.

[2] H. Bass – Covering theory for graphs of groups, *J. Pure Appl. Algebra* **89** (1993), no. 1-2, p. 3–47.

[3] G.M. Bergman – On groups acting on locally finite graphs, *Ann. of Math.* **88** (1968), p. 335–340.

[4] M. Bestvina & M. Feighn – Bounding the complexity of simplicial actions on trees, *Invent. Math.* **103** (1991), p. 145–186.

[5] B.H. Bowditch – Cut points and canonical splittings of hyperbolic groups, *Acta Math.* **180** (1998), p. 145–186.

[6] ———, Connectedness properties of limit sets, *Trans. Amer. Math. Soc.* **351** (1999), p. 3673–3686.

[7] ———, *Treelike structures arising from continua and convergence groups*, Mem. Amer. Math. Soc., vol. 662, American Mathematical Society, Providence, RI, 1999.

[8] ———, Splittings of finitely generated groups over two-ended subgroups, *Trans. Amer. Math. Soc.* **354** (2002), no. 3, p. 1049–1078.

[9] D.E. Cohen – *Groups of cohomological dimension one*, Lect. Notes in Math., vol. 245, Springer-Verlag, Berlin, 1972.

[10] W. Dicks & M.J. Dunwoody – *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics, vol. 17, Cambridge University Press, 1989.

[11] M.J. Dunwoody – Accessibility and groups of cohomological dimension one, *Proc. London Math. Soc. (3)* **38** (1979), p. 193–215.

[12] ———, The accessibility of finitely presented groups, *Invent. Math.* **81** (1985), p. 449–457.

[13] M.J. Dunwoody & M. Roller – Splittings of groups over polycyclic by finite subgroups, *Bull. London Math. Soc.* **25** (1993), p. 29–36.

[14] M.J. Dunwoody & M.E. Sageev – JSJ-splittings for finitely presented groups over slender groups, *Invent. Math.* **135** (1999), p. 25–44.

[15] M.J. Dunwoody & E. Swenson – The Algebraic Torus Theorem, *Invent. Math.* **140** (2000), p. 605–637.

[16] M. Forester – On uniqueness of JSJ decompositions of finitely generated groups, *Comment. Math. Helv.*, to appear.

[17] ———, Deformation and rigidity of simplicial group actions on trees, *Geom. Topol.* **6** (2002), p. 219–267.

[18] M.H. Freedman, J. Hass & P. Scott – Closed geodesics on surfaces, *Bull. London Math. Soc.* **14** (1982), p. 385–391.

[19] ———, Least area incompressible surfaces in 3-manifolds, *Invent. Math.* **71** (1983), p. 609–642.

[20] K. Fujiwara & P. Papasoglu – JSJ-decompositions for finitely presented groups and complexes of groups, Preprint, 1997.

[21] R. Geoghegan – Topological Methods in Group Theory, book to appear.

[22] V. Guirardel – A very short proof of Forester's rigidity result, *Geom. Topol.* **7** (2003), p. 321–328.

[23] G. Higman – Subgroups of finitely presented groups, *Proc. Roy. Soc. London Ser. A* **262** (1961), p. 455–475.

[24] C.H. Houghton – Ends of locally compact groups and their quotient spaces, *J. Aust. Math. Soc.* **17** (1974), p. 274–284.

[25] W. Jaco & P. Shalen – *Seifert fibered spaces in* 3*-manifolds*, vol. 21, Mem. Amer. Math. Soc., no. 220, American Mathematical Society, Providence, RI, 1979.

[26] K. Johannson – *Homotopy equivalences of* 3*-manifolds with boundary*, Lect. Notes in Math., vol. 761, Springer-Verlag, 1979.

[27] P.H. Kropholler – An analogue of the torus decomposition theorem for certain Poincaré duality groups, *Proc. London Math. Soc. (3)* **60** (1990), p. 503–529.

[28] ———, A group-theoretic proof of the torus theorem, in *Geometric group theory, Vol. 1 (Sussex, 1991)*, London Math. Soc. Lecture Note Ser., vol. 181, Cambridge Univ. Press, Cambridge, 1993, p. 138–158.

[29] P.H. Kropholler & M.A. Roller – Splittings of Poincaré duality groups, *Math. Z.* **197** (1988), no. 3, p. 421–438.

[30] ______, Relative ends and duality groups, *J. Pure Appl. Algebra* **61** (1989), p. 197–210.

[31] ______, Splittings of Poincaré duality groups III, *J. London Math. Soc. (2)* **39** (1989), no. 2, p. 271–284.

[32] B. Leeb & P. Scott – The characteristic submanifold in higher dimensions, *Comment. Math. Helv.* **75** (2000), p. 201–215.

[33] G. Mess – Examples of Poincaré duality groups, *Proc. Amer. Math. Soc.* **110** (1990), no. 4, p. 1145–1146.

[34] W.D. Neumann & G.A. Swarup – Canonical decompositions of 3-manifolds, *Geom. Topol.* **1** (1997), p. 21–40.

[35] G. Niblo – The singularity obstruction for group splittings, *Topology Appl.* **119** (2002), no. 1, p. 17–31.

[36] E. Rips & Z. Sela – Cyclic splittings of finitely presented groups and canonical JSJ-decomposition, *Ann. of Math. (2)* **146** (1997), p. 53–109.

[37] J.H. Rubinstein & S. Wang – π_1-injective surfaces in graph manifolds, *Comment. Math. Helv.* **73** (1998), p. 499–515.

[38] M. Sageev – Codimension-1 subgroups and splittings of groups, *J. Algebra* **189** (1997), no. 2, p. 377–389.

[39] P. Scott – JSJ-decompositions of 3-manifolds, unpublished manuscript.

[40] ______, Ends of pairs of groups, *J. Pure Appl. Algebra* **11** (1977), p. 179–198.

[41] ______, A new proof of the Annulus and Torus Theorems, *Amer. J. Math.* **102** (1980), p. 241–277.

[42] ______, The Symmetry of Intersection Numbers in Group Theory, *Geom. Topol.* **2** (1998), p. 11–29, Correction (ibid), 1998.

[43] P. Scott & G.A. Swarup – An Algebraic Annulus Theorem, *Pacific J. Math.* **196** (2000), p. 461–506.

[44] ______, Splittings of groups and intersection numbers, *Geom. Topol.* **4** (2000), p. 179–218.

[45] ______, Canonical splittings of groups and 3-manifolds, *Trans. Amer. Math. Soc.* **353** (2001), p. 4973–5001.

[46] ______, Regular Neighbourhoods and canonical decompositions for groups, *Electron. Res. Announc. Amer. Math. Soc.* **8** (2002), p. 20–28.

[47] P. Scott & C.T.C. Wall – Topological methods in group theory, in *Homological Group Theory*, London Math. Soc. Lecture Notes Ser., vol. 36, Cambridge University Press, 1979, p. 137–214.

[48] Z. Sela – Acylindrical accessibility for groups, *Invent. Math.* **129** (1997), p. 527–565.

[49] ______, Structure and rigidity in (Gromov) hyperbolic groups and rank 1 Lie groups, *Geometric and Functional Anal.* **7** (1997), p. 561–593.

[50] J.-P. Serre – *Arbres, amalgames, SL_2*, Astérisque, vol. 46, Société Mathématique de France, Paris, 1977.

[51] ______, *Trees*, Springer-Verlag, Berlin-New York, 1980, Translated from the French by John Stillwell.

[52] J.R. Stallings – On torsion-free groups with infinitely many ends, *Ann. of Math.* **88** (1968), p. 312–334.

[53] ______, *Group theory and three-dimensional manifolds*, Yale Mathematical Monographs, vol. 4, Yale University Press, New Haven, Conn., 1971.

[54] E. Swenson – Axial pairs and convergence groups on S^1, *Topology* **39** (2000), p. 229–237.

[55] J. Tits – Sur le groupe des automorphismes d'un arbre, in *Essays on topology and related topics (Mémoires dédiés à Georges de Rham)* (A. Haefliger & R. Narasimhan, eds.), Springer-Verlag, New York-Berlin, 1970, p. 188–211.

[56] F. Waldhausen – On the determination of some bounded 3-manifolds by their fundamental groups alone, in *Proc. of the International sympo. on Topology and its applications*, Herceg-Novi, Yugoslavia, Beograd, 1969, p. 331–332.

INDEX