ISSN 1364-0380

11

Geometry & Topology
Volume 2 (1998) 11–29
Published: 19 March 1998

The Symmetry of Intersection Numbers in Group Theory

Peter Scott

Mathematics Department
University of Michigan
Ann Arbor, MI 48109, USA

Email: `pscott@math.lsa.umich.edu`

Abstract

For suitable subgroups of a finitely generated group, we define the intersection number of one subgroup with another subgroup and show that this number is symmetric. We also give an interpretation of this number.

AMS Classification numbers Primary: 20F32

Secondary: 20E06, 20E07, 20E08, 57M07

Keywords: Ends, amalgamated free products, trees

Proposed: Jean-Pierre Otal
Seconded: Cameron Gordon, Walter Neumann

Received: 21 February 1997
Revised: 13 March 1998

If one considers two simple closed curves L and S on a closed orientable surface F, one can define their intersection number to be the least number of intersection points obtainable by isotoping L and S transverse to each other. (Note that the count is to be made without any signs attached to the intersection points.) By definition, this number is symmetric, ie the roles of L and S are interchangeable. This can be regarded as a definition of the intersection number of the two infinite cyclic subgroups Λ and Σ of the fundamental group of F which are carried by L and S. In this paper, we show that an analogous definition of intersection number of subgroups of a group can be given in much greater generality and proved to be symmetric. We also give an interpretation of these intersection numbers.

In [7], Rips and Sela considered a torsion free finitely presented group G and infinite cyclic subgroups Λ and Σ such that G splits over each. (A group G splits over a subgroup C if either G has a HNN decomposition $G = A*_C$, or G has an amalgamated free product structure $G = A *_C B$, where $A \neq C \neq B$.) They effectively considered the intersection number $i(\Lambda, \Sigma)$ of Λ with Σ, and they proved that $i(\Lambda, \Sigma) = 0$ if and only if $i(\Sigma, \Lambda) = 0$. Using this, they proved that G has what they call a JSJ decomposition. If $i(\Lambda, \Sigma)$ was not zero, it follows from their work that G can be expressed as the fundamental group of a graph of groups with some vertex group being a surface group H which contains Λ and Σ. Now it is intuitively clear (and we discuss it further at the end of section 2 of this paper) that the intersection number of Λ with Σ is the same whether it is measured in G or in H. Also the intersection numbers of Λ and Σ in H are symmetric because of their topological interpretation. So it follows at the end of all their work that the intersection numbers of Λ and Σ in G are also symmetric. In 1994, Rips asked if there was a simpler proof of this symmetry which does not depend on their proof of the JSJ splitting. The answer is positive, and the ideas needed for the proof are all essentially contained in earlier papers of the author. This paper is a belated response to Rips' question. The main idea is to reduce the natural, but not clearly symmetric, definition of intersection number to counting the intersections of suitably chosen sets. The most general possible algebraic situation in which to define intersection numbers seems to be that of a finitely generated group G and two finitely generated subgroups Λ and Σ, not necessarily cyclic, such that the number of ends of each of the pairs (G, Λ) and (G, Σ) is more than one. Note that any infinite cyclic subgroup Λ of $\pi_1(F)$ satisfies $e(\pi_1(F), \Lambda) = 2$. This is because F is closed and orientable so that the cover of F with fundamental group Λ is an open annulus which has two ends. In order to handle the general situation, we will need the concept of an almost invariant set, which is closely related to the theory of ends. We should note that Kropholler and Roller [6] introduced an intersection cohomology class in the special case of $PD(n-1)$–subgroups of PDn–groups. Their ideas are closely related to ours, and we will discuss the connections at the start of the last section of this paper. Finally, we should point out that since Rips asked the above question about symmetry of intersection numbers, Dunwoody and Sageev [2] have given a proof of the existence of a JSJ decomposition for any finitely presented group which is very much simpler and more elementary than that of Rips and Sela.

The preceding discussion is a little misleading, as the intersection numbers which we define are not determined simply by a choice of subgroups. In fact, we define intersection numbers for almost invariant sets. A special case occurs when one has a group G and subgroups Λ and Σ such that G splits over each, as a splitting of G has a well defined almost invariant set associated. This is discussed in section 2. Thus we can define the intersection number of two splittings of G. In the case of cyclic subgroups of surface groups corresponding to simple closed curves, these curves determine splittings of the surface group over each cyclic subgroup, and the intersection number we define for these splittings is the same as the topological intersection number of the curves.

In the first section of this paper, we discuss in more detail intersection numbers of closed curves on surfaces. In the second section we introduce the concept of an almost invariant set and prove the symmetry results advertised in the title. In the third section, we discuss the interpretation of intersection numbers when they are defined, and how our ideas are connected with those of Kropholler and Roller.

Acknowledgments This paper was written while the author was visiting the Mathematical Sciences Research Institute in Berkeley in 1996/7. Research at MSRI is supported in part by NSF grant DMS-9022140. He is also grateful for the partial support provided by NSF grants DMS-9306240 and DMS-9626537.

1 The symmetry for surface groups

In this section, we will discuss further the special case of two essential closed curves L and S on a compact surface F. This will serve to motivate the definitions in the following section, and also show that the results of that section do indeed answer the question of Rips. It is not necessary to assume that F is closed or orientable, but we do need to assume that L and S are two-sided on F. As described in the introduction in the case of simple curves, one defines their intersection number to be the least number of intersection points obtainable by homotoping L and S transverse to each other, where the count is to be made without any signs attached to the intersection points. (One should also insist that L and S be in general position, in order to make the count correctly.) Of course, this number is symmetric, ie the roles of L and S are interchangeable. We will show in section 2 that one can define these intersection numbers in an algebraically natural way. There is also an idea of self-intersection number for a curve on a surface and we will discuss a corresponding algebraic idea.

For the next discussion, we will restrict our attention to the case when L and S are simple and introduce the algebraic approach to defining intersection numbers taken by Rips and Sela in [7]. Let G denote $\pi_1(F)$. Suppose that L and S cannot be made disjoint and choose a basepoint on $L \cap S$. Suppose that L represents the element λ of G. This element λ cannot be trivial, nor can L be parallel to a boundary component of F, because of our assumption that L and S cannot be made disjoint. Thus L induces a splitting of G over the infinite cyclic subgroup Λ of G which is generated by λ. Let σ denote the element of G represented by S. Define $d(\sigma, \lambda)$ to be the length of σ when written as a word in cyclically reduced form in the splitting of G determined by L. Similarly, define $d(\lambda, \sigma)$ to be the length of λ when written as a word in cyclically reduced form in the splitting of G determined by S. For convenience, suppose also that L and S are separating. Then each of these numbers is equal to the intersection number of L and S described above and therefore $d(\lambda, \sigma) = d(\sigma, \lambda)$. What is interesting is that this symmetry is not obvious from the purely algebraic point of view, but it is obvious topologically because the intersection of two sets is symmetric.

In the above discussion, we restricted attention to simple closed curves on a surface F, because the algebraic analogue is clear. If F is closed, then not only does a simple closed curve on F determine a splitting of $\pi_1(F)$ over the infinite cyclic subgroup carried by the curve, but any splitting of $\pi_1(F)$ over an infinite cyclic subgroup is induced in this way by some simple closed curve on F. Hence the algebraic situation described above exactly corresponds to the topological situation when F is closed.

Now we continue with further discussion of the intersection number of two closed curves L and S which need not be simple. As in [3], it will be convenient to assume that L and S are shortest closed geodesics in some Riemannian metric on F so that they automatically intersect minimally. Instead of defining the intersection number of L and S in the "obvious" way, we will interpret our intersection numbers in suitable covers of F, exactly as in [3] and [4]. Let F_Λ denote the cover of F with fundamental group equal to Λ. Then L lifts to F_Λ

and we denote its lift by L again. Let l denote the pre-image of this lift in the universal cover $\widetilde{F}$ of F. The full pre-image of L in $\widetilde{F}$ consists of disjoint lines which we call L–lines, which are all translates of l by the action of G. Similarly, we define F_Σ, the line s and S–lines in $\widetilde{F}$. Now we consider the images of the L–lines in F_Σ. Each L–line has image in F_Σ which is a possibly singular line or circle. Then we define $d(L,S)$ to be the number of images of L–lines in F_Σ which meet S. Similarly, we define $d(S,L)$ to be the number of images of S–lines in F_Λ which meet L. It is shown in [3], using the assumption that L and S are shortest closed geodesics, that each L–line in F_Σ crosses S at most once, and similarly for S–lines in F_Λ. It follows that $d(L,S)$ and $d(S,L)$ are each equal to the number of points of $L \cap S$, and so they are equal to each other. (This assumes that L and S are in general position.)

Here is an argument which shows that $d(L,S)$ and $d(S,L)$ are equal without reference to the situation in the surface F. Recall that the L–lines are translates of l by elements of G. Of course, there is not a unique element of G which sends l to a given L–line. In fact, the L–lines are in natural bijective correspondence with the cosets $g\Lambda$ of Λ in G. (Our groups act on the left on covering spaces.) The images of the L–lines in F_Σ are in natural bijective correspondence with the double cosets $\Sigma g\Lambda$, and $d(L,S)$ counts the number of these double cosets such that the line gl crosses s. Similarly, $d(S,L)$ counts the number of the double cosets $\Lambda h\Sigma$ such that the line hs crosses l. Note that it is trivial that gl crosses s if and only if l crosses $g^{-1}s$. Now we use the bijection from G to itself given by sending each element to its inverse. This induces a bijection between the set of all double cosets $\Sigma g\Lambda$ and the set of all double cosets $\Lambda h\Sigma$ by sending $\Sigma g\Lambda$ to $\Lambda g^{-1}\Sigma$. It follows that it also induces a bijection between those double cosets $\Sigma g\Lambda$ such that gl crosses s and those double cosets $\Lambda h\Sigma$ such that hs crosses l, which shows that $d(L,S)$ equals $d(S,L)$ as required.

This argument has more point when one applies it to a more complicated situation than that of curves on surfaces. In [4], we considered least area maps of surfaces into a 3–manifold. The intersection number which we used there was defined in essentially the same way but it had no obvious topological interpretation such as the number of double curves of intersection. We proved that our intersection numbers were symmetric by the above double coset argument, in [4] just before Theorem 6.3.

2 Intersection Numbers in General

In order to handle the general case, we will need the idea of an almost invariant set. This idea was introduced by Cohen in [1] and was first used in the relative context by Houghton in [5]. We will introduce this idea and explain its connection with the foregoing.

Let E and F be sets. We say that E and F are almost equal, and write $E \overset{a}{=} F$, if the symmetric difference $(E - F) \cup (F - E)$ is finite. If E is contained in some set W on which a group G acts on the right, we say that E is almost invariant if $Eg \overset{a}{=} E$, for all g in G. An almost invariant subset E of W will be called non-trivial if it is infinite and has infinite complement. The connection of this idea with the theory of ends of groups is via the Cayley graph Γ of G with respect to some finite generating set of G. (Note that in this paper groups act on the left on covering spaces and, in particular, G acts on its Cayley graph on the left.) Using $\mathbb{Z}_2$ as coefficients, we can identify 0–cochains and 1–cochains on Γ with sets of vertices or edges. A subset E of G represents a set of vertices of Γ which we also denote by E, and it is a beautiful fact, due to Cohen [1], that E is an almost invariant subset of G if and only if δE is finite, where δ is the coboundary operator. If H is a subgroup of G, we let $H\backslash G$ denote the set of cosets Hg of H in G, 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 have the idea of an almost invariant subset of $H\backslash G$.

Now we again consider the situation of simple closed curves L and S on a compact surface F and let $\widetilde{F}$ denote the universal cover of F. Pick a generating set for G which can be represented by a bouquet of circles embedded in F. We will assume that the wedge point of the bouquet does not lie on L or S. The pre-image of this bouquet in $\widetilde{F}$ will be a copy of the Cayley graph Γ of G with respect to the chosen generating set. The pre-image in F_Λ of the bouquet will be a copy of the graph $\Lambda\backslash\Gamma$, the quotient of Γ by the action of Λ on the left. Consider the closed curve L on F_Λ. Let D denote the set of all vertices of $\Lambda\backslash\Gamma$ which lie on one side of L. Then D has finite coboundary, as δD equals exactly the edges of $\Lambda\backslash\Gamma$ which cross L. Hence D is an almost invariant subset of $\Lambda\backslash G$. Let X denote the pre-image of D in Γ, so that X equals the set of vertices of Γ which lie on one side of the line l. There is an algebraic description of X in terms of canonical forms for elements of G as follows. Suppose that L separates F, so that $G = A *_\Lambda B$. Also suppose that L and D are chosen so that all the vertices of Γ labelled with an element of Λ do not lie in X. Pick right transversals T and T' for Λ in A and B respectively, both of which contain the identity e of G. (A right transversal of Λ in A consists of a choice of coset representative for each coset $a\Lambda$.) Each element of G can be expressed uniquely in the form $a_1b_1\ldots a_nb_n\lambda$, where $n \geq 1$, λ lies in Λ, each a_i lies in $T-\{e\}$ except that a_1 may be trivial, and each b_i lies in $T'-\{e\}$ except that b_n may be trivial. Then X consists of those elements for which a_1 is non-trivial. If Λ is non-separating in F, there is a similar description for X. See Theorem 1.7 of [11] for details. Similarly, we can define a set E in F_Σ and its pre-image Y in $\widetilde{F}$ which equals the set of vertices of Γ which lie on one side of the line s. Now finally the connection between the earlier arguments and almost invariant sets can be given. For we can decide whether the lines l and s cross by considering instead the sets X and Y. The lines l and s together divide G into the four sets $X\cap Y, X^*\cap Y, X\cap Y^*$ and $X^*\cap Y^*$, where X * denotes $G-X$, and l crosses s if and only if each of these four sets projects to an infinite subset of $\Sigma\backslash G$. Equally, s crosses l if and only if each of these four sets projects to an infinite subset of $\Lambda\backslash G$. As we know that l crosses s if and only if s crosses l, it follows that these conditions are equivalent. We will show that this symmetry holds in a far more general context.

Note that in the preceding example the subset X of G is Λ–invariant under the left action of Λ on G, ie $\lambda X = X$, for all λ in Λ.

For the most general version of this symmetry result, we can consider any finitely generated group G. Note that the subgroups of G which we consider need not be finitely generated.

Definition 2.1 If G is a finitely generated group and H is a subgroup, then a subset X of G is *H–almost invariant* if X is invariant under the left action of H, and simultaneously the quotient set $H\backslash X$ is almost invariant under the right action of G. In addition, X is a *non-trivial* H–almost invariant subset of G if $H\backslash X$ and $H\backslash X^*$ are both infinite.

Note that if X is a non-trivial H–almost invariant subset of G, then $e(G,H)$ is at least 2, as $H\backslash X$ is a non-trivial almost invariant subset of $H\backslash G$.

Definition 2.2 Let X be a Λ–almost invariant subset of G and let Y be a Σ–almost invariant subset of G. We will say that X *crosses* Y if each of the four sets $X\cap Y, X^*\cap Y, X\cap Y^*$ and $X^*\cap Y^*$ projects to an infinite subset of $\Sigma\backslash G$.

Note that it is obvious that if Y is trivial, then X cannot cross Y. Our first and most basic symmetry result is the following. This is essentially proved in Lemma 2.3 of [9], but the context there is less general.

Lemma 2.3 *If G is a finitely generated group with subgroups Λ and Σ, and X is a non-trivial Λ–almost invariant subset of G and Y is a non-trivial Σ–almost invariant subset of G, then X crosses Y if and only if Y crosses X.*

Remark 2.4 If X and Y are both trivial, then neither can cross the other, so the above symmetry result is clear. However, this symmetry result fails if only one of X or Y is trivial. Here is a simple example. Let Λ and Σ denote infinite cyclic groups with generators λ and σ respectively, and let G denote the group $\Lambda \times \Sigma$. We identify G with the set of integer points in the plane. Let $X = \{(m,n) \in G : n > 0\}$, and let $Y = \{(m,n) \in G : m = 0\}$. Then X is a non-trivial Λ–almost invariant subset of G and Y is a trivial Σ–almost invariant subset of G. One can easily check that Y crosses X, although X cannot cross Y as Y is trivial.

Proof Suppose that X does not cross Y. By replacing one or both of X and Y by its complement if needed, we can assume that $X \cap Y$ projects to a finite subset of $\Sigma\backslash G$. The fact that Y is non-trivial implies that $\Sigma\backslash Y$ is an infinite subset of $\Sigma\backslash G$, so there is a point z in $\Sigma\backslash Y$ which is not in the image of $X \cap Y$. Now we need to use some choice of generators for G and consider the corresponding Cayley graph Γ of G. The vertices of Γ are identified with G and the action of G on itself on the left extends to an action on Γ. We consider z and the image of $X \cap Y$ in the quotient graph $\Sigma\backslash\Gamma$. As $X \cap Y$ has finite image, there is a number d such that each point of its image can be joined to z by a path of length at most d. As the projection of Γ to $\Sigma\backslash\Gamma$ is a covering map, it follows that each point of $X \cap Y$ can be joined to some point lying above z by a path of length at most d. As any point above z lies in X^*, it follows that each point of $X \cap Y$ can be joined to some point of X^* by a path of length at most d. Hence each point of $X \cap Y$ lies at most distance d from δX. Thus the image of $X \cap Y$ in $\Lambda\backslash\Gamma$ lies within the d–neighbourhood of the compact set $\delta(\Lambda\backslash X)$, and so must itself be finite. It follows that Y does not cross X, which completes the proof of the symmetry result. □

At the start of this section, we explained how to connect the topological intersection of simple closed curves on a surface with crossing of sets. One can construct many other interesting examples in much the same way.

Example 2.5 As before, let F denote a closed surface with fundamental group G, and let $\widetilde{F}$ denote the universal cover of F. Pick a generating set of G which can be represented by a bouquet of circles embedded in F, so that $\widetilde{F}$ contains a copy of the Cayley graph Γ of G with respect to the chosen generators. Let

F_1 denote a cover of F which is homeomorphic to a four punctured torus and let Λ denote its fundamental group. For example, if F is the closed orientable surface of genus three, we can consider a compact subsurface F' of F which is homeomorphic to a torus with four open discs removed, and take the cover F_1 of F such that $\pi_1(F_1) = \pi_1(F')$. For notational convenience, we identify F_1 with $S^1 \times S^1$ with the four points $(1,1),(1,i),(1,-1)$ and $(1,-i)$ removed. Now we choose 1–dimensional submanifolds of F_1 each consisting of two circles and each separating F_1 into two pieces. Let L denote $S^1 \times \{e^{\pi i/4}, e^{5\pi i/4}\}$ and let S denote $S^1 \times \{e^{3\pi i/4}, e^{7\pi i/4}\}$. As before, we let D denote all the vertices of the graph $\Lambda\backslash\Gamma$ in F_1 which lie on one side of L, and let E denote all the vertices of the graph $\Lambda\backslash\Gamma$ in F_1 which lie on one side of S. Let X and Y denote the pre-images of D and E in G. Now D is an almost invariant subset of $\Lambda\backslash G$, as δD equals exactly the edges of $\Lambda\backslash\Gamma$ which cross L, and E is almost invariant for similar reasons. Hence X and Y are each Λ–almost invariant subsets of G. Clearly X and Y cross. An important feature of this example is that although X and Y cross, the boundaries L and S of the corresponding surfaces in F_1 are disjoint. This is quite different from the example with which we introduced almost invariant sets, but this is a much more typical situation.

Definition 2.6 Let Λ and Σ be subgroups of a finitely generated group G. Let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$ and let X and Y denote the pre-images in G of D and E respectively. We define $i(D,E)$ to equal the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y.

For this definition to be interesting, we need to show that $i(D,E)$ is finite, which is not obvious from the definition in this general situation. In fact, it may well be false if one does not assume that the groups Λ and Σ are finitely generated, although we have no examples. From now on, we will assume that Λ and Σ are finitely generated.

Lemma 2.7 *Let Λ and Σ be finitely generated subgroups of a finitely generated group G. Let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. Then $i(D,E)$ is finite.*

Proof This is again proved by using the Cayley graph, so it appears to depend on the fact that G is finitely generated. However, we have no examples where $i(D,E)$ is not finite when G is not finitely generated. The proof we give is essentially contained in that of Lemmas 4.3 and 4.4 of [8]. Start by considering the finite graph δD in $\Lambda\backslash\Gamma$. As Λ is finitely generated, we can add edges and vertices to δD to obtain a finite connected subgraph $\delta_1 D$ of $\Lambda\backslash\Gamma$ which contains δD and has the property that its inclusion in $\Lambda\backslash\Gamma$ induces a surjection of its fundamental group to Λ. Thus the pre-image of $\delta_1 D$ in Γ is a connected graph which we denote by $\delta_1 X$. Similarly, we obtain a finite connected graph $\delta_1 E$ of $\Sigma\backslash\Gamma$ which contains δE and has connected pre-image $\delta_1 Y$ in Γ. As usual, we will denote the pre-images of D and E in G by X and Y respectively.

Next we claim that if gX crosses Y then $g(\delta_1 X)$ intersects $\delta_1 Y$. (The converse need not be true.) Suppose that $g(\delta_1 X)$ and $\delta_1 Y$ are disjoint. Then $g(\delta_1 X)$ cannot meet δY. As $g(\delta_1 X)$ is connected, it must lie in Y or Y^*. It follows that $g(\delta X)$ lies in Y or Y^*, so that one of the four sets $X\cap Y, X^*\cap Y, X\cap Y^*$ and $X^*\cap Y^*$ must be empty, which implies that gX does not cross Y.

Now we can show that $i(D,E)$ must be finite. Recall that $i(D,E)$ is defined to be the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y. The preceding paragraph implies that $i(D,E)$ is bounded above by the number of double cosets $\Sigma g\Lambda$ such that $g(\delta_1 X)$ meets $\delta_1 Y$. Let P and Q be finite subgraphs of $\delta_1 X$ and $\delta_1 Y$ which project onto $\delta_1 D$ and $\delta_1 E$ respectively. If $g(\delta_1 X)$ meets $\delta_1 Y$, then there exist elements λ of Λ and σ of Σ such that $g(\lambda P)$ meets σQ. Thus $\sigma^{-1}g\lambda P$ meets Q. Now there are only finitely many elements of G which can translate P to meet Q, and it follows that $i(D,E)$ is bounded above by this number. $\square$

We have just shown that, as in the preceding section, the intersection numbers we have defined are symmetric, but we will need a little more information.

Lemma 2.8 *Let G be a finitely generated group with subgroups Λ and Σ, let D denote a non-trivial almost invariant subset of $\Lambda\backslash G$, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. Then the following statements hold:*

1) $i(D,E) = i(E,D)$,
2) $i(D,E) = i(D^*,E) = i(D,E^*) = i(D^*,E^*)$,
3) *if D' is almost equal to D and E' is almost equal to E, and X, X' and Y, Y' denote their pre-images in G, then X crosses Y if and only if X' crosses Y', so that $i(D,E) = i(D',E')$.*

Proof The first part is proved by using the bijection from G to itself given by sending each element to its inverse. This induces a bijection between all double cosets $\Sigma g\Lambda$ and $\Lambda h\Sigma$ by sending $\Sigma g\Lambda$ to $\Lambda g^{-1}\Sigma$, and it further induces a bijection between those double cosets $\Sigma g\Lambda$ such that gX crosses Y and those double cosets $\Lambda h\Sigma$ such that hY crosses X.

The second part is clear from the definitions.

For the third part, we note that, as E and E' are almost equal, so are their complements in $\Sigma\backslash G$, and it follows that X crosses Y if and only if it crosses Y'. Hence the symmetry proved in Lemma 2.3, shows that Y crosses X if and only Y' crosses X. Now the same argument reversing the roles of D and E yields the required result. $\square$

At this point, we have defined in a natural way a number which can reasonably be called the intersection number of D and E, but have not yet defined an intersection number for subgroups of G. First note that if $e(G,\Lambda)$ is equal to 2, then all choices of non-trivial almost invariant sets in $\Lambda\backslash G$ are almost equal or almost complementary. Let D denote some choice here. Suppose that $e(G,\Sigma)$ is also equal to 2, and let E denote a non-trivial almost invariant subset of $\Sigma\backslash G$. The third part of the preceding lemma implies that $i(D,E)$ is independent of the choices of D and E and so depends only on the subgroups Λ and Σ. This is then the definition of the intersection number $i(\Lambda,\Sigma)$. In the special case when G is the fundamental group of a closed orientable surface and Λ and Σ are cyclic subgroups of G, it is automatic that $e(G,\Lambda)$ and $e(G,\Sigma)$ are each equal to 2. The discussion of the previous section clearly shows that this definition coincides with the topological definition of intersection number of loops representing generators of these subgroups, whether or not those loops are simple. Note that one can also define the self-intersection number of an almost invariant subset D of $\Lambda\backslash G$ to be $i(D,D)$, and hence can define the self-intersection number of a subgroup Λ of G such that $e(G,\Lambda)=2$. Again this idea generalises the topological idea of self-intersection number of a loop on a surface.

If one considers subgroups Λ and Σ such that $e(G,\Lambda)$ or $e(G,\Sigma)$ is greater than 2, there are possibly different ideas for their intersection number depending on which almost invariant sets we pick. (It is tempting to simply define $i(\Lambda,\Sigma)$ to be the minimum possible value for $i(D,E)$, where D is a non-trivial Λ–almost invariant subset of G and E is a non-trivial Σ–almost invariant subset of G. But this does not seem to be the "right" definition.) However, there is a natural way to choose these almost invariant sets if we are given splittings of G over Λ and Σ. As discussed in the previous section in the case of surface groups, the standard way to do this when $G = A *_\Lambda B$ is in terms of canonical forms for elements of G as follows. Pick right transversals T and T' for Λ in A and B respectively, both of which contain the identity e of G. Then each element can be expressed uniquely in the form $a_1b_1\ldots a_nb_n\lambda$, where $n \geq 1$, λ lies in Λ, each a_i lies in $T-\{e\}$ except that a_1 may be trivial, and each b_i lies in $T'-\{e\}$ except that b_n may be trivial. Let X denote the subset of G consisting of elements for which a_1 is non-trivial, and let D denote $\Lambda\backslash X$. It is easy to check directly that X is Λ–almost invariant. One must check that $\lambda X = X$, for all λ in Λ and that $Dg \overset{a}{=} D$, for all g in G. The first equation is trivial, and the second is easily checked when g lies in A or B, which implies that it holds for all g in G. Note also that the definition of X is independent of the choices of transversals of Λ in A and B. Then D is the almost invariant set determined by the given splitting of G. This definition seems asymmetric, but if instead we consider the Λ–almost invariant subset of G consisting of elements whose canonical form begins with a non-trivial element of B, we will obtain an almost invariant subset of $\Lambda\backslash G$ which is almost equal to the complement of D. There is a similar description of D when $G = A*_\Lambda$. For details see Theorem 1.7 of [11]. The connection between D and the given splitting of G can be seen in several ways. From the topologists' point of view, one sees this as described earlier for surface groups. From the point of view of groups acting on trees, there is also a very natural description. One identifies a splitting of G with an action of G on a tree T without inversions, such that the quotient $G\backslash T$ has a single edge. Let e denote the edge of T with stabiliser Λ, let v denote the vertex of e with stabiliser A, and let E denote the component of $T-\{e\}$ which contains v. Then we can define $X = \{g \in G : ge \subset E\}$. It is easy to check directly that this set is the same as the set X defined above using canonical forms.

In the preceding paragraph, we showed how to obtain a well defined intersection number of given splittings over Λ and Σ. An important point to notice is that this intersection number is not determined by the subgroups Λ and Σ of G only. It depends on the given splittings. In the case when G is a surface group, this is irrelevant as there can be at most one splitting of a surface group over a given infinite cyclic subgroup. But in general, a group G with subgroup Λ can have many different splittings over Λ.

Example 2.9 Here is a simple example to show that intersection numbers depend on splittings, not just on subgroups. First we note that the self-intersection number of any splitting is zero. Now construct a group G by amalgamating four groups G_1, G_2, G_3 and G_4 along a common subgroup Λ. Thus G can be expressed as $G_{12} *_\Lambda G_{34}$, where G_{ij} is the subgroup of G generated by G_i and G_j, but it can also be expressed as $G_{13} *_\Lambda G_{24}$ or $G_{14} *_\Lambda G_{23}$.

The intersection number of any distinct pair of these splittings of G is non-zero, but all the splittings being considered are splittings over the same group Λ.

A question which arose in our introduction in connection with the work of Rips and Sela was how the intersection number of two subgroups of a group G alters if one replaces G by a subgroup. In general, nothing can be said, but in interesting cases one can understand the answer to this question. The particular case considered by Rips and Sela was of a finitely presented group G which is expressed as the fundamental group of a graph of groups with some vertex group being a group H which contains infinite cyclic subgroups Λ and Σ. Further H is the fundamental group of a surface F and Λ and Σ are carried by simple closed curves L and S on F. A point deliberately left unclear in our earlier discussion of their work was that F is not a closed surface. It is a compact surface with non-empty boundary. The curves L and S are not homotopic to boundary components and so define splittings of H. The edges in the graph of groups which are attached to H all carry some subgroup of the fundamental group of a boundary component of F. This implies that L and S also define splittings of G. It is clear from this picture that the intersection number of Λ and Σ should be the same whether measured in G or in H, as it should equal the intersection number of the curves L and S, but this needs a little more thought to make precise. As usual, the first point to make is that we are really talking about the intersection numbers of the splittings defined by L and S, rather than intersection numbers of Λ and Σ. For the number of ends $e(H,\Lambda)$ and $e(H,\Sigma)$ are infinite when F is a surface with boundary. As G is finitely presented, we can attach cells to the boundary of F to construct a finite complex K with fundamental group G. Now the identification of the intersection number of the given splittings of G with the intersection number of L and S proceeds exactly as at the start of this section, where we showed how to identify the intersection number of the given splittings of H with the intersection number of L and S.

3 Interpreting intersection numbers

It is natural to ask what is the meaning of the intersection numbers defined in the previous section. The answer is already clear in the case of a surface group with cyclic subgroups. In this section, we will give an interpretation of the intersection number of two splittings of a finitely generated group G over finitely generated subgroups. We start by discussing the connection with the work of Kropholler and Roller.

In [6], Kropholler and Roller introduced an intersection cohomology class for $PD(n-1)$–subgroups of a PDn–group. The pairs involved always have two ends, so the work of the previous section defines an intersection number in this situation. The connection between our intersection number and their intersection cohomology class is the following. Recall that if one has subgroups Λ and Σ of a finitely generated group G, such that $e(G,\Lambda)$ and $e(G,\Sigma)$ are each equal to 2, then one chooses a non-trivial Λ–almost invariant subset X of G and a non-trivial Σ–almost invariant subset Y of G and defines our intersection number $i(\Lambda,\Sigma)$ to equal the number of double cosets $\Sigma g\Lambda$ such that gX crosses Y. Their cohomology class encodes the information about which double cosets have this crossing property. Thus their invariant is much finer than the intersection number and it is trivial to deduce the intersection number from their cohomology class.

To interpret the intersection number of two splittings of a group G, we need to discuss the Subgroup Theorem for amalgamated free products. Let G be a finitely generated group, which splits over finitely generated subgroups Λ and Σ. We will write $G = A_1 *_\Lambda (B_1)$ to denote that either G has the HNN structure $A_1 *_\Lambda$ or G has the structure $A_1 *_\Lambda B_1$. Similarly, we will write $G = A_2 *_\Sigma (B_2)$. The Subgroup Theorem, see [11] and [12] (or [13]) for discussions from the topological and algebraic points of view, yields a graph of groups structure $\Phi_1(\Sigma)$ for Σ, with vertex groups lying in conjugates of A_1 or B_1 and edge groups lying in conjugates of Λ. Typically this graph will not be finite or even locally finite. However, as Σ is finitely generated, there is a finite subgraph Ψ_1 which still carries Σ. If we reverse the roles of Λ and Σ, we will obtain a graph of groups structure $\Phi_2(\Lambda)$ for Λ, with vertex groups lying in conjugates of A_2 or B_2 and edge groups lying in conjugates of Σ, and there is a finite subgraph Ψ_2 which still carries Λ. We show below that, in most cases, the intersection number of Λ and Σ measures the minimal possible number of edges of these finite subgraphs. Notice that if we consider the special case when G is the fundamental group of a closed surface and Λ and Σ are infinite cyclic subgroups, this statement is clear. Now the symmetry of intersection numbers implies the surprising fact that the minimal number of edges for Ψ_1 and Ψ_2 are the same.

There is an alternative point of view which we will use for our proof. The splitting $A_2 *_\Sigma (B_2)$ of G corresponds to an action of G on a tree T such that the quotient $G\backslash T$ has one edge. The edge stabilisers in this action on T are all conjugate to Σ and the vertex stabilisers are conjugate to A_2 or B_2 as appropriate. If one has a subgroup Λ of G, the quotient $\Lambda\backslash T$ will be the graph underlying $\Phi_2(\Lambda)$. There is a Λ–invariant subtree T' of T, such that the graph $\Lambda\backslash T'$ is the graph underlying Ψ_2. Whichever point of view you take, it is necessary to connect it with the ideas about almost invariant sets which we have already discussed. Here is our interpretation of intersection numbers.

Theorem 3.1 *Let G be a finitely generated group, which splits over finitely generated subgroups Λ and Σ, such that if U and V are any conjugates of Λ and Σ respectively, then $U \cap V$ has infinite index in both U and V. Then the intersection number of the two splittings equals the minimal number of edges in each of the graphs Ψ_1 and Ψ_2.*

Remark 3.2 This result is clearly false if the condition on conjugates is omitted. For example, if $\Lambda = \Sigma$, then $\Psi_1(\Sigma)$ and $\Psi_2(\Lambda)$ will each consist of a single vertex, but the intersection number of the two splittings need not be zero.

The proof will use the following sequence of lemmas.

We start with a general result about minimal G–invariant subtrees of a tree T on which a group G acts. If every element of G fixes each point of a non-trivial subtree T' of T, then any vertex of T' is a minimal G–invariant subtree of T. Otherwise, there is a unique minimal G–invariant subtree of T. An orientation of an edge e of T consists of a choice of one vertex as the initial vertex $i(e)$ of e and the other as the terminal vertex $t(e)$. An oriented path in T consists of a finite sequence of oriented edges $e_1, e_2, \ldots, e_k$ of T, such that $t(e_j) = i(e_{j+1})$, for $1 \leq j \leq k-1$. If we consider two oriented edges e and e' of T we say that they are coherently oriented if there is an oriented path which begins with one and ends with the other. Finally, given an edge e of T and an element g of G, we will say that e and ge are coherently oriented if for some (and hence either) orientation on e and the induced orientation on ge, the edges e and ge are coherently oriented.

Lemma 3.3 *Suppose that a group G acts on a tree T without inversions and without fixing a point. Let T' denote the minimal G–invariant subtree. Then an edge e of T lies in T' if and only if there exists an element g of G such that e and ge are distinct and coherently oriented.*

Proof First consider an edge e not lying in T'. Orient e so that it is the first edge of an oriented path λ in T which starts with e, has no edge in T', and ends at a vertex of T'. Thus ge, with the induced orientation, is the first edge of an oriented path $g\lambda$ in T which starts with ge, has no edge in T', and ends at a vertex of T'. Now the unique path in T which joins e and ge must consist either of λ and $g\lambda$ together with a path in T' or of an initial segment of e together with an initial segment of ge. In either case, it follows that e and ge are not coherently oriented.

Now we consider an edge e of T' and its image $\overline{e}$ in $G\backslash T'$.

If $\overline{e}$ is non-separating in $G\backslash T'$, let μ denote an oriented path in $G\backslash T'$ which joins the ends of $\overline{e}$ and meets $\overline{e}$ only in its endpoints. Then the loop formed by $\mu \cup \overline{e}$ lifts to an oriented path in T', which shows that there is g in G such that e and ge are distinct and coherently oriented.

If $\overline{e}$ separates $G\backslash T'$, we can write the graph $G\backslash T'$ as $\Gamma_1 \cup \overline{e} \cup \Gamma_2$, where each Γ_i is connected and meets $\overline{e}$ in one endpoint only. Now consider the graph of groups structure given by $G\backslash T'$. By contracting each Γ_i to a point, we obtain an amalgamated free product structure of G as $G_1 *_C G_2$, where $C = stab(e)$ and each G_i is the fundamental group of the graph of groups Γ_i. Let T_i denote the tree on which G_i acts with quotient Γ_i. Then the complement in T' of the edge e and its translates consists of disjoint copies of T_1 and T_2. We identify T_i with the copy of T_i which meets e. Note that T_1 and T_2 are disjoint. Now it is clear that $G_1 \neq C \neq G_2$. For if $G_1 = C$, then $G = G_2$, which implies that T_2 is a G–invariant subtree of T', contradicting the minimality of T'. As $G_1 \neq C$, there is an element g_1 of G_1 such that $g_1 e \neq e$, and similarly there is an element g_2 of G_2 such that $g_2 e \neq e$. For each i, there is a path λ_i in T_i which begins at e and ends at $g_i e$. As T_1 and T_2 are disjoint, so are λ_1 and λ_2. It follows that of the three edges $e, g_1 e, g_2 e$, at least one pair is coherently oriented, which completes the proof of the lemma. $\square$

The following result is clear.

Lemma 3.4 *Suppose that a group G acts on a tree T without inversions and without fixing a point. Let e denote an edge of T, let E denote a component of $T - \{e\}$ and let g denote an element of G. Then e and ge are distinct and coherently oriented if and only if either $gE \subsetneqq E$ or $gE^* \subsetneqq E^*$.*

Next we need to connect this with almost invariant sets, although the following result does not use the almost invariance property.

Lemma 3.5 *Suppose that a group G acts on a tree T without inversions and without fixing a point and suppose that the quotient graph $G\backslash T$ has only one edge. Let e denote an edge of T, let E denote a component of $T - \{e\}$ and let $Y = \{k \in G : ke \subset E\}$. Then the following statements hold for all elements g of G :*

1) $gY \subset Y$ *if and only if* $gE \subset E$, *and* $gY^* \subset Y^*$ *if and only if* $gE^* \subset E^*$.

2) $gY = Y$ *if and only if* $gE = E$, *and* $gY^* = Y^*$ *if and only if* $gE^* = E^*$.

3) $gY \subsetneqq Y$ *if and only if* $gE \subsetneqq E$, *and* $gY^* \subsetneqq Y^*$ *if and only if* $gE^* \subsetneqq E^*$.

Proof Suppose that $gE \subset E$. If k lies in Y, then $ke \subset E$, so that $gke \subset gE \subset E$. Thus gk also lies in Y. It follows that $gY \subset Y$.

Conversely, suppose that $gY \subset Y$ and consider an edge f of E. As $G\backslash T$ has only one edge, $f = ke$ for some k in G. As f lies in E, k lies in Y, and hence gk also lies in Y by our assumption that $gY \subset Y$. Thus $gke \subset E$, so that $gf \subset E$. Thus implies that $gE \subset E$ as required.

The proof for the second equivalence in part 1 is essentially the same.

The equivalences in part 2 follow by applying part 1 for g and g^{-1}. Now the equivalences in part 3 are clear. □

Next we connect the above inclusions with crossing of sets.

Lemma 3.6 *Suppose that a finitely generated group G splits over a finitely generated subgroup Λ with corresponding Λ–almost invariant set X and also splits over a finitely generated subgroup Σ with corresponding Σ–almost invariant set Y. Suppose further that if U and V are any conjugates of Λ and Σ respectively, then $U \cap V$ has infinite index in U. Then X crosses Y if and only if there is an element λ in Λ such that either $\lambda Y \subsetneqq Y$ or $\lambda Y^* \subsetneqq Y^*$.*

Proof We claim that there exists $\lambda_1 \in \Lambda$ such that either $\lambda_1 Y \subsetneqq Y$ or $\lambda_1 Y^* \subsetneqq Y$, and there exists $\lambda_2 \in \Lambda$ such that either $\lambda_2 Y \subsetneqq Y^*$ or $\lambda_2 Y^* \subsetneqq Y^*$. Assuming this, either $\lambda_1 Y \subsetneqq Y$ or $\lambda_2 Y^* \subsetneqq Y^*$, and our proof is complete, or we have $\lambda_1 Y^* \subsetneqq Y$ and $\lambda_2 Y \subsetneqq Y^*$. The last possibility implies that $\lambda_2\lambda_1 Y^* \subsetneqq \lambda_2 Y \subsetneqq Y^*$, again completing the proof.

To prove our claim, we pick a finite generating set for G, and consider the Cayley graph Γ of G with respect to this generating set. As Y is a Σ–almost invariant set associated to a splitting $A_2 *_\Sigma (B_2)$ of G over Σ, we can choose Γ and Y so that, for every g in G, $g\delta Y$ is disjoint from or coincides with δY. A simple way to arrange this is to take as generators of G the union of a set of generators of Σ and of A_2 and B_2, so that $\Gamma(G)$ contains a copy of the Cayley graph $\Gamma(\Sigma)$ of Σ and $\Sigma\backslash\Gamma$ contains $\Sigma\backslash\Gamma(\Sigma)$ which is a wedge of circles. (Note that this uses the hypothesis that Σ is finitely generated.) Let v denote the wedge point, and let E denote the collection of vertices of $\Sigma\backslash\Gamma$ which can be joined to v by a path whose interior is disjoint from v such that the last edge is labelled by an element of A. Then clearly δE consists of exactly those edges of $\Sigma\backslash\Gamma$ which have one end at v and are labelled by an element of A. Further, if we let Y denote the pre-image of E in G, then, for every g in G, $g\delta Y$ is disjoint from or coincides with δY.

In order to prove that λ_1 exists, we argue as follows. As $\Lambda\cap\Sigma$ has infinite index in Λ, and as δX is Λ–invariant, it follows that δX must contain points which are arbitrarily far from δY on each side of δY. Recall that $\Lambda\backslash X$ is an almost invariant subset of $\Lambda\backslash G$, so that it has finite coboundary which equals $\Lambda\backslash\delta X$. Hence there is a number d such that any point of $\Lambda\backslash\delta X$ can be joined to the image of δY in $\Lambda\backslash\Gamma$ by a path of length at most d. It follows that any point of δX can be joined to $\lambda\delta Y$, for some λ in Λ, by a path in Γ of length at most d. Hence there is a translate of δY which contains points on one side of δY and another translate which contains points on the other side of δY. Hence there are elements λ_1 and λ_2 of Λ such that $\lambda_1\delta Y$ lies on one side of δY and $\lambda_2\delta Y$ lies on the other. Without loss of generality, we can suppose that $\lambda_1\delta Y$ lies on the side containing Y so that either $\lambda_1 Y \subsetneqq Y$ or $\lambda_1 Y^* \subsetneqq Y$. As $\lambda_2\delta Y$ lies on the side of δY containing Y^*, either $\lambda_2 Y \subsetneqq Y^*$ or $\lambda_2 Y^* \subsetneqq Y^*$. This completes the proof of the claim made at the start of the proof. □

Now we can give the proof of Theorem 3.1.

Proof Recall that G splits over finitely generated subgroups Λ and Σ such that if U and V are any conjugates of Λ and Σ, then $U\cap V$ has infinite index in both U and V. Also G acts on a tree T so as to induce the given splitting over Σ. Let e denote an edge of T with stabiliser Σ and consider the action of Λ on T. Our hypothesis on conjugates of Λ and Σ implies, in particular, that Λ is not contained in any conjugate of Σ so that Λ cannot fix an edge of T. Thus there is a unique minimal Λ–invariant subtree T' of T. Lemma 3.3 shows that an edge he of T lies in T' if and only if there is λ in Λ such that he and λhe are distinct and coherently oriented. Lemma 3.4 shows that this occurs if and only if either $\lambda hE \subsetneqq hE$ or $\lambda hE^* \subsetneqq hE^*$, and Lemma 3.5 shows that this occurs if and only if $\lambda hY \subsetneqq hY$ or $\lambda hY^* \subsetneqq hY^*$. Finally Lemma 3.6 shows that this occurs if and only if X crosses hY. We conclude that an edge he of T lies in T' if and only if X crosses hY. Thus the edges of T which lie in the minimal Λ–invariant subtree T' naturally correspond to the cosets $h\Sigma$ such that X crosses hY. Hence the number of edges in $\Psi_2(\Lambda)$ equals the number of double cosets $\Lambda h\Sigma$ such that X crosses hY, which was defined to be the intersection number of the given splittings. Similarly, one can show that the intersection number of the given splittings equals the minimal possible number of edges in the graph $\Psi_1(\Sigma)$. This completes the proof of Theorem 3.1. $\square$

References

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

[2] **M Dunwoody**, **M Sageev**, *JSJ splittings for finitely presented groups over slender groups*, preprint

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

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

[5] **C H Houghton**, *Ends of locally compact groups and their quotient spaces*, J. Aust. Math. Soc. 17 (1974) 274–284

[6] **P H Kropholler**, **M A Roller**, *Splittings of Poincaré duality groups*, Math. Zeit. 197 (1988) no. 3, 421–438

[7] **E Rips**, **Z Sela**, *Cyclic splittings of finitelypresented groups and the canonical JSJ decomposition*, preprint

[8] **P Scott**, *Ends of pairs of groups*, J. Pure Appl. Algebra 11 (1977) 179–198

[9] **P Scott**, *A new proof of the Annulus and Torus Theorems*, Amer. J. Math. 102 (1980) 241–277

[10] **P Scott**, **G A Swarup**, *An Algebraic Annulus Theorem*, preprint

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

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

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