ISSN 1364-0380 **179**

Geometry & Topology
Volume 4 (2000) 179–218
Published: 9 August 2000

Splittings of groups and intersection numbers

Peter Scott
Gadde A Swarup

Mathematics Department, University of Michigan
Ann Arbor, Michigan 48109, USA
and
Mathematics Department, University of Melbourne
Parkville, Victoria 3052, Australia

Email: `pscott@math.lsa.umich.edu` and `gadde@ms.unimelb.edu.au`

Abstract

We prove algebraic analogues of the facts that a curve on a surface with self-intersection number zero is homotopic to a cover of a simple curve, and that two simple curves on a surface with intersection number zero can be isotoped to be disjoint.

AMS Classification numbers Primary: 20E06, 20E08

Secondary: 20F32, 57M07

Keywords: Amalgamated free product, splitting, intersection number, ends

Proposed: Jean-Pierre Otal
Seconded: Walter Neumann, Joan Birman

Received: 18 May 1999
Revised: 6 April 2000

In this paper, we will discuss an algebraic version of intersection numbers which was introduced by Scott in [14]. 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 [7], 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 [14], Scott extended the ideas of [7] 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 [14], 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 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 1. There is also an idea of intersection number of such things, which we give in Definition 1.3. This too was introduced by Scott in [14]. Our second main result, Theorem 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 2.

In a separate paper [17], we use the ideas about intersection numbers of splittings developed in [14] 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 [10] and [11]). 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 [18] 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 [14] provides some of the foundational material.

This paper is organised as follows. In section 1, we recall from [14] the basic definitions of intersection numbers in the algebraic context. We also prove a technical result which was essentially proved by Scott [13] 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 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 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 [8] 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 [9] on splittings of groups in general. We point out that these ideas lead to an alternative approach to the algebraic Torus Theorem [5]. We end the section with a brief discussion of an error in [14]. 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 [14] 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.

1 Preliminaries and statements of main results

We will start by recalling from [14] 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 [6], 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 [6] and [7]. 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 [6], 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 [2], 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, ie, 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 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 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 [14] 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 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 1.4 The following facts about the intersection number are proved in [14].

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

(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 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 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 [14], 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 [15] 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 $a_1 b_1 a_2 \ldots a_n b_n h$ 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 $a_1 t^{\epsilon_1} a_2 t^{\epsilon_2} ... a_n t^{\epsilon_n} a_{n+1}$ where a_{n+1} lies in A and, for $1 \leq i \leq n$, $\epsilon_i = 1$ 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 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 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 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 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, [19] or [20], 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 [3], 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 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 $e \leq f$, then $\overline{f} \leq \overline{e}$.*

(2) *If $e,\ f \in E$, there are only finitely many $g \in E$ such that $e \leq g \leq f$.*

(3) *If e, $f \in E$, at least one of the four relations $e \leq f$, $e \leq \overline{f}$, $\overline{e} \leq f$, $\overline{e} \leq \overline{f}$ holds.*

(4) *If e, $f \in E$, one cannot have $e \leq f$ and $e \leq \overline{f}$.*

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 $e \leq f$ 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 1.4. We give the proof here because the proof in [14] is not complete, and we will need to apply the methods of proof later in this paper.

Lemma 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 $\{g \in G : gX$ and Y 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 [15]. 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 [13], 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 2.5 and 2.8. The first part of the theorem directly corresponds to the result proved in [13], and the second part is a simple generalisation which will be needed later.

Theorem 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 $H_1, \ldots, H_k$ be finitely generated subgroups of a finitely generated group G. Let X_i, $1 \le i \le k$, be a non-trivial H_i–almost invariant set*

in G such that $E = \{gX_i, gX_i^* : 1 \leq i \leq k, g \in G\}$ 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 [13], 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 $U \leq V$. 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 1.12 is that if two of the four sets $U^{(*)} \cap V^{(*)}$ are small, then one of them is empty. Thus, as in [13], we define $U \leq V$ if and only if $U \cap V^*$ is empty or the only small set of the four. Note that if $U \subset V$ then $U \leq V$. 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 [13].

Lemma 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. $\square$

Now we can prove that our definition of $\leq$ yields a partial order on E. Our proof is essentially the same as in Lemma 2.4 of [13].

Lemma 1.14 *If a relation $\leq$ is defined on E by the condition that $U \leq V$ if and only if $U \cap V^*$ is empty or the only small set of the four sets $U^{(*)} \cap V^{(*)}$, then $\leq$ is a partial order.*

Proof We need to show that $\leq$ is transitive and that if $U \leq V$ and $V \leq U$ then $U = V$.

Suppose first that $U \leq V$ and $V \leq U$. 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 1.12 implies that one of these two sets must be empty. As $U \leq V$, our definition of $\leq$ implies that $U \cap V^*$ is empty. Similarly, the fact that $V \leq U$ 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 $U \leq V \leq W$. We must show that $U \leq W$.

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 $\leq$ now shows that $U \leq W$, 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 $U \leq V \leq W$, then either $U \leq W$ or $W \subset U$. Now we consider two cases.

First suppose that $U \subset V \leq W$, so that either $U \leq W$ or $W \subset U$. If $W \subset U$, then $W \subset V$, so that $W \leq V$. As $V \leq W$ and $W \leq V$, it follows from the first paragraph of the proof of this lemma that $V = W$. Hence, in either case, $U \leq W$.

Now consider the general situation when $U \leq V \leq W$. Again either $U \leq W$ or $W \subset U$. If $W \subset U$, then we have $W \subset U \leq V$. Now the preceding paragraph implies that $W \leq V$. Hence we again have $V \leq W$ and $W \leq V$ so that $V = W$. Hence $U \leq W$ still holds. This completes the proof of the lemma. □

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 1.10.

Lemma 1.15 *E together with $\leq$ satisfies the following conditions.*

(1) *If U, $V \in E$ and $U \leq V$, then $V^* \leq U^*$.*
(2) *If U, $V \in E$, there are only finitely many $Z \in E$ such that $U \leq Z \leq V$.*
(3) *If U, $V \in E$, at least one of the four relations $U \leq V$, $U \leq V^*$, $U^* \leq V$, $U^* \leq V^*$ holds.*
(4) *If U, $V \in E$, one cannot have $U \leq V$ and $U \leq V^*$.*

Proof Conditions (1) and (3) are obvious from the definition of $\leq$ and the hypotheses of Theorem 1.12.

To prove (4), we observe that if $U \leq V$ and $U \leq V^*$, 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 $Z \leq X$. 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 $Z \leq X$ but Z is not contained in X, then Z and X are not nested. Now Lemma 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 $Z \leq X$, 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 $Z \leq X$ then Z lies in the d_1–neighbourhood of X. Similarly, there is $d_2 > 0$ such that if $Z \leq X^*$, 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 $U \leq V$, the set $U \cap V^*$ lies in the d–neighbourhood of each of U, U^*, V and V^*.

Now suppose we are given $U \leq V$ 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 $U \leq Z \leq V$, 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 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 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 $\leq$ be defined on E by the condition that $U \leq V$ if $U \cap V^*$ is empty or the only small set of the four sets $U^{(*)} \cap V^{(*)}$. Lemmas 1.14 and 1.15 show that $\leq$ 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 $\{g : gX \leq X \text{ or } gX^* \leq X\}$, and note that condition 3) of Lemma 1.15 shows that $W^* = \{g : gX \leq X^* \text{ or } gX^* \leq X^*\}$. Recall that there is $d_1 > 0$ such that if $Z \leq X$ 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 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 1.12.

For the second part, we will simply comment on the modifications needed to the preceding proof. The statement of Lemma 1.13 remains true though the proof needs a little modification. The statement and proof of Lemma 1.14 apply unchanged. The statement of Lemma 1.15 remains true, though the proof needs some minor modifications. Finally the proof of the first part of Theorem 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 1.12. □

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 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 T has stabiliser H. Let T' denote the fixed set of H, ie, 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', ie $NT' = T'$. 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 $n \geq 1$. As G acts transitively on T, we have $T = T' \cup g_1T' \cup \ldots \cup g_nT'$. 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 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 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 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 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 $\delta \geq 0$ 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 $\delta \geq 0$. 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 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 [15] 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 [15] 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 $k^nX = X$, 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 $kY' = k^2Y = Y \subset kY^* = (Y')^*$. 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 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 2.3. □

Lemma 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 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 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 $H_1, ..., H_n$ with associated H_i–almost invariant subsets X_i of G, and let $E = \{gX_i, gX_i^* : g \in G, 1 \leq i \leq n\}$. 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 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 $1 \leq i \neq j \leq n$. 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 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 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 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, $1 \leq i \leq n$, 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', $1 \leq i \leq n$, 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 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 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 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 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 2.5. □

Now we turn to the proof of Theorem 2.8 that splittings exist. It will be convenient to make the following definitions. We will use H^g to denote gHg^{-1}.

Definition 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 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 $Comm_G(H)$ 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 2.8 *Let G be a finitely generated group with a finitely generated subgroup H, such that $e(G,H) \geq 2$. 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 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 1.12. As before, the assumption that $i(P,P) = 0$ implies that E is almost nested. However, in order to apply Theorem 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 2.5, we simply applied Lemma 2.4. However, here the situation is somewhat more complicated. Lemma 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 $Comm_G(H)$. Thus it would suffice to arrange that E is nested with respect to $\mathcal{K}$, ie, that gX and X are nested so long as g lies in $\mathcal{K}$. Now Proposition 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 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 2.8 apart from the proofs of Lemma 2.10 and Proposition 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 $e(G,H) \geq 2$, and let X denote a non-trivial H–almost invariant subset of G.

Recall that our problem in the proof of Theorem 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 $\mathcal{K} = \{g \in G : gX \sim X$ 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 2.8.

Lemma 2.10 *If $\mathcal{K} = \{g \in G : gX \sim X$ or $X^*\}$, then $H \subset \mathcal{K} \subset Comm_G(H)$.*

Proof The first inclusion is clear. The second is proved in essentially the same way as the proof of the first part of Lemma 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 $\mathcal{K} \subset Comm_G(H)$, as required. $\square$

Another way of describing our difficulty in applying Theorem 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 1.11 assures us that "most" of the translates are nested. The following result gives us a much stronger finiteness result.

Lemma 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 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 $1 \leq i \leq n$. 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. $\square$

Now we come to the key result.

Lemma 2.12 *Let G, H, X, $\mathcal{K}$ be as above. Then there are a finite number of finite index subgroups $H_1, ..., H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1 \leq i \leq m$, where $N(H_i)$ denotes the normaliser of H_i in G.*

Proof Consider an element g in $\mathcal{K}$. Lemma 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 2.11 that $\{g \in \mathcal{K} : gX \text{ and } X \text{ are not nested}\}$ 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 $\{h \in \mathcal{K} : h(gX) \text{ and } gX \text{ are not nested}\}$ 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 $H_1, ..., H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1 \leq i \leq m$, as required. □

Using this result, we can prove the following.

Lemma 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 [15] 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 [15] 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 $H_1,...,H_m$ of H, such that $\mathcal{K}$ is contained in the union of the groups $N(H_i)$, $1\leq i\leq m$. 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 2.8.

Proposition 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 $\{g_1 = e, g_2, ..., g_n\}$ 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 2.14. □

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 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 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 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 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 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 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 [18] that s is hyperbolic with respect to t.

Remark 3.2 shows that $si(H\backslash X, H\backslash X) \leq i(H\backslash X, H\backslash X)$. Recall that Theorem 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 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 $Comm_G(H) = Comm_G(K)$. However, Y may be very different from X. This leads to some splitting results when we place further restrictions on H.

Proposition 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 $g\delta Y \cap Y$ or $g\delta Y \cap Y^*$ 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, $1 \leq i \leq n$. 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 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 1.12. □

Definition 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 [16]. 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 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 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 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 [8], 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 3.5, the above Corollary and then Theorem 2.8 give an alternative proof of their splitting result. Thus Theorem 2.8 may be considered as a generalization of their splitting theorem. We next reformulate in our language a conjecture of Kropholler and Roller [9]:

Conjecture 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 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 [4], and Sageev [12] proved it for quasiconvex subgroups of hyperbolic groups. The paper of Dunwoody and Roller [4] contains information useful in the general case. The second step in their proof, which uses a theorem of Bergman [1], proves the following result, stated in our language. (There is an exposition of Bergman's argument and parts of [4] in the later versions of [5].)

Theorem 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 $Comm_G(H) = Comm_G(H')$.*

This combined with Proposition 2.14 gives:

Corollary 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 $Comm_G(H)$ and is nested with respect to the subgroup $\mathcal{K} = \{g \in G : gX \sim X \text{ or } X^*\}$ of $Comm_G(H)$.*

Now Theorem 1.12 yields the following generalization of Stallings' Theorem [21] already noted by Dunwoody and Roller [4]:

Theorem 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 3.11 leads to the following partial solution of the above conjecture of Kropholler and Roller:

Theorem 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 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 $Comm_G(H)$ 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 $Comm_G(H)$ 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 $Comm_G(H)$, we have a contradiction to our hypothesis that $e(G,K) = 1$ with $K = H \cap H^g$. This completes the proof. □

We note another application of groups of surface type which provides an approach to the Algebraic Torus Theorem [5] similar to ours in [16]. We will omit a complete discussion of this approach, but will prove the following proposition to illustrate the ideas.

Proposition 3.14 *If (G,H) is of surface type and if H has infinite index in $Comm_G(H)$, 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 $Comm_G(H)$ 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 $Comm_G(H)$. 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^*, ie, $gX \sim X$ or $gX \sim X^*$. Recall from Lemma 2.10, that if $\mathcal{K}$ denotes $\{g \in G : gX \sim X \text{ or } gX \sim X^*\}$, then $\mathcal{K} \subset Comm_G(H)$. It follows that in our present situation $\mathcal{K}$ must equal $Comm_G(H)$. By Lemma 2.12, we see that there are a finite number of subgroups $H_1, ..., H_m$ 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 $\mathcal{K} = Comm_G(H)$, 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 [15] 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 [15] 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. □

The arguments of [16] can be extended to show:

Theorem 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, $e(G,K)\geq 3$, 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 [14]. 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 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 [14] goes through because of our assumption that $i(H\backslash X,K\backslash Y)=si(H\backslash X,K\backslash Y)$. The mistake in [14] occurs in the proof of Lemma 3.6 of [14] 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. □

Acknowledgement The first author is partially supported by NSF grants DMS 034681 and 9626537.

References

[1] **G M Bergman**, *On groups acting on locally finite graphs*, Annals of Mathematics, 88 (1968) 335–340

[2] **D E Cohen**, *Groups of cohomological dimension one*, Lecture Notes in Math. 245, Springer–Verlag, Berlin (1972)

[3] **M J Dunwoody**, *Accessibility and groups of cohomological dimension one*, Proc. London Math. Soc. 38 (1979) 193–215

[4] **M J Dunwoody**, **M A Roller**, *Splittings of groups over polycyclic-by-finite groups*, Bull. London Math. Soc. 25 (1993) 29–36

[5] **M J Dunwoody**, **E Swenson**, *The Algebraic Torus Theorem*, Preprints, 1997, 1998, 1999

[6] **M H Freedman**, **J Hass**, **P Scott**, *Closed geodesics on surfaces*, Bull. London Math. Soc. 14 (1982) 385–391

[7] **M Freedman**, **J Hass**, **P Scott**, *Least area incompressible surfaces in 3–manifolds*, Invent. Math. 71 (1983) 609–642

[8] **P H Kropholler**, **M A Roller**, *Splittings of Poincaré Duality groups*, Math. Zeit. 197 (1988) 421–438

[9] **P H Kropholler**, **M A Roller**, *Relative ends and duality groups*, J. Pure Appl. Algebra, 61 (1989) 197–210

[10] **B Leeb**, **P Scott**, *A Geometric Characteristic Splitting in all Dimensions*, to appear in Comm. Math. Helv.

[11] **W D Neumann**, **G A Swarup**, *Canonical Decompositions of 3–manifolds*, Geometry and Topology, 1 (1997) 21–40

[12] **M Sageev**, *Codimension–1 subgroups and splittings of groups, J. Algebra 189 (1997) 377–389*

[13] **P Scott**, *A New Proof of the Annulus and Torus Theorems*, American J. of Math. 102(1980) 241–277

[14] **P Scott**, *The Symmetry of Intersection Numbers in Group Theory*, Geometry and Topology 2 (1998) 11–29, Correction (ibid) (1998) 333–335

[15] **P Scott**, **C T C Wall**, *Topological methods in group theory*, Homological Group Theory, London Math. Soc. Lecture Notes Series, 36 (1979) 137– 214

[16] **P Scott**, **G A Swarup**, *An Algebraic Annulus Theorem*, to appear in the Pacific Journal of Mathematics

[17] **P Scott**, **G A Swarup**, *Canonical splittings of groups and 3–manifolds*, Preprint (2000)

[18] **Z Sela**, *Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups II*, Geom. Funct. Anal. 7 (1997) 561–594

[19] **J-P Serre**, *Arbres, amalgames, SL_2*, Astérisque No. 46, Société Mathématique de France, Paris (1977)

[20] **J-P Serre**, *Trees*, translated from French by John Stillwell, Springer–Verlag, Berlin–New York (1980) ISBN 3-540-10103-9

[21] **J R Stallings**, *Group theory and 3–dimensional manifolds*, Yale University Press, New Haven (1971)