

Splittings of Non-Finitely Generated Groups

by

Robin M. Lassonde

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2012

Doctoral Committee:

Professor G. Peter Scott, Chair
Professor Richard Canary
Professor Ralf J. Spatzier
Professor John Stembridge
Professor Martin Strauss

ACKNOWLEDGEMENTS

Thanks to my advisor Peter Scott for posing many of the questions answered in this dissertation, for many productive discussions, and for your mathematical mentorship over the past three years.

Thanks to my dissertation committee for helping me through the final hurdle of graduate school, and especially to my readers Peter Scott and Dick Canary for making many helpful suggestions.

Thanks to my parents Beth Kean, Greg Lassonde, Sarah Davis and Raleigh McLemore for being spectacular role models and encouraging me to always do my best, to all my family and friends for your support over the years, and to my loving husband Brian Lassiter.

TABLE OF CONTENTS

| | |
|--|-----------|
| ACKNOWLEDGEMENTS | ii |
| LIST OF FIGURES | v |
| CHAPTER | |
| 1. Introduction and Contents | 1 |
| 1.1 Introduction | 1 |
| 1.2 Contents | 2 |
| 2. History | 7 |
| 2.3 Ends | 7 |
| 2.4 Splittings | 8 |
| 2.5 Stallings' Theorem | 8 |
| 2.6 Almost Invariant Sets | 8 |
| 2.7 JSJ Decompositions | 9 |
| 3. Preliminaries and Main Ideas | 11 |
| 3.8 Splittings | 11 |
| 3.9 Almost Invariant Sets and Crossing | 13 |
| 3.10 Almost Invariant Sets Arising from Splittings | 14 |
| 3.11 Dunwoody's Theorem | 16 |
| 3.12 Almost Inclusion and Small Corners | 17 |
| 3.13 Example: Simple Closed Curves on a Surface | 19 |
| 3.14 Sandwiching | 21 |
| 3.15 Compatibility | 23 |
| 3.16 Sandwiching is Necessary | 24 |
| 3.17 Turning Almost Inclusion Into Inclusion | 24 |
| 3.18 Algebraic Regular Neighborhoods | 25 |
| 4. Symmetry of Crossing | 29 |
| 5. Examples of Infinite Intersection Number | 32 |
| 6. Almost Inclusion | 35 |
| 7. Compatibility and Intersection Number Zero | 42 |

| | |
|--|-----------|
| 8. CAT(0) Cubical Complexes and Positive Intersection Number | 47 |
| 8.1 Producing Almost Invariant Sets from a CAT(0) cubical complex . | 48 |
| 8.2 Ultrafilters | 50 |
| 8.3 Sageev's Cubing | 50 |
| 8.4 Minimal Cubings | 53 |
| 8.5 Putting the X_j 's in Very Good Position | 61 |
| 9. Existence of Algebraic Regular Neighborhoods | 65 |
| 10. Uniqueness of Algebraic Regular Neighborhoods | 70 |
| 11. Mixed Almost Invariant Sets when G is Finitely Generated | 75 |
| BIBLIOGRAPHY | 79 |

LIST OF FIGURES

Figure

| | | |
|-----|--|----|
| 3.1 | Separating loop | 20 |
| 3.2 | Non-separating loop | 20 |
| 3.3 | Two loops that do not intersect | 21 |
| 3.4 | Two loops that intersect trivially | 21 |
| 3.5 | Regular neighborhood | 28 |
| 4.1 | Proof of key lemma | 31 |
| 5.1 | Infinite intersection number | 33 |
| 5.2 | Infinite self-intersection number | 33 |
| 7.1 | Branching | 43 |
| 8.1 | Dual graph is part of the cubing | 48 |

CHAPTER 1

Introduction and Contents

1.1 Introduction

In this paper, we investigate the intersection number of splittings of a group. Splittings are analogues of codimension–1 embedded submanifolds. More precisely, a splitting of a group G is a one–edged graph of groups structure for G . In this paper, I describe splittings in terms of G –trees and prefer not to use graphs of groups. For example, from the perspective of graphs of groups, it is not completely obvious how to define an isomorphism of splittings. On the other hand, from the perspective of G –trees, clearly we should define two splittings to be isomorphic precisely when their G –trees are isomorphic.

Almost invariant sets are analogues of codimension–1 immersed submanifolds. In [23], Scott defined the intersection number of two almost invariant subsets of a finitely generated group, and proved that the definition is symmetric. Shortly after, Scott and Swarup showed that if two splittings have intersection number zero, then one can find a common refinement of their Bass–Serre trees [24]. The same authors further developed these concepts in [25] to construct algebraic regular neighborhoods of finite collections of almost invariant sets. In [18], Niblo–Sageev–Scott–Swarup showed how, given finitely many almost invariant sets, to produce a CAT(0) cubical complex, in which meeting of hyperplanes corresponds to crossing of almost invariant

sets. The above results applied only to finitely generated groups, and often required further that the stabilizer of each almost invariant set (and, in particular, the subgroups associated to the splittings in question) be finitely generated. Here, we use the properties of splittings to remove the finite generation assumptions in the case where the almost invariant sets come from splittings. Instead of using the Cayley graph, which is only useful when G is finitely generated, we turn to the Bass–Serre trees for the splittings. We also show how to use the aforementioned CAT(0) cubical complex to make a more direct construction of algebraic regular neighborhoods.

1.2 Contents

Chapter 2 provides a history of splittings. Chapter 3 gives an explanation of the main concepts used in this paper. It is recommended that the reader review Chapter 3 before proceeding.

For the next two paragraphs, let G be a group, X be a nontrivial H -almost invariant subset of G , and Y a nontrivial K -almost invariant subset of G . In [23], Scott defined “ X crosses Y ” and proved that if G is finitely generated, then this definition is symmetric. By counting the number of group elements $g \in G$ such that gX crosses Y , Scott gave a well-defined, symmetric intersection number of X and Y . To define the intersection number of two splittings of G , pick associated almost invariant sets and take their intersection number. In this paper, I prove that if Y arises from a splitting of G , and if X crosses Y , then Y crosses X , without any assumption of finite generation.

Proposition 4.2. *Let G be any group with subgroups H and K . Suppose X is any nontrivial H -almost invariant set, and Y is a K -almost invariant set arising from a splitting of G over K . If X crosses Y , then Y crosses X .*

Hence if both X and Y arise from splittings, then crossing is symmetric. In particular, for any group (not necessarily finitely generated), the intersection number of two splittings is well-defined. The key argument used to prove this is laid out in Lemma 4.1.

Also in [23], Scott proved that if G , H and K are all finitely generated, then the intersection number of X and Y is finite. I give two examples showing that the assumption that G be finitely generated is crucial: when G is the free group on countably many generators, it is possible for two splittings over the trivial group to have an infinite intersection number (Example 5.1), and also possible for a $\{1\}$ -almost invariant set to have infinite self-intersection number (Example 5.2). I also explain an example, due to Guirardel, of infinite intersection when G is finitely generated but H and K are not (Example 5.3).

For this paragraph, let G be a group, and let σ_j be a splitting of G over H_j , for $j = 1, \dots, n$. In [24], Scott and Swarup proved that if G and all the H_j 's are finitely generated, and the σ_j 's have pairwise intersection number zero, then the σ_j 's are compatible. In this paper, I generalize Scott and Swarup's result, replacing the finite generation assumptions by a sandwiching assumption.

Theorem 7.5. *Let σ_j be a splitting of G over H_j and assume $\{\sigma_1, \dots, \sigma_n\}$ satisfies sandwiching. If $i(\sigma_j, \sigma_k) = 0$, for all j and k , then the splittings $\{\sigma_1, \dots, \sigma_n\}$ are compatible.*

Sandwiching is automatic if none of the splittings is a trivially ascending HNN extension (see Corollary 7.3), and if the σ_j 's do not satisfy sandwiching, then the σ_j 's cannot be compatible (see Chapter 3.16). My proof of the theorem mirrors the proof in [24], replacing the coboundary arguments by new arguments using G -trees for splittings. In particular:

- Scott and Swarup used Cayley graph arguments to show that “almost inclusion” defines a partial order on the set of all translates of all the almost invariant sets (arising from the splittings with intersection number zero) and their complements. I show that the fact that almost inclusion defines a partial order can be deduced directly once one has symmetry of crossing (see Corollary 6.5).
- Scott and Swarup used Cayley graph arguments to prove interval finiteness for finite collections of almost invariant sets. Their arguments require both G and the associated subgroups to be finitely generated. In the case where the almost invariant sets arise from splittings, I show how to deduce interval finiteness from sandwiching (see Proposition 7.4).

For the remainder of this chapter, let $\{X_j | j = 1 \dots n\}$ be a finite collection of nontrivial H_j -almost invariant subsets of a group G , let Σ denote the set of all translates of the X_j 's and their complements, and let \leq denote almost inclusion on Σ .

In [21], Sageev constructed a CAT(0) cubical complex from the partially ordered set (Σ, \subset) . In [18], Niblo, Sageev, Scott and Swarup generalized Sageev's construction, using the partial order of almost inclusion on Σ , instead of inclusion, to get a “minimal” cubing. Their results assumed the ambient group G , as well as the H_j 's, to be finitely generated. In this paper, I remove the finite generation assumptions and instead assume that all the X_j 's come from splittings that collectively satisfy sandwiching.

Theorem 8.7. *Let G be any group with a finite collection $\{\sigma_j | j = 1, \dots, n\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j = 1, \dots, n\}$ satisfies sandwiching. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$.*

Then there exists a CAT(0) cubical complex L , with a bijective correspondence between Σ and the set of oriented hyperplanes of L , such that two elements of Σ cross if, and only if, their corresponding hyperplanes cross in some square. Moreover, the hyperplane corresponding to any $A \in \Sigma$ determines a $\text{Stab}(A)$ -almost invariant set that is $\text{Stab}(A)$ -almost equal to A .

The main challenge in adapting the Cayley graph arguments from [18] to the non-finitely generated case is to show that the cubing is nonempty. In particular:

- Constructing an ultrafilter on (Σ, \leq) (see the first half of proof of Theorem 8.11), and
- Proving that this ultrafilter satisfies the descending chain condition (see the second half of proof of Theorem 8.11).

One application of minimal cubings is putting the X_j 's in “very good position,” i.e. perturbing each X_j such that inclusion and almost inclusion give the same partial order on Σ .

In [25], Scott and Swarup defined the algebraic regular neighborhood of a collection of almost invariant subsets of G , assuming G is finitely generated. They proved the existence of an algebraic regular neighborhood for a finite family of almost invariant sets having finitely generated stabilizers, and uniqueness of algebraic regular neighborhoods for arbitrary collections of almost invariant sets. Scott and Swarup presumed that the minimal cubing from [18] could be turned into an algebraic regular neighborhood of the X_j 's. However, they were unable to prove that the object they constructed satisfied the definition of an algebraic regular neighborhood, because edge stabilizers might not be finitely generated. I include the missing arguments in Chapter 11. This result will be used in a forthcoming paper by Guirardel, Scott and

Swarup on relative versions of the algebraic torus theorem and other results. These results concern splittings “adapted” to a family of subgroups, a concept that was introduced by Müller in [17].

In Theorem 9.1, I show how to turn a minimal cubing into an algebraic regular neighborhood under my hypotheses (i.e. that the X_j ’s come from splittings satisfying sandwiching, and neither G nor any of the H_j ’s need to be finitely generated). In Theorem 10.1, I prove that algebraic regular neighborhoods are unique, even for a possibly infinite collection of splittings, and without any finite generation assumptions.

CHAPTER 2

History

Here I review the history leading up to the theory described above.

2.3 Ends

In 1931, Freudenthal defined the number of ends of a topological space [6]. Roughly speaking, the ends of a space are the space's "connected components at infinity." To count the number of ends of a locally finite CW complex, simply remove a finite number of open cells, and count the number of infinite components remaining. The number of ends is the supremum over all such removals. In 1944, Hopf realized that the number of ends of the Cayley graph of a finitely generated group does not depend on the choice of generating set; hence one can define the number of ends of a finitely generated group to be the number of ends of its Cayley graph [12]. In fact, if a finitely generated group acts freely and cocompactly on a locally finite space, the number of ends of the space is the same as the number of ends of a group. Several years later, Specker introduced a purely algebraic definition for the number of ends of any group [31]. Usually when we think of the number of ends of a group, we are thinking of the geometric interpretation (which only works for finitely generated groups). However, it is good to know that the definition can be extended to non-finitely generated groups.

2.4 Splittings

Group splittings were defined around the same time as ends. Schreier introduced amalgamated free products in 1927 [22], and two decades later, Higman–Neumann–Neumann introduced HNN extensions [11]. Both amalgamated free products and HNN extensions were initially described in terms of normal forms for words. In 1977, Serre discovered that splittings can be described as group actions on trees. This topic is known as “Bass–Serre theory” [29, 30]. Immediately after, Scott and Wall noted that one can use the Seifert–van-Kampen Theorem to realize any graph of groups as a graph of spaces [27]. From Scott’s point of view, there is no reason to distinguish between amalgamated free products and HNN extensions, so he called both “splittings.”

2.5 Stallings’ Theorem

A decade before Serre’s discovery that group actions on trees correspond to splittings, Stallings made a connection between splittings and ends of groups. Stallings’ theorem states that a finitely generated group G has at least two ends if, and only if, G splits over a finite subgroup [33, 32]. He also showed that groups of cohomological dimension one are free. Swan extended Stallings’ results to non-finitely generated groups [34].

2.6 Almost Invariant Sets

Also in the 1970’s, Cohen coined the term “almost invariant set” over the trivial subgroup, as a way of keeping track of the ends of a group [2]. A group G has at least two ends if, and only if, G has a nontrivial $\{1\}$ -almost invariant subset. Twenty years earlier had Specker used almost invariant sets in his paper [31]; however, this fact has been entirely overlooked in the history of almost invariant sets, in part because

Specker was interested in something more general. Cohen also observed that a subset of a finitely generated group is almost invariant if, and only if, the subset has finite coboundary in the Cayley graph of the group. Houghton formally defined the number of ends of a pair (G, H) of groups [13], and one can make a similar observation that a subset of a finitely generated group G is H -almost invariant if, and only if, the subset is stabilized under left multiplication by elements of H and its image has finite coboundary in the Cayley graph of G quotiented out by H .

Dunwoody used Bass–Serre theory and Cohen’s almost invariant sets to produce a beautiful geometric proof of Stallings’ theorem, in which one takes a suitable “end” of a finitely generated group G and uses it to directly construct a Bass–Serre tree (see Chapter 3.11, [3], and [27, Section 6]). Sageev showed how to construct a CAT(0) cubical complex from an almost invariant set [21], and Niblo used Sageev’s construction to produce another geometric proof of Stallings’ theorem [19].

2.7 JSJ Decompositions

Classical JSJ theory for 3–manifolds was initiated by Waldhausen [35] in 1969, and was developed in 1979 by Jaco and Shalen [14], and Johannson [15]. The basic idea is to describe all annuli and tori in a compact, irreducible, orientable 3–manifold. If one restricts to closed 3–manifolds, then only tori are needed. Since 1990, geometric group theorists have been developing group theoretic analogues of classical JSJ theory, including Kropholler [16], Sela [28], Rips–Sela [20], Bowditch [1], Dunwoody–Sageev [4], Dunwoody–Swenson [5], Fujiwara–Papasoglu [7], Scott–Swarup [25], and Guirardel–Levitt [9, 10]. The classical JSJ decomposition for 3–manifolds is unique. Ideally one would like group theoretic JSJ decompositions that are unique, and can simultaneously handle splittings over multiple types of edge groups (for example, virtually polycyclic groups of arbitrary lengths). However, previous attempts have

had limited success developing decompositions addressing both of these issues. The JSJ decompositions constructed by Scott and Swarup in [25] are built from algebraic regular neighborhoods and are unique, and I am hopeful that the results on algebraic regular neighborhoods that I prove in this paper will lead to more generalized JSJ decompositions.

CHAPTER 3

Preliminaries and Main Ideas

The purpose of this chapter is to:

1. Give enough background information to enable to reader to understand the statements of all the results in this paper, and
2. Provide an idea of how everything fits together.

3.8 Splittings

A G -tree is a simplicial tree equipped with a (simplicial) G -action, such that the action does not invert any edges. A *splitting* of a group G is a G -tree T with no global fixed points, such that the quotient graph $G \backslash T$ has exactly one edge. There are two cases:

1. $G \backslash T$ consists of one edge with distinct endpoints. Pick an edge e of T , let H denote the stabilizer of e , and let A and B denote the stabilizers of the endpoints of e . We have inclusions $i_1: H \hookrightarrow A$ and $i_2: H \hookrightarrow B$. We call σ an *amalgamated free product* and may write $\sigma: G \cong A *_H B$. Note that using Bass-Serre theory [30], we could reconstruct T using only the inclusions $i_1: H \hookrightarrow A$ and $i_2: H \hookrightarrow B$. One presentation for G is $\langle A, B | i_1(h) = i_2(h), \text{ for all } h \in H \rangle$.
2. $G \backslash T$ consists of a loop with one edge and one vertex. Pick an edge e of T , let H

denote the stabilizer of e , and let A denote the stabilizer of one of the endpoints of e . We have an inclusion $i_1: H \hookrightarrow A$. As there is only one orbit of vertices, the other endpoint of e is a translate of the first endpoint by some $t \in G$. The stabilizer of this vertex is tAt^{-1} , so we have an inclusion $H \hookrightarrow tAt^{-1} \cong A$. We can view this second inclusion as $i_2: H \hookrightarrow A$, where $i_2(h) := t^{-1}i_1(h)t$. We call σ an *HNN extension* and may write $\sigma: G \cong A*_H$. Note that using Bass–Serre theory [30], we could recover T using only the two inclusions $i_1, i_2: H \hookrightarrow A$. One presentation for G is $\langle A, t | i_2(h) = t^{-1}i_1(h)t, \text{ for all } h \in H \rangle$.

In either case, σ is called a *splitting of G over H* . The subgroup H is well-defined up to conjugacy in G .

If the G -tree for a splitting is a line on which G acts by translations only, we call the splitting a *trivially ascending HNN extension*. Note that each edge in the line has the same stabilizer, denoted H , so that H acts trivially on the line. Equivalently, H is normal in G and $\frac{G}{H} \cong \mathbb{Z}$. Equivalently, the splitting has the form $A*_H$ where both inclusions $i_1: H \hookrightarrow A$ and $i_2: H \hookrightarrow A$ are isomorphisms.

Definition 3.1. Two splittings of a group G are *isomorphic* if there exists a G -equivariant isomorphism between the trees for the two splittings.

To describe how two splittings of G cross, for each splitting we will construct a subset of G , and then look at how the two subsets and their translates cross. We show how to construct the subsets in Chapter 3.10. In most of this paper, existence results (in particular, Theorems 7.5, 8.12, and 9.1) only work for a finite collection of splittings, while uniqueness results (namely, Corollary 6.6 and Theorem 10.1) do not require such an assumption.

3.9 Almost Invariant Sets and Crossing

Many concepts used here consider subgroups of G “up to finite index” and subsets of G “up to finitely many cosets.” Here are a few key definitions capturing this idea. Note that in this paper, unless otherwise specified, all cosets are *right cosets*, i.e. of the form Hg for some $g \in G$.

Definition 3.2. Let H and K be subgroups of a group G .

- H and K are *commensurable* if $H \cap K$ has finite index in H and in K .
- A subset of G is *H -finite* if it is contained in only finitely many right cosets Hg_i of H in G .
- Two subsets A and B of G are *H -almost equal*, written $A \stackrel{H\text{-a}}{=} B$, if their symmetric difference is H -finite.

An almost invariant subset of a group is a subset which does not change by much when you multiply on the right by an element of the group. Specifically:

Definition 3.3. Let G be any group, H a subgroup of G , and X a subset of G .

We say X is an *H -almost invariant* subset of G if the following two properties are satisfied:

1. H stabilizes X , i.e. $hX = X$, for all $h \in H$, and
2. $Xg \stackrel{H\text{-a}}{=} X$, i.e. the symmetric difference of Xg and X is H -finite, for all $g \in G$.

Call X *trivial* if it or its complement is H -almost equal to the empty set.

Let X be H -almost invariant and Y be K -almost invariant. We call the four sets $X \cap Y$, $X \cap Y^*$, $X^* \cap Y$, and $X^* \cap Y^*$ the *corners of the pair* (X, Y) .

Definition 3.4. Let X and Y be subsets of G . Then X and Y are *nested* if X or X^* is a subset of Y or Y^* , i.e. a corner of the pair (X, Y) is empty. Otherwise, X and Y are *not nested*.

We would like a similar notion that works “up to finitely many cosets.”

Definition 3.5. Let X be an H -almost invariant subset of G , and Y a K -almost invariant subset of G . The pair (X, Y) is *almost nested* if a corner of the pair (X, Y) is K -finite.

Otherwise, X *crosses* Y , i.e. no corner of the pair (X, Y) is K -finite.

A couple of facts justify this terminology:

1. If X and Y arise from splittings of G (see Chapter 3.10), or if X and Y do not necessarily come from splittings but G is finitely generated, then X crosses Y if, and only if, Y crosses X . For the proof in the case where X and Y come from splittings, see Proposition 4.2. For the proof in the case where G is finitely generated, see [23, Lemma 2.3].
2. If Y is both K -almost invariant and K' -almost invariant, then K and K' must be commensurable (see Lemma 6.1); in particular, K -finiteness is the same as K' -finiteness. Also note that if Y is K -almost invariant, then $K \subset \text{Stab}(Y)$, so that K must be a finite-index subgroup of $\text{Stab}(Y)$.

3.10 Almost Invariant Sets Arising from Splittings

Scott and Swarup [25] noted that given a splitting of a group, one can produce an almost invariant subset as follows. Let σ be a splitting of G , and let T be a G -tree for σ . Pick a base vertex v and a (directed) edge e of T . Define a subset X of G by:

$$X := \{g \in G \mid e \text{ points away from } gv\}.$$

Let H denote the stabilizer of e . Such X is in fact an H -almost invariant subset of G (for a proof, see Corollary 8.2). Some almost invariant sets arise from splittings, and others do not. Note that if X is an almost invariant set arising from σ , then so is each translate of X or its complement.

We will often go back and forth between a splitting σ and an almost invariant set X arising from σ . Given a splitting σ , we can construct X by picking a base vertex and edge in the tree for σ . Given an almost invariant set X arising from a splitting σ , let Σ denote the set of all translates of X and its complement, partially ordered by inclusion. We can apply Dunwoody's theorem (see Chapter 3.11) to produce a G -tree. This tree will yield a splitting isomorphic to σ (see Proposition 6.3 for a proof). The choice of base edge e and base vertex v is necessary to define X , but it does not particularly matter which one we choose. A change in e will result in X being replaced by a translate of X or its complement (and hence does not change the set Σ). A change in v will result in an almost invariant set $X' \stackrel{H\text{-a}}{=} X$ (see Corollary 8.2), and hence yields a splitting isomorphic to σ (see Proposition 6.3).

We will use the following convention: “ X is an H -almost invariant set arising from the splitting σ ” implicitly means that H is equal to the stabilizer of X (as opposed to a proper subgroup of $\text{Stab}(X)$), and σ is a splitting of G over H .

Now we are ready to define the intersection number of two splittings σ and τ of G . Let X be an H -almost invariant set arising from σ , and Y a K -almost invariant set arising from τ . To compute the intersection number of X and Y , count the number of $g \in G$ such that gX crosses Y , then eliminate double-counting. If $h \in H$ and $k \in K$, then $hX = X$ and $kY = Y$, so gX crossing Y is the same as $k^{-1}ghX$ crossing Y . Define the *intersection number of σ and τ* by:

$$i(\sigma, \tau) := \text{number of double cosets } KgH \text{ such that } gX \text{ crosses } Y.$$

In this paper, we will mostly only care whether the intersection number of two splittings is non-zero, i.e. whether any translates of X and Y cross each other.

3.11 Dunwoody's Theorem

Dunwoody's Theorem takes a partially ordered set satisfying tree-like properties, and produces a tree [3]. Some applications include Dunwoody's proof of Stallings' theorem [3, 27], reconstructing the Bass–Serre tree for a splitting by using an almost invariant set that came from the splitting (see Chapter 3.10), and constructing a common refinement for trees representing two splittings that have intersection number zero (see Chapter 3.15).

Take any simplicial tree and let Σ denote its (directed) edge set. Reversing the direction of an edge gives a free involution $*$ on Σ , and we can describe an undirected edge as a pair $\{e, e^*\}$. The set Σ satisfies the following tree-like properties:

- If there is an edge path starting with the edge e and ending with f , then there is an edge path starting with f^* and ending with e^* .
- For any two undirected edges, there is a simple edge path connecting them.
- Σ has no loops.

One can define a partial order on Σ as follows:

$$e \leq f \iff \text{there exists a simple edge path starting at } e \text{ and ending at } f.$$

Dunwoody's theorem states that any partially ordered set (Σ, \leq) satisfying analogous properties can be turned into the edge set of a tree.

Theorem 3.6 (Dunwoody's Theorem). *Let (Σ, \leq) be a partially ordered set equipped with a free involution $*$ on Σ . Suppose the following conditions are satisfied:*

1. For all $A, B \in \Sigma$, if $A \leq B$, then $B^* \leq A^*$.
2. For all $A, B \in \Sigma$ with $A \leq B$, there are only finitely many $C \in \Sigma$ with $A \leq C \leq B$.
3. For all $A, B \in \Sigma$, at least one of the four relations $A^{(*)} \leq B^{(*)}$ holds.
4. There are no two elements $A, B \in \Sigma$ such that simultaneously $A \leq B$ and $A \leq B^*$.

Then there exists a tree T with (directed) edge set Σ , and such that $A \leq B$ if, and only if, there exists a simple edge path whose first edge is A and whose last edge is B .

The key idea in the proof of Dunwoody's theorem is constructing the vertices of T . Let each element of Σ be a directed edge, and make a vertex wherever there are two edges with nothing in between. Specifically, define the vertices of the tree to be equivalence classes of elements of Σ :

$$[e] = [f] \iff \begin{array}{l} e \leq f^*, \text{ AND} \\ \text{if } e \leq a \leq f^*, \text{ then } a = e \text{ or } a = f^*. \end{array}$$

Then one must prove that everything works out.

3.12 Almost Inclusion and Small Corners

Given an H -almost invariant set X , we will often want to refer to the set of all translates of X and its complement. Denote this set by $\Sigma(X)$.

$$\Sigma(X) := \{gX, gX^* | g \in G\}.$$

For the remainder of the paragraph, let J be some indexing set, and for each $j \in J$, let X_j be an H_j -almost invariant subset of G . Let Σ denote $\bigcup_{j \in J} \Sigma(X_j)$. Assume

that we are in a situation where crossing is symmetric – for example, assume G is finitely generated or assume that each X_j comes from a splitting. Given $A, B \in \Sigma$, we say a corner of the pair (A, B) is *small* if it is $\text{Stab}(A)$ –finite. By the above remarks, this is equivalent to the corner being $\text{Stab}(B)$ –finite. In this paper, we will only use the term “small” when we know that crossing is symmetric.

Inclusion partially orders Σ ; however, we would prefer a partial order that is not affected by changing some X_j by finitely many H_j –cosets. The obvious thing to do is declare $A \leq B$ precisely when $A \cap B^*$ is small. However, we run into a potential difficulty: if two corners of a given pair (A, B) are small, how do we decide which inequality to choose? If two corners are small and one of them is empty, then we choose to only pay attention to the empty corner. For example, if $A \cap B^*$ and $A^* \cap B$ are small, and $A^* \cap B$ is empty (i.e. $B \subset A$), then we declare $B \leq A$.

Definition 3.7. We say Σ is in *good position* if for all $A, B \in \Sigma$, whenever two corners of the pair (A, B) are small, one is empty.

If Σ is in good position, then one can define a partial order \leq on Σ as follows.

Definition 3.8 (Almost Inclusion). Let G be any group, and Σ any collection of almost invariant subsets of G . Assume crossing is symmetric (for example, assume the elements of Σ arise from splittings of G). Define a relation \leq on Σ by:

$$A \leq B \iff A \cap B^* \text{ is empty, or is the only small corner of the pair } (A, B).$$

If Σ is in good position, then \leq is a partial order on Σ (see Corollary 6.5 for a proof). I show that good position is automatic if the X_j ’s arise from non-isomorphic splittings (see Corollary 6.4).

3.13 Example: Simple Closed Curves on a Surface

To gain more intuition about splittings, we look at a few concrete examples. Let S be a closed, orientable surface of genus at least two. Let G denote the fundamental group of S . Let γ be a π_1 -injective simple closed curve on S . Let $\pi: \tilde{S} \rightarrow S$ denote the universal cover of S . The preimage $\pi^{-1}(\gamma)$ is a collection of disjoint lines. Pick one of these lines, and call it l . Let H denote the stabilizer of l , so that $H \cong \mathbb{Z}$. Construct a tree T as follows. The vertices of T are the regions of $\tilde{S} - \pi^{-1}(\gamma)$. Whenever two regions are adjacent, attach an edge with an endpoint in each region. T is called the *dual graph* to $\pi^{-1}(\gamma)$ in \tilde{S} . Since γ does not intersect itself, T is in fact a tree. After choosing basepoints, G acts on T via deck transformations, with no fixed points or edge inversions. The stabilizer of an edge is isomorphic to \mathbb{Z} , so we have a splitting σ of G over \mathbb{Z} . If γ separates S into two components S' and S'' , then σ is an amalgamated free product $\pi_1(S') *_\mathbb{Z} \pi_1(S'')$ (see Figure 3.1). If γ does not separate S , then σ is an HNN extension $\pi_1(S - \gamma) *_\mathbb{Z}$ (see Figure 3.2).

Next, consider the curves γ_1 and γ_2 from Figures 3.1 and 3.2 simultaneously, as in Figure 3.3. We have associated splittings σ_1 and σ_2 of $\pi_1(S)$. Since γ_1 and γ_2 do not cross each other, the dual graph T to $\pi^{-1}(\gamma_1 \cup \gamma_2)$ is a tree. Moreover, T is a common refinement of the dual tree to $\pi^{-1}(\gamma_1)$ and the dual tree to $\pi^{-1}(\gamma_2)$. Hence trees for σ_1 and σ_2 have a common refinement. This is an example of compatible splittings (see Chapter 3.15).

What if we “poke a finger” out of γ_2 , as in Figure 3.4? This gives the same two splittings σ_1 and σ_2 of $\pi_1(S)$ as in the previous paragraph. However, from the way γ_1 and γ_2 are drawn, the dual graph to $\pi^{-1}(\gamma_1 \cup \gamma_2)$ is no longer a tree. In order to find a common refinement of the trees for σ_1 and σ_2 , it is helpful to first pull γ_1 and γ_2 tight to geodesics. When dealing with arbitrary splittings (not just those induced

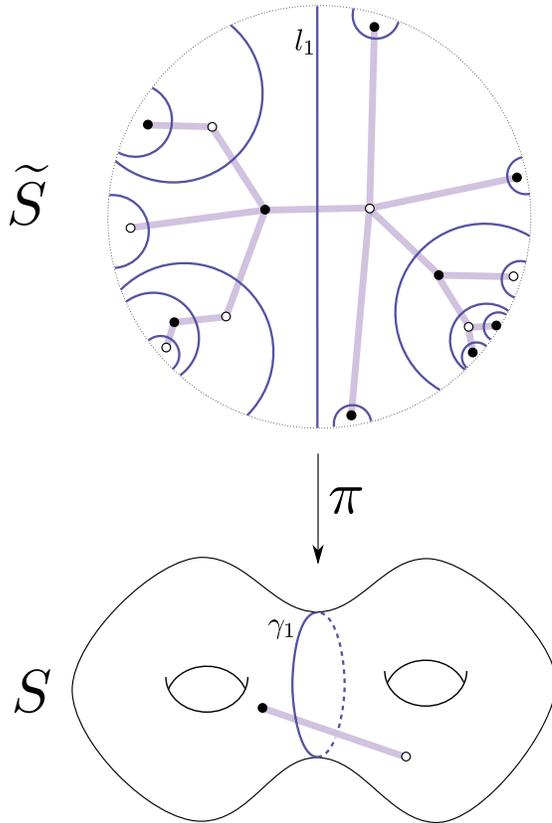


Figure 3.1: Schematic picture: $\pi^{-1}(\gamma_1)$, the union of l_1 and all its translates, is a collection of disjoint lines. The complement $\tilde{S} - \pi^{-1}(\gamma_1)$ has two types of components: those that project to the left of γ_1 , and those that project to the right of γ_1 . Correspondingly, the dual tree to $\pi^{-1}(\gamma_1)$ has two orbits of vertices. The action of the fundamental group of S on the tree gives an amalgamated free product.

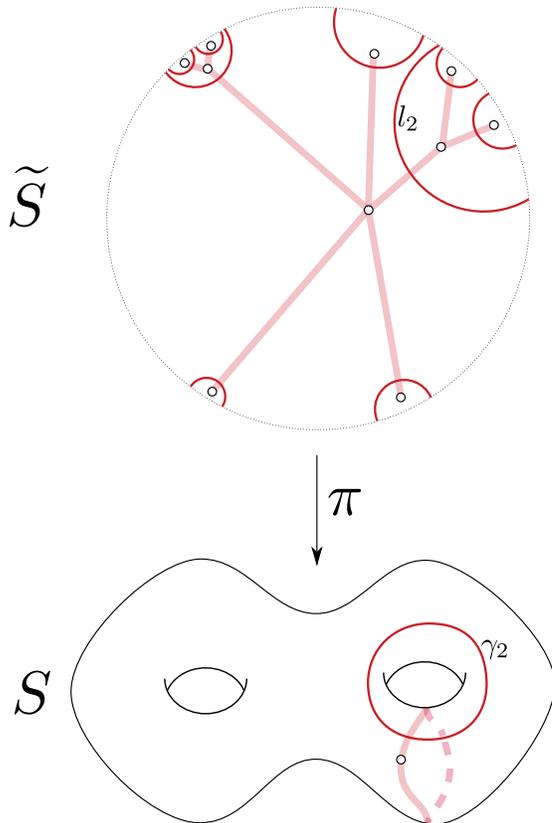


Figure 3.2: Schematic picture: $\pi^{-1}(\gamma_2)$, the union of l_2 and all its translates, is a collection of disjoint lines. $S - \gamma$ has only one region, so the dual tree to $\pi^{-1}(\gamma_2)$ has only one orbit of vertices. The action of the fundamental group of S on the tree gives an HNN extension.

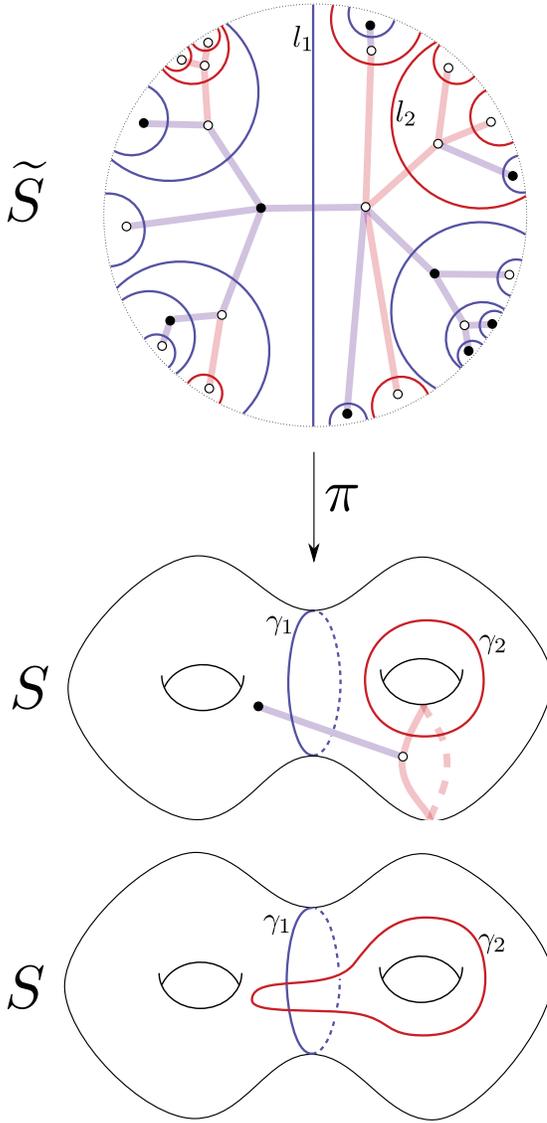


Figure 3.3: Schematic picture: $\pi^{-1}(\gamma_1 \cup \gamma_2)$, the union of l_1 , l_2 , and all their translates, is a collection of distinct lines, each projecting to either γ_1 or γ_2 . Correspondingly, the dual tree to $\pi^{-1}(\gamma_1 \cup \gamma_2)$ has two orbits of edges. The complement $S - \pi^{-1}(\gamma_1 \cup \gamma_2)$ has two types of components: those that project to the left of γ_1 , and those that project to the right of γ_1 . Correspondingly, the dual tree to $\pi^{-1}(\gamma_1 \cup \gamma_2)$ has two orbits of vertices. This tree is a compatibility tree for the splittings.

Figure 3.4: Take the previous example, but deform γ_2 slightly. The splittings induced by γ_1 and γ_2 still have intersection number zero, but in this example, the dual graph to $\pi^{-1}(\gamma_1 \cup \gamma_2)$ is not a tree. Since γ_1 and the new γ_2 have “inessential crossing,” they are not the best curves to use for the splittings.

by simple closed curves on surfaces), we’ll need some sort of algebraic tool to choose nice representatives for splittings.

3.14 Sandwiching

At times we will need to assume that either X crosses all translates of Y , or X can be sandwiched between two translates of Y or Y^* .

Definition 3.9 (Modified from [26]). Let $\{X_j | j \in J\}$ denote any collection of almost invariant subsets of G , where $\bigcup_{j \in J} \Sigma(X_j)$ is partially ordered by \leq .

- X_j is sandwiched by X_k if either there exist $A, B \in \Sigma(X_k)$ such that $A \leq X_j \leq B$,
or X_j crosses every element of $\Sigma(X_k)$.
- $\{X_j | j \in J\}$ satisfies sandwiching if for all $j, k \in J$, we have X_j is sandwiched by X_k .
- If X_j is an almost invariant set arising from a splitting σ_j ,
then we say $\{\sigma_j | j \in J\}$ satisfies sandwiching. Note that by Corollary 8.2, it does not matter which X_j we choose to represent σ_j .

“Most” collections of splittings satisfy sandwiching. In fact, if none of the X_j ’s yields a trivially ascending HNN extension, then $\{X_j | j \in J\}$ automatically satisfies sandwiching (see Corollary 7.3).

On the other hand, if a collection of pairwise non-isomorphic splittings satisfies sandwiching, and one of the splittings is a trivially ascending HNN extension, then that splitting must “cross everything.” More precisely, suppose Y is a K -almost invariant set arising from a trivially ascending HNN extension in the collection, and let X be a H -almost invariant set arising any splitting in the collection. I claim that, unless X is a translate of Y or Y^* , we must have X crosses Y . If X does not cross Y , then since the collection satisfies sandwiching, we can find translates A and B of Y such that $A \leq X \leq B$. Since all translates of Y are K -almost equal to Y , this implies that X is K -almost equal to Y . It now follows from Proposition 6.3 that X and Y arose from isomorphic splittings of G . As we assumed the collection of splittings to be pairwise non-isomorphic, X must be a translate of Y or Y^* .

3.15 Compatibility

We call two splittings of G “compatible” if their G -trees have a common refinement. Here is the formal definition, which works for an arbitrary number of splittings.

Definition 3.10. Let $\{\sigma_j | j \in J\}$ be any collection of splittings of G . A *compatibility tree* for $\{\sigma_j | j \in J\}$ is a G -tree T with a bijective correspondence between the edge orbits of T and the set J , such that for each $j \in J$, collapsing all edges except the σ_j -edges yields a tree for σ_j . We say $\{\sigma_j | j \in J\}$ is *compatible* if $\{\sigma_j | j \in J\}$ has a compatibility tree.

I now present two consequences of this definition. First of all, any splitting is compatible with (any splitting isomorphic to) itself. To see this, take a G -tree for the splitting, and subdivide each edge in two.

Secondly, if $\{\sigma_j | j \in J\}$ is compatible, then only finitely many of the σ_j 's can belong to any given isomorphism class. To see this, let T be a compatibility tree for $\{\sigma_j | j \in J\}$, and suppose that for some infinite subset J_0 of J , the splittings $\{\sigma_j | j \in J_0\}$ are pairwise isomorphic. Fix a vertex v in T , and for each $j \in J_0$, pick an edge e_j whose orbit corresponds to σ_j . Let $X_j := \{g \in G | e_j \text{ points away from } gv\}$. By Corollary 8.2, without loss of generality (after possibly replacing some of the e_j 's by a different representative from the same edge orbit), each X_j has the same stabilizer H , and $X_j \stackrel{H\text{-a}}{=} X_{j'}$, for all $j, j' \in J_0$. This implies that $\bigcup_{j \in J_0} e_j$ is an infinite path of valence two edges in T , and hence must be all of T . This forces each edge orbit to consist of only one edge, a contradiction to the fact that T is a refinement of G -trees for splittings. Hence a collection of compatible splittings can have only finitely many splittings in any given isomorphism class.

I prove that if $i(\sigma, \tau) = 0$ and σ and τ satisfy sandwiching, then σ and τ are compatible (see Theorem 7.5). The main idea of the proof is to note that if $i(\sigma, \tau) = 0$, then for all $A, B \in \Sigma$, we have A and B are almost nested, i.e. one of $A \leq B$, $A \leq B^*$, $A^* \leq B$, or $A^* \leq B^*$. Then apply Dunwoody's theorem (see Chapter 3.11).

Note that if two splittings are compatible, then we can find corresponding almost invariant sets that are nested (instead of just almost nested) as follows. Suppose that X is an H -almost invariant set arising from σ , that Y is a K -almost invariant set arising from τ , and that T is a compatibility tree for σ and τ . Since $i(\sigma, \tau) = 0$, we have for all $A, B \in \Sigma$, one of $A^{(*)} \leq B^{(*)}$. Apply Corollary 8.2 to T to get $X' \stackrel{H-a}{=} X$ and $Y' \stackrel{H-a}{=} Y$, such that for all $A', B' \in \Sigma(X') \cup \Sigma(Y')$, one of $A'^{(*)} \subset B'^{(*)}$.

3.16 Sandwiching is Necessary

If a G -tree has two edge orbits and no fixed points, then for any given edge, we can find two edges on either side of it belonging to the other edge orbit. This shows that if two splittings are compatible, then the splittings necessarily satisfy sandwiching.

Guirardel produced an example of two splittings with intersection number zero that do not satisfy sandwiching (and hence are not compatible). These are splittings of the free group on two generators, and over non-finitely generated subgroups. One of the splittings is a trivially ascending HNN extension. See [26] for the construction.

3.17 Turning Almost Inclusion Into Inclusion

What happens if we try to apply ideas from the “intersection number zero implies compatible” theorem to splittings having positive intersection number?

Let $\sigma_1, \sigma_2, \dots, \sigma_n$ be splittings of G collectively satisfying sandwiching, and such that no two of the splittings are isomorphic. Let X_j be an almost invariant set arising from σ_j . Using Definition 3.8, we have a partial order \leq on $\bigcup_{j=1}^n \Sigma(X_j)$, such that

if A is a subset of B , then $A \leq B$. It turns out that we can build a CAT(0) cubical complex, then use Theorem 8.12 to replace each X_j by another almost invariant set $X'_j \stackrel{H_j\text{-a}}{=} X_j$ such that $\bigcup_{j=1}^n \Sigma(X'_j)$ is in “very good position.”

Definition 3.11. Take Σ as in Definition 3.8. We say Σ is in *very good position* if for all $A, B \in \Sigma$, we have $A \subset B$ if, and only if, $A \leq B$.

Details for this construction are laid out in Chapter 8. The same CAT(0) cubical complex can also be used to construct an algebraic regular neighborhood of the σ_j 's (see Chapter 9).

3.18 Algebraic Regular Neighborhoods

The notion of “algebraic regular neighborhood” is a generalization of PL regular neighborhood, up to homotopy. Let $\gamma_1, \dots, \gamma_n$ be π_1 -injective simple closed curves on a closed surface S , yielding splittings $\sigma_1, \dots, \sigma_n$ of $\pi_1(S)$. Let N be a regular neighborhood of the γ_j 's. Assume each component of the boundary of N is π_1 -injective. We will construct a bipartite graph dual to the boundary of N . For each component of N , add a V_0 -vertex. For each component of $S - N$, add a V_1 -vertex. For each component of the boundary of N , add an edge connecting the corresponding V_0 - and V_1 -vertices. The preimage of Γ in \tilde{S} is a tree. Call this tree an algebraic regular neighborhood of $\{\sigma_1, \dots, \sigma_n\}$. Each V_0 -vertex orbit encloses (to be defined below) some of the σ_j 's. Each simple closed curve disjoint from all the γ_j 's can be homotoped to be disjoint from N , and hence is enclosed in a V_1 -vertex. See Figure 3.5 for a concrete example.

In [25], the authors defined an algebraic regular neighborhood of a family of almost invariant sets as a graph of groups, and they defined what it means for a vertex to enclose an almost invariant set. Here, to avoid confusion about base points, we define

algebraic regular neighborhood as a G -tree. Also, since this paper concerns almost invariant sets that arise from splittings, we define what it means for the orbit of a vertex of a G -tree to enclose a splitting.

Definition 3.12. Let T be a G -tree, V a vertex of T , and σ a splitting of G . *The orbit of V encloses σ* if T can be refined by inserting an edge at each vertex in the orbit of V , such that the new edges form a tree for σ . More precisely, there exists a G -tree T' and a G -orbit of edges in T' (call these edges σ edges) such that both of the following hold:

1. There exists a G -equivariant isomorphism

$$T' / \text{all non-}\sigma \text{ edges collapsed} \cong \text{tree for } \sigma.$$

2. There exists a G -equivariant isomorphism

$$T' / \text{all } \sigma \text{ edges collapsed} \cong T.$$

Given a collection $\{\sigma_j | j \in J\}$ of splittings, call σ_j an *isolated splitting* if it has intersection number zero with each other splitting σ_k . Given a G -tree, call a vertex *isolated* if the vertex has valence two in both the G -tree and the tree's quotient under the G -action. In an algebraic regular neighborhood, each isolated splitting in $\{\sigma_j | j \in J\}$ should be enclosed by the orbit of an isolated V_0 -vertex. Now we are ready to formally define an algebraic regular neighborhood.

Definition 3.13 (Reformulated from [25]). Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j \in J\}$ satisfies sandwiching. An *algebraic regular neighborhood of $\{\sigma_j | j \in J\}$* is a bipartite G -tree T (denote the two vertex colors by V_0 and V_1) satisfying the following five conditions:

1. Each σ_j is enclosed by some V_0 -vertex orbit in T , and each V_0 -vertex orbit encloses some σ_j .
2. If σ is a splitting of G over H , where σ is sandwiched by σ_j and $i(\sigma, \sigma_j) = 0$, for all $j \in J$, then σ is enclosed by some V_1 -vertex orbit in T .
3. T is a minimal G -tree.
4. There exists a bijection

$$f: \{j \in J \mid \sigma_j \text{ is isolated}\} \rightarrow G\text{-orbits of isolated } V_0\text{-vertices of } T$$

such that $f(j)$ encloses σ_j .

5. Every non-isolated V_0 -vertex orbit in T encloses some non-isolated σ_j .

In Chapter 9, I prove the existence of algebraic regular neighborhoods for any finite collection of splittings satisfying sandwiching. In Chapter 10, I prove uniqueness of algebraic regular neighborhoods for possibly infinite collections of splittings satisfying sandwiching.

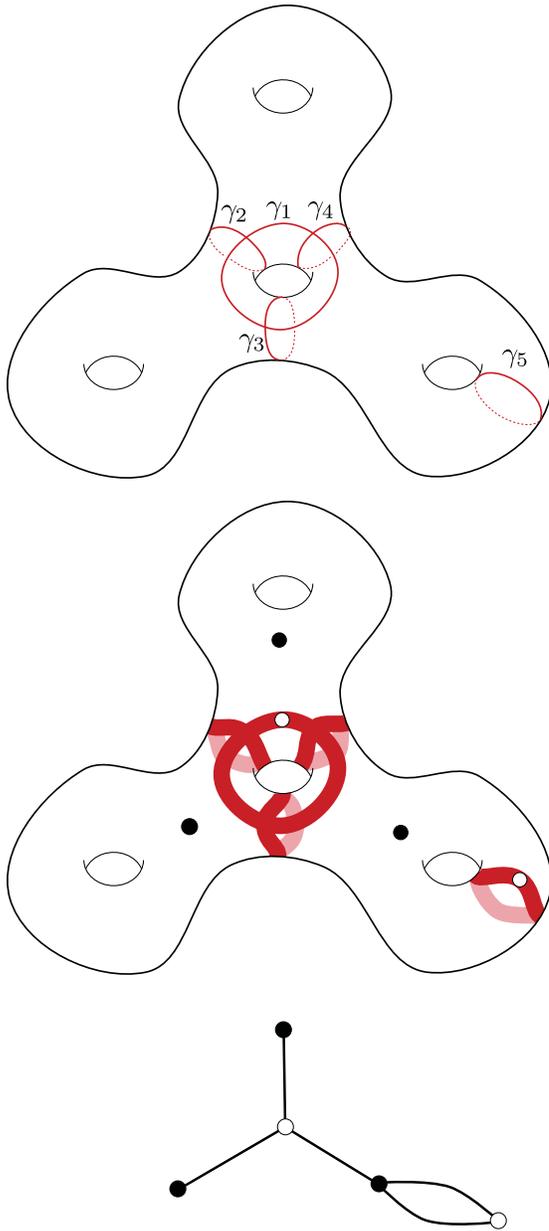


Figure 3.5: Let $\sigma_1, \sigma_2, \sigma_3, \sigma_4,$ and σ_5 denote the induced splittings of the fundamental group of S . Shown in bold is a PL regular neighborhood of $\{\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5\}$; call it N . Construct the dual graph to the boundary of N , making a V_0 -vertex for each component of N and a V_1 -vertex for each component of $S - N$. The resulting graph, call it Γ , is bipartite. Γ has two V_0 -vertices. One V_0 -vertex encloses $\sigma_1, \sigma_2, \sigma_3,$ and σ_4 . The other V_0 -vertex encloses σ_5 . The preimage of Γ in \tilde{S} is an algebraic regular neighborhood of $\{\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5\}$.

CHAPTER 4

Symmetry of Crossing

Let G be any group with subgroups H and K . Let X and Y be H - and K -almost invariant subsets of G , respectively. Recall Definition 3.5:

X crosses $Y \iff$ all four corners of the pair (X, Y) are K -infinite.

If G is finitely generated and neither X nor Y is trivial, an argument using coboundary in the Cayley graph for G shows that the relation “ X crosses Y ” is symmetric [23, Lemma 2.3]. For non-finitely generated G , this argument utterly fails, and so it seems plausible that crossing of almost invariant sets is not symmetric. However, below I prove that if X and Y come from splittings of G , then crossing is symmetric. Here is the key lemma.

Lemma 4.1. *Let G be any group with subgroups H and K . Suppose Y is a K -almost invariant set arising from a splitting of G over K . Further, suppose Y contains some nonempty subset X' that is stabilized by H (equivalently, Y contains at least one H -coset). Take any $g_0 \in Y^*$. Then*

$$Hg_0 \cap Y^* = (H \cap K)g_0 \cap Y^*$$

(so that $Hg_0 \cap Y^*$ is both H - and K -finite).

Proof. Clearly $(H \cap K)g_0 \cap Y^*$ is a subset of $Hg_0 \cap Y^*$. In the remaining part of the proof, we show that $Hg_0 \cap Y^*$ is a subset of $(H \cap K)g_0 \cap Y^*$.

Since Y comes from a splitting, there exists a G -tree T with an edge e and a vertex w , such that $K = \text{Stab}(e)$ and

$$Y = \{g \in G \mid e \text{ points away from } gw\}.$$

Measure the distance between two vertices in T by counting the number of edges in a simple path connecting them. Since $Hg_0 \cap Y^*$ is contained in a single H -coset, and since X is stabilized by H , we can choose $x \in X'$ such that the path $[xw, g_0w]$ realizes the minimum distance from $X'w$ to $(Hg_0 \cap Y^*)w$. Let D denote this distance. See Figure 4.1 for the basic picture.

Take any $h \in H$ such that $hg_0 \in Y^*$. We will show that $h \in K$. Since h stabilizes X' , multiplying the path $[xw, g_0w]$ on the left by h gives another path from $X'w$ to $(Hg_0 \cap Y^*)w$ of length D . As $X' \subset Y$, both paths must pass through e . Let d_1 denote the distance from xw to $e(0)$. Let d_2 denote the distance from $hxxw$ to $e(0)$. I claim that $d_1 = d_2$. If $d_1 < d_2$, then $[xw, hg_0w]$ would be a path from $X'w$ to $(Hg_0 \cap Y^*)w$ with length strictly less than D . Similarly, if $d_2 < d_1$, then $[hxxw, g_0w]$ would be a path from $X'w$ to $(Hg_0 \cap Y^*)w$ with length strictly less than D . Hence $d_1 = d_2$. It follows that $e = he$. As K is the stabilizer of e , this implies $h \in K$, as desired. This concludes the proof that $Hg_0 \cap Y^*$ is equal to $(H \cap K)g_0 \cap Y^*$. \square

We can use the lemma to prove symmetry of crossings for almost invariant sets that come from splittings.

Proposition 4.2. *Let G be any group with subgroups H and K . Suppose X is any nontrivial H -almost invariant set, and Y is a K -almost invariant set arising from a splitting of G over K . If X crosses Y , then Y crosses X .*

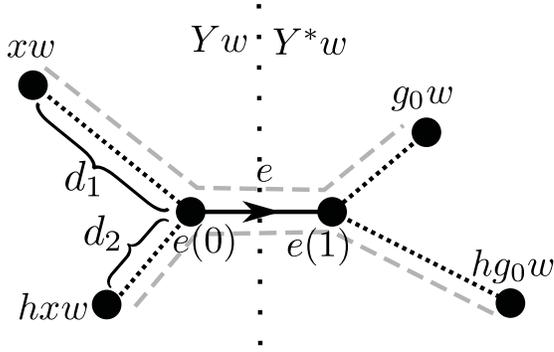


Figure 4.1: Proof of Lemma 4.1.

Proof. If Y does not cross X , then one of the corners of (X, Y) is H -finite. Without loss of generality (after possibly replacing X by X^* or Y by Y^*), $X \cap Y^*$ is H -finite. This means that we can choose finitely many $g_i \in X \cap Y^*$ such that $X \subset Y \cup Hg_1 \cup \dots \cup Hg_r$. Let $X' := X - \coprod_{i=1}^r Hg_i$, so that $X' \subset Y$. Since X is nontrivial, X' is nonempty. X' is also stabilized by H . As Y comes from a splitting, Lemma 4.1 proves that $Hg_i \cap Y^* = (H \cap K)g_i \cap Y^*$, for all i . Hence $X \cap Y^*$ is K -finite, so that X does not cross Y . \square

Corollary 4.3. *Intersection number of a pair of splittings (see Definition 3.10) is well-defined, even if the ambient group is not finitely generated.*

CHAPTER 5

Examples of Infinite Intersection Number

Scott and Swarup have shown that the intersection number of the two splittings of a finite group over finitely generated subgroups is finite [23, Lemma 2.7]. In the spirit of this paper, one might ask if we can eliminate one of the finite generation requirements. The answer is, definitively, “no.”

Example 5.1. This is an example of two splittings of a non-finitely generated group over the trivial group, where the intersection of the splittings is infinite. Let S be an infinite strip with countably many punctures:

$$S := \left[-\frac{1}{2}, \frac{1}{2}\right] \times \mathbb{R} - \{0\} \times \mathbb{Z}.$$

Take l_1 and l_2 as shown in Figure 5.1. Let l_1^+ be a regular neighborhood of the part of S lying above l_1 , and define l_1^- , l_2^+ , and l_2^- similarly. By Van Kampen’s theorem, we have the following two splittings of $G := \pi_1(S)$ over the trivial group.

$$\sigma : G \cong \pi_1(l_1^+) *_{\{1\}} \pi_1(l_1^-) = \mathbb{F}^{\mathbb{Z}} * \mathbb{F}^{\mathbb{Z}}$$

$$\tau : G \cong \pi_1(l_2^+) *_{\{1\}} \pi_1(l_2^-) = \mathbb{F}^{\mathbb{Z}} * \mathbb{F}^{\mathbb{Z}}$$

Here, the intersection number of σ and τ is visibly infinite.

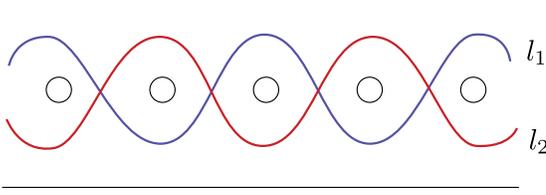


Figure 5.1: Two curves yielding a pair of splittings with infinite intersection number.

Example 5.2. We have a similar example exhibiting infinite self-intersection number for a $\{1\}$ -almost invariant set not that is not associated to a splitting. Take S and G as in the previous example. The curve in Figure 5.2 yields a $\{1\}$ -almost invariant subset of G with infinite self-intersection number.

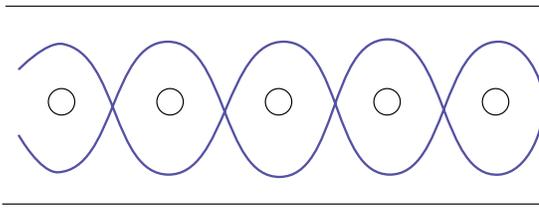


Figure 5.2: A curve yielding a $\{1\}$ -almost invariant set with infinite self-intersection number.

Example 5.3 (Guirardel). In [8, Lemma 8.4] Guirardel gave an example of two splittings of \mathbb{F}_3 with infinite intersection number. For completeness, we include the example here. Let $G = \langle a, b, c \rangle$. Let σ be the HNN extension

$$\sigma : G \cong \langle b, c \rangle *_{\{1\}}$$

where the loop represents conjugation by a . Let K denote the non-finitely generated subgroup $\langle b^i a b^{-i} \mid i \in \mathbb{Z} \rangle$, and let τ be the splitting

$$\tau : G \cong \langle a, b \rangle *_{K} \langle K, c \rangle$$

Use Theorem 8.7 to construct the CAT(0) cubical complex for σ and τ . Let \mathcal{K} be a hyperplane associated to τ , such that $K = \text{Stab}(\mathcal{K})$. Each edge e in \mathcal{K} belongs to a unique square in the cubical complex, and hence meets transversally a unique hyperplane \mathcal{H}_e associated to σ . Let T_0 denote the dual tree to the hyperplanes $\{\mathcal{H}_e \mid e \in \mathcal{K}\}$. Since σ is a splitting over the trivial group, H acts freely on T_0 . As H

is not finitely generated, it follows that this action is not cocompact. The intersection number of σ and τ is equal to the number of edges in the quotient of T_0 by H . Hence this intersection number is infinite.

CHAPTER 6

Almost Inclusion

Let $\{X_j | j \in J\}$ be a collection of H_j -almost invariant subsets of a group G , arising from pairwise non-isomorphic splittings σ_j over H_j . Let $\Sigma := \{gX_j, gX_j^* | g \in G, j \in J\}$. Recall Definitions 3.7 and 3.8. In this chapter, we prove that Σ is in good position (Corollary 6.4), and hence \leq defines a partial order on Σ (Corollary 6.5).

The following lemma, proved in a preprint by Scott and Swarup, shows that if an H -almost invariant set is H -almost equal to a K -almost invariant set, then H and K are commensurable.

Lemma 6.1. *Let G be any group with a nontrivial H -almost invariant set X and a nontrivial K -almost invariant set Y . If $X \stackrel{H\text{-a}}{=} Y$, then H and K are commensurable subgroups of G .*

Proof. $X \stackrel{H\text{-a}}{=} Y$ immediately implies that $Xg \stackrel{H\text{-a}}{=} Yg$, for all $g \in G$. As X is H -almost invariant, we have $Xg \stackrel{H\text{-a}}{=} X$, and hence $Yg \stackrel{H\text{-a}}{=} Y$, for all $g \in G$. As Y is K -almost invariant, K stabilizes Y , so each of Y and Y^* is a union of cosets Kg of K in G .

Since Y is nontrivial, we can choose $u, v \in G$ such that $Ku \subset Y$ (equivalently, $K \subset Yu^{-1}$) and $Kv \subset Y^*$. Recall that by the preceding paragraph, $Y(u^{-1}v) \stackrel{H\text{-a}}{=} Y$. Note that Kv lies in the symmetric difference of $Y(u^{-1}v)$ and Y , so Kv must be H -

finite. Hence K is also H -finite. We can write $K \subset \coprod_{i=1}^r Hg_i$, where r is minimal. We have $K = \coprod_{i=1}^r (K \cap H)g_i$. As K is the union of finitely many $(K \cap H)$ -cosets, it follows that $[K : K \cap H] < \infty$.

A similar argument shows that $K \cap H$ is finite index in H . Hence H and K are commensurable subgroups of G . \square

Next we show that if X and Y arise from splittings, then their stabilizers are actually equal.

Lemma 6.2 (Modified from Lemma 2.2 of [24]). *Let X be an H -almost invariant subset arising from a splitting of G over H , and let Y be a K -almost invariant set arising from a splitting of G over K .*

1. *If two corners of the pair (X, Y) are H -finite, then $H = K$.*
2. *If two corners of the pair (X, gX) are H -finite, then g normalizes H .*

Proof. To prove the first part of the lemma, suppose two corners of the pair (X, Y) are H -finite. Without loss of generality, $X \cap Y^*$ and $X^* \cap Y$ are H -finite (if not, replace X by X^*), so that X and Y are H -almost equal.

As Y comes from a splitting, it follows that hY and Y are nested, for all $h \in G$. We will now show that $H \subset K$. Let $h \in H$. If $hY \subset Y^*$ (or $hY^* \subset Y$), then $X \stackrel{H\text{-a}}{=} X^*$, a contradiction to G being H -infinite. If $hY \subset Y$ but $hY \neq Y$, then we get an infinite chain of inclusions

$$\dots \subset h^n Y \dots \subset hY \subset Y.$$

As H and K are commensurable, some power of h lies in K , so that $h^n Y = Y$ for some n . This implies $hY = Y$, so that $h \in K$. Similarly, if $Y \subset hY$, we must also have $Y = hY$ and $h \in K$.

A similar argument shows that $K \subset H$. Hence $H = K$.

To prove the second part of the lemma, apply the first part using $Y := gX$. The first part of the lemma gives $H = K = gHg^{-1}$, so that g normalizes H . \square

Now we show that if X and Y are H - and K -almost invariant sets arising from non-isomorphic splittings of G , then it is impossible to have $X \stackrel{H\text{-a}}{=} Y$.

Proposition 6.3 (Modified from Lemma 2.3 of [24]). *Let X and Y be H - and K -almost invariant sets arising from splittings σ and τ of G over subgroups H and K , respectively. If two corners of the pair (X, Y) are small, then σ and τ are isomorphic splittings.*

Further, at least one of the following holds (after possibly replacing X by X^):*

1. $X \mapsto Y$ induces a G -equivariant, order-preserving isomorphism from $(\Sigma(X), \subset)$ to $(\Sigma(Y), \subset)$; or
2. The two splittings are of the form $G = A *_H B$, where H has index 2 in A , and there exists $a \in A$ such that $X \mapsto aY$ induces a G -equivariant, order-preserving isomorphism from $\Sigma(X)$ to $\Sigma(Y)$.

Proof. By replacing X by X^* if necessary, without loss of generality, $X \cap Y^*$ and $X^* \cap Y$ are H -finite, i.e. $X \stackrel{H\text{-a}}{=} Y$. By Lemma 6.2, we have $H = K$.

A corner of (X, gX) is small if, and only if, the corresponding corner of (Y, gY) is small. If, for all $g \in G - H$, only one corner of (X, gX) is small (and hence empty), then the corresponding corner of (Y, gY) must also be empty. Then $X \mapsto Y$ induces a G -equivariant, order-preserving isomorphism from $\Sigma(X)$ to $\Sigma(Y)$, and hence the splittings are isomorphic by Dunwoody's theorem (see Chapter 3.11 and [3]).

If there exists $g \in G - H$ such that two corners of (X, gX) are small, then the trees for σ and τ must each have some vertices of valence two. There are two cases:

1. σ is a trivially ascending HNN extension, $G \cong H *_H$. Then T_σ and T_τ are lines, H and K are normal in G , and $G = \langle H, t \rangle$ for some $t \in G$. Thus $X \mapsto Y$ induces a G -equivariant, order-preserving isomorphism from $\Sigma(X)$ to $\Sigma(Y)$.
2. σ is an amalgamated free product of the form $G \cong A *_H B$, where H has index 2 in A . We can write $A = \langle H, a \rangle$. Then $X \stackrel{H-a}{=} aX^*$ and $Y \stackrel{H-a}{=} aY^*$. If the corresponding corners of (X, aX^*) and (Y, aY^*) are empty, then $X \mapsto Y$ induces a G -equivariant, order-preserving isomorphism from $\Sigma(X)$ to $\Sigma(Y)$. Otherwise, the isomorphism comes from assigning $X \mapsto aY$.

□

Corollary 6.4. *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Then $\Sigma := \bigcup_{j \in J} \Sigma(X_j)$ is in good position.*

Proof. If there exists $g \in G$ and distinct $j, k \in J$ such that two corners of the pair (X_j, gX_k) are small, then σ_j and σ_k are isomorphic splittings (by Proposition 6.3), a contradiction to the hypotheses. □

Since Σ is in good position, we can define a partial order on Σ as follows.

Corollary 6.5. *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Let $\Sigma := \bigcup_{\sigma \in J} \Sigma(X_j)$. Define a binary relation \leq on Σ by*

$$A \leq B \Leftrightarrow A \cap B^* \text{ is empty or the only small corner of the pair } (A, B).$$

Then \leq is a partial order on Σ .

Note that “small” means “ $\text{Stab}(A)$ -finite” or equivalently “ $\text{Stab}(B)$ -finite” (see Proposition 4.2).

Proof. Reflexivity is obvious. We need to show antisymmetry and transitivity.

To show antisymmetry, suppose $A \leq B$ and $B \leq A$. Then both $A \cap B^*$ and $B \cap A^*$ are small corners of the pair (A, B) . Since two corners are small, the first inequality now implies $A \cap B^*$ is empty, while the second implies $B \cap A^*$ is empty. Hence $A = B$. Thus \leq satisfies antisymmetry.

To show transitivity, suppose $A \leq B$ and $B \leq C$, where A , B , and C are all distinct. We need to show that $A \leq C$. Since $B \leq C$, we can subtract finitely many $\text{Stab}(B)$ -cosets from B to obtain $B' \subset C$. Since $A \leq B$ and B is $\text{Stab}(B)$ -almost equal to B' , we have $A \cap B'^*$ is $\text{Stab}(B)$ -finite. By Lemma 4.1, since A arises from a splitting, $A \cap B'^*$ is also $\text{Stab}(A)$ -finite. Hence we can subtract finitely many $\text{Stab}(A)$ -cosets from A to obtain $A' \subset B'$. It follows that $A' \subset C$. Since A is $\text{Stab}(A)$ -almost equal to A' and since $A' \subset C$, we have $A \cap C^*$ is a small corner of the pair (A, C) .

Thus the only way we could possibly fail to have $A \leq C$ is if another corner were small. If two corners of the pair (A, C) are small, then Proposition 6.3 proves that A , B and C all must have come from isomorphic splittings of G . Since we assumed no two distinct j 's have isomorphic σ_j 's, it follows that A , B and C are all translates of X_j or X_j^* , for the same j . So we must have $A \subset B \subset C$. This completes the proof that \leq satisfies transitivity. \square

Now that we've put a partial order \leq on Σ , we show that the partial order is unique.

Corollary 6.6 (Uniqueness of the partial order). *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. Suppose that $\{\sigma'_j | j \in J\}$ is another collection of splittings of G , where $\sigma_j \cong \sigma'_j$, for all $j \in J$. For each j , let X_j be an H_j -almost invariant set arising from σ_j , and let X'_j be an H'_j -almost*

invariant set arising from σ'_j . Let $\Sigma := \bigcup_{j \in J} \Sigma(X_j)$, and let $\Sigma' := \bigcup_{j \in J} \Sigma(X'_j)$. Then there exists a G -equivariant, order-preserving isomorphism from (Σ, \leq) to (Σ', \leq) .

Proof. By Proposition 6.3, for all j , there exist $g_j \in G$ such that $X_j \mapsto g_j X'_j$ or $X_j^* \mapsto g_j X'_j$ induces a G -equivariant, order-preserving isomorphism from $\Sigma(X_j)$ to $\Sigma(X'_j)$. Together, these induce a G -equivariant isomorphism from Σ to Σ' . We need to show that this isomorphism is order-preserving. As no two of the σ_j 's are isomorphic, whenever $A \in \Sigma(X_j)$ and $B \in \Sigma(X_k)$ ($k \neq j$), at most one corner of (A, B) is small. If no corner of (A, B) is small, then no corner of (A', B') is small. If exactly one corner of (A, B) is small, then the same corner of (A', B') must be the only small corner of (A', B') . Hence $\Sigma \rightarrow \Sigma'$ is order-preserving. \square

Next we spell out this uniqueness result in the case when the splittings happen to be compatible. In this case, we allow some of the σ_j 's to be isomorphic to each other.

Corollary 6.7 (Uniqueness of compatibility trees). *Let G be any group with a finite collection $\{\sigma_j | j \in J\}$ of splittings. Suppose $\{\sigma_j | j \in J\}$ is compatible, and let T and T' be compatibility trees. Then there exists a G -equivariant isomorphism from T to T' .*

Proof. First, we prove the result in the case where no two distinct j 's have isomorphic splittings. Fix a vertex v in T . For each j , pick a σ_j -edge e_j in T , and define a subset X_j of G by

$$X_j := \{g \in G | e \text{ points away from } gv\}.$$

Fix a vertex v' in T' . For each j , pick a σ_j -edge e'_j in T' whose stabilizer is the same as $\text{Stab}(e_j)$, such that

$$X_j \stackrel{H_j\text{-a}}{=} \{g \in G | e'_j \text{ points away from } gv'\}$$

(we can do this by Corollary 8.2). Let X'_j denote the set $\{g \in G \mid e'_j \text{ points away from } gv'\}$. Apply Corollary 6.6 to get a G -equivariant, order preserving isomorphism from $\bigcup_{j \in J} \Sigma(X_j)$ to $\bigcup_{j \in J} \Sigma(X'_j)$. Dunwoody's theorem (see Chapter 3.11) now gives a G -equivariant isomorphism from T to T' .

Second, we prove the result in the case where $\{\sigma_j \mid j \in J\}$ possibly has duplicate splittings. For each isomorphism class $\{\sigma_j \mid j \in I\}$ of splittings, discard all but one representative; call it σ_I . Note that the edge in T (or T') corresponding to X_I must be contained in an interval of $|I|$ edges, one for each σ_j in the isomorphism class, where the interior vertices of the interval each have valence two. Collapse the edge orbits of T and T' corresponding to the discarded splittings. To recover an isomorphism from T to T' , for each isomorphism class I , subdivide each σ_I edge in the collapsed T and the collapsed T' into an interval of $|I|$ edges. \square

CHAPTER 7

Compatibility and Intersection Number Zero

Take any finite collection of non-isomorphic splittings of G satisfying sandwiching (see Definition 3.9). Here we show that if the splittings have pairwise intersection number zero, then the splittings are compatible (this is Theorem 7.5). This is a special case of very good position, when the intersection number of each pair of splittings is zero.

The sandwiching assumption is necessary; see Chapter 3.16. For more intuition about sandwiching, we begin by proving that sandwiching is automatic if none of the splittings is a trivially ascending HNN extension (see Chapter 3.8). The key fact used is that if X arises from a splitting that is not trivially ascending HNN, then all four types of nesting occur between X and its translates:

Lemma 7.1. *Let σ be a splitting of G over H , where σ is not a trivially ascending HNN extension. Let X be an H -almost invariant set arising from σ . Then, by varying g , all four of $gX^{(*)} \subset X^{(*)}$ occur.*

This result is a strengthening of [25, Lemma 5.5], which assumes that σ is not any ascending HNN extension.

Proof. Since X arises from a splitting, there is a G -tree T with an edge e and a

vertex w , and exactly one orbit of edges, such that

$$X = \{g \in G \mid e \text{ points away from } gw\}.$$

It suffices to show that there exist translates of e such that $g_1e < g_2e$ and $g_3e < g_4\bar{e}$.

There are two cases:

1. T is a line, so since σ is not trivially ascending HNN, σ must have the form

$G \cong A *_H B$, where $|A : H| = |B : H| = 2$. To get $g_1e < g_2e$, take two translates of e separated by 1 edge. To get $g_3e < g_4\bar{e}$, take two adjacent translates of e .

2. T has branching, hence there exist three distinct translates of e such that the geodesics between any two pair of them all meet at exactly one vertex, and that either two of the translates point toward the vertex and one points away, or vice-versa. See Figure 7.1. To get $g_1e \leq g_2e$, take two of these translates of e where one is pointing toward the vertex and the other away. To get $g_3e \leq g_4\bar{e}$, take two translates pointing toward (or two pointing away from) the vertex.

□

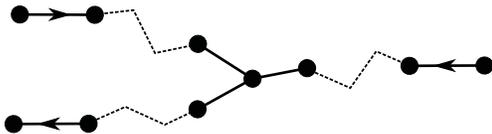


Figure 7.1: If a G -tree T has branching and exactly one edge orbit, then for any edge, we can find three of its translates such that either two point toward each other and the other one points away, or vice-versa.

If X and Y are almost invariant sets arising from splittings of G , where neither splitting is a trivially ascending HNN extension, then either X crosses all translates of Y , or all four types of almost nesting occur between X and translates of Y :

Lemma 7.2. *Let σ and τ be splittings of G over H and K , respectively, where neither σ nor τ is a trivially ascending HNN extension. Let X and Y be almost invariant*

sets arising from σ and τ , respectively. Suppose that there exists $g_0 \in G$ such that X and g_0Y do not cross. Then, by varying $g \in G$, all four of $X^{(*)} \leq gY^{(*)}$ occur.

Proof. Without loss of generality (after possibly replacing X by X^* or Y by Y^*), $X \leq g_0Y$. Obtain each of the four cases as follows:

1. $X \leq g_0Y$ is already given.
2. To show there exists g with $X \leq gY^*$, apply Lemma 7.1 to get $g_0Y \subset g_1Y^*$, so that $X \leq g_0Y \subset g_1Y^*$.
3. To show there exists g with $X^* \leq gY^*$, apply Lemma 7.1 to get $g_2X^* \subset X$. Now $g_2X^* \subset X \leq g_0Y \subset g_1Y^*$, so that $X^* \leq g_2^{-1}g_1Y^*$.
4. To show there exists g with $X^* \leq gY$, apply Lemma 7.1 to get $g_0Y \subset g_3Y$. Now $g_2X^* \subset X \leq g_0Y \subset g_3Y$, so that $X^* \leq g_2^{-1}g_3Y$.

□

Note that if we assume that X and g_0Y are nested (instead of almost nested), then the same proof shows that all four inclusions $X^{(*)} \subset gY^{(*)}$ occur.

Corollary 7.3. *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of splittings, where no σ_j is a trivially ascending HNN extension. Then $\{\sigma_j | j \in J\}$ satisfies sandwiching.*

Most of the results in the rest of the paper will require the sandwiching assumption. The key reason we need sandwiching is to get interval finiteness:

Proposition 7.4. *Let σ_j be a splitting of G over H_j , and assume $\{\sigma_1, \dots, \sigma_n\}$ satisfies sandwiching. Let X_j be an H_j -almost invariant set arising from σ_j . Let $\Sigma = \{gX_j, gX_j^* | g \in G, j = 1, \dots, n\}$. Then for all $A, B \in \Sigma$, there are only finitely many $C \in \Sigma$ such that $A \leq C \leq B$.*

Proof. Fix $A, B \in \Sigma$. If $A \not\leq B$, then there is no C such that $A \leq C \leq B$; so assume $A \leq B$. Since $\bigcup_{j \in J} \Sigma(X_j)$ satisfies sandwiching, for each $j \in J$ we can choose $A_j, B_j \in \Sigma(X_j)$ such that

$$A_j \leq A \leq B \leq B_j.$$

If $C \in \Sigma(X_j)$ and $A \leq C \leq B$, then $A_j \leq C \leq B_j$. But since the X_j 's arise from splittings, for each j there are only finitely many such C . As we are only considering finitely many splittings, there are only finitely many $C \in \Sigma$ satisfying $A \leq C \leq B$. \square

Theorem 7.5. *Let σ_j be a splitting of G over H_j and assume $\{\sigma_1, \dots, \sigma_n\}$ satisfies sandwiching. If $i(\sigma_j, \sigma_k) = 0$, for all j and k , then the splittings $\{\sigma_1, \dots, \sigma_n\}$ are compatible.*

Proof. First we prove the theorem for the case when no σ_j is isomorphic to any other.

Let X_j be an H_j -almost invariant set arising from σ_j . By Corollary 6.4, \leq is a partial order on $\Sigma = \{gX_j, gX_j^* | g \in G, j = 1, \dots, n\}$. We can see that the four conditions of Dunwoody's theorem (see Chapter 3.11) are satisfied:

1. For all $A, B \in \Sigma$, if $A \leq B$, then $B^* \leq A^*$.
2. For all $A, B \in \Sigma$ with $A \leq B$, there are only finitely many $C \in \Sigma$ with $A \leq C \leq B$ (see Proposition 7.4).
3. For all $A, B \in \Sigma$, at least one of $A^{(*)} \leq B^{(*)}$ (because $i(\sigma_j, \sigma_k) = 0$, for all j and k).
4. We cannot have simultaneously $A \leq B$ and $A \leq B^*$.

Construct Dunwoody's tree, T_Σ , with edge set Σ . Each edge is a σ_j -edge for unique

j . For all j , we have a G -equivariant isomorphism $T_\Sigma/(\text{all but } j\text{-edges collapsed}) \rightarrow T_{\sigma_j}$. Hence T_Σ is a compatibility tree for $\{\sigma_1, \dots, \sigma_n\}$.

Second, we prove the theorem in the case when we possibly have duplicate splittings. Discard all but one splitting from each isomorphism class. Apply the above procedure. Then subdivide the resulting tree, as in the proof of Corollary 6.7. \square

CHAPTER 8

CAT(0) Cubical Complexes and Positive Intersection Number

A *cubical complex* C is a CW-complex whose cells are standard Euclidean cubes of varying dimensions, such that the intersection of any two cells is either empty or a common face of both. C is called a *CAT(0) cubical complex* if, in addition, C is simply connected, and the link of any vertex (i.e. 0-cube) is a flag complex. Another word for “CAT(0) cubical complex” is *cubing*.

In this chapter, we start with any finite collection of pairwise non-isomorphic splittings of any group G , and construct a CAT(0) cubical complex. G acts naturally on the complex, and each hyperplane orbit will correspond to one of the splittings. Furthermore, hyperplanes cross precisely when their associated splittings cross. Essentially, we are showing how to make Niblo–Sageev–Scott–Swarup’s “minimal cubing” construction from [18] work without requiring G or the subgroups over which G splits to be finitely generated. Their “minimal cubing” construction, in turn, was a generalization of Sageev’s cubing construction in [21]. For applications of the cubing construction, see Theorem 8.12 and Chapter 9.

If the splittings are induced by simple closed curves on a surface as in Figure 8.1, then the dual graph to the preimage of the curves in the universal cover of the surface is part of the 1-skeleton of the cubing. However, in general the 1-skeleton of the

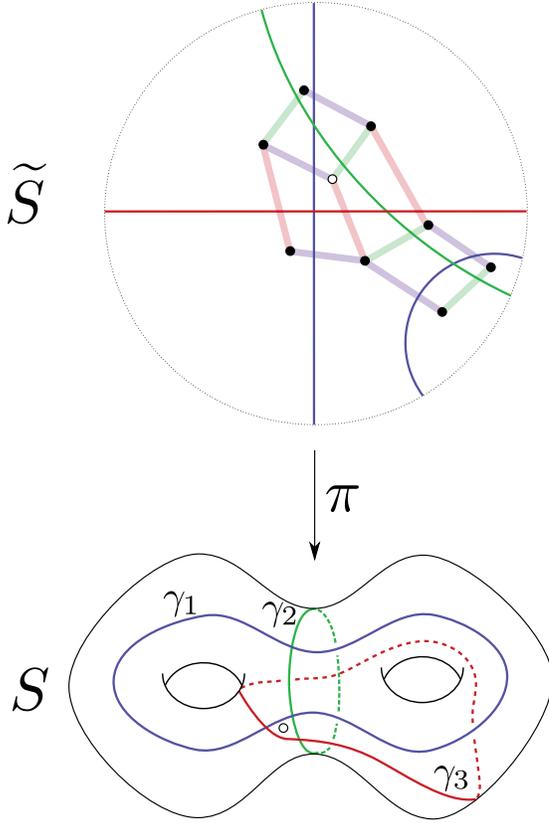


Figure 8.1: Since the curves γ_1 , γ_2 , and γ_3 cross pairwise, we can find three lines in $\pi^{-1}(\gamma_1 \cup \gamma_2 \cup \gamma_3)$ that cross pairwise. Shown in \tilde{S} is a part of the dual graph to $\pi^{-1}(\gamma_1 \cup \gamma_2 \cup \gamma_3)$. We see the 1-skeleton of half of a 3-cube, with an attached 2-cube. The 1-skeleton of the other half of the 3-cube (and, in fact, the 3-cube itself) is present in the cubing, but not in this dual graph.

cubing contains vertices that are not in this dual graph. See Figure 8.1.

We will briefly review all the basic constructions. For more details, see [18] and [21, Sections 2 and 3].

8.1 Producing Almost Invariant Sets from a CAT(0) cubical complex

In [25], Scott and Swarup showed how to produce an almost invariant set from a G -tree. Then in [18], Niblo–Sageev–Scott–Swarup generalized this construction by producing an almost invariant set from any CAT(0) cubical complex on which G acts. We include the formal statement and proof of this result below. Note that a tree is precisely a 1-dimensional CAT(0) cubical complex, and hyperplanes in a tree are midpoints of edges.

Lemma 8.1 (Lemma 1.17 from [18]). *Let G be any group acting on a cubing C . Let \mathcal{H} be a hyperplane in C with stabilizer H , and suppose that H preserves each of*

\mathcal{H}^+ and \mathcal{H}^- . Then for any vertex v , the set $X_v := \{g \in G \mid gv \in \mathcal{H}^+\}$ is H -almost invariant. Moreover, for any vertices v and w , the set X_v is H -almost equal to X_w .

In [18], the authors assume G is finitely generated, but their proof does not actually use that assumption.

Proof. First, we show that X is H -almost invariant. Clearly $hX_v = X_v$, for all $h \in H$. We also need $X_v a$ is H -almost equal to X_v , for all $a \in G$. We have:

$$\begin{aligned} X_v &= \{g \in G \mid gv \in \mathcal{H}^+\}, \text{ so that} \\ X_v a &= \{ga \in G \mid gv \in \mathcal{H}^+\} \\ &= \{g \in G \mid ga^{-1}v \in \mathcal{H}^+\} \end{aligned}$$

To show the symmetric difference of $X_v a$ and X_v is H -finite, first we consider one half of the symmetric difference:

$$\begin{aligned} X_v - X_v a &= \{g \in G \mid gv \in \mathcal{H}^+ \text{ and } ga^{-1}v \notin \mathcal{H}^+\} \\ &= \{g \in G \mid \mathcal{H} \text{ separates } gv \text{ from } ga^{-1}v\} \\ &= \{g \in G \mid g^{-1}\mathcal{H} \text{ separates } v \text{ from } a^{-1}v\} \end{aligned}$$

There are only finitely many (say, m) hyperplanes separating v from $a^{-1}v$. If $g, g' \in G$ with $g^{-1}\mathcal{H} = g'^{-1}\mathcal{H}$, then $g'g^{-1} \in \text{Stab}(\mathcal{H}) = H$, and so Hg' and Hg^{-1} are actually the same coset. We conclude that $X_v - X_v a$ is contained in at most m cosets Hg . Similarly, $X_v a - X_v$ is H -finite. Hence X_v is H -almost invariant.

Second, let v and w be vertices of C . We need to show that X_v is H -almost equal to X_w . We have:

$$\begin{aligned} g \in X_v - X_w &\iff gv \in \mathcal{H}^+ \text{ and } gw \notin \mathcal{H}^+ \\ &\iff g^{-1}\mathcal{H} \text{ separates } v \text{ from } w \end{aligned}$$

As in the argument above, the set of all such g is H -finite. Similarly, $X_w - X_v$ is H -finite. Hence X_v is H -almost equal to X_w . \square

Corollary 8.2. *Let G be any group and T a G -tree. Let e be an edge in T with stabilizer H . Then for any vertex v , the set $X_v := \{g \in G \mid e \text{ points away from } gv\}$ is H -almost invariant. Moreover, for any vertices v and w , the set X_v is H -almost equal to X_w .*

8.2 Ultrafilters

A *partially ordered set with complementation*, or *pocset*, is a partially ordered set (Σ, \leq) , equipped with a free involution $*$ on Σ behaving like complementation, i.e. $A \leq B$ implies $B^* \leq A^*$. This terminology was introduced by Sageev and Roller.

Definition 8.3. Let (Σ, \leq) be a pocset. An *ultrafilter* on (Σ, \leq) is a subset V of the power set of Σ such that both of the following conditions are satisfied:

1. For all $A \in \Sigma$, either $A \in V$ or $A^* \in V$ (but not both), and
2. If $A \in V$ and $A \leq B$, then $B \in V$.

We say an ultrafilter V satisfies the *descending chain condition (DCC)* if every chain $A_1 \geq A_2 \geq \dots$ stabilizes after finitely many steps.

Note that if V is an ultrafilter on Σ , then for any $g \in G$, the translate $gV := \{gA \mid A \in V\}$ is also an ultrafilter on Σ . Also note that $V - \{A\} \cup \{A^*\}$ is an ultrafilter if, and only if, A is a minimal element of (V, \leq) .

8.3 Sageev's Cubing

In [21], Sageev constructed a cubing C from a finite collection $\{X_j \mid j = 1, \dots, n\}$ of H_j -almost invariant subsets of a group G , using the partial order of inclusion. We will now briefly review this construction. Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. The vertices of Sageev's

cubing are a subset of all ultrafilters on (Σ, \subset) . Let C' be the complex with a vertex for each ultrafilter on (Σ, \subset) , and an edge connecting each pair of ultrafilters that differ by exactly one complementary pair (A, A^*) . If V is a vertex and $V \cup \{A^*\} - \{A\}$ is also a vertex, we say the (directed) edge from V to $V \cup \{A^*\} - \{A\}$ *exits* A . See below for the definition of “basic vertex.” Define the one-skeleton of C to be the connected component of C' containing all the basic vertices. We define higher skeleta of C inductively: whenever you see the boundary of an n -cube, attach an n -cube. Sageev showed that C is a CAT(0) cubical complex [21]. This is Sageev’s cubing.

Definition 8.4 (basic vertex). Let G be a group with a finite collection $\{X_j | j = 1, \dots, n\}$ of H_j -almost invariant subsets. Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. Let g be any element of G . Define V_g as follows:

$$V_g := \{A \in \Sigma | g \in A\}.$$

We call V_g a *basic vertex*. Some authors may refer to basic vertices as basic ultrafilters, principal vertices, or principal ultrafilters.

The following two lemmas prove that the vertices of Sageev’s cubing can be characterized as ultrafilters on (Σ, \subset) satisfying DCC.

Lemma 8.5. *Let G be a group with any finite collection $\{X_j | j = 1, \dots, n\}$ of H_j -almost invariant subsets. Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. For each $g \in G$, the basic vertex V_g is an ultrafilter on (Σ, \subset) and satisfies DCC.*

Proof. Fix $g \in G$. We first show that V_g satisfies conditions 1 and 2 of Definition 8.3.

1. Let $A, B \in \Sigma$ be arbitrary. Either $g \in A$ or $g \in A^*$, so either $A \in V_g$ or $A^* \in V_g$ (but not both).
2. If $A \in V_g$ and $A \subset B$, then $g \in B$, so $B \in V_g$.

Hence V_g is an ultrafilter on (Σ, \subset) .

To show V_g satisfies DCC, take a descending chain $B_1 \supset B_2 \supset \dots$ in V_g . If the B_k are not all equal to begin with, then without loss of generality (after passing to a subsequence), $B_1 - B_2$ is nonempty, and there exists some fixed j such that $B_k \in \Sigma(X_j)$, for all k . Fix $g_0 \in B_1 - B_2$. We claim (as proved in [21, Lemma 3.4]) that

$$\mathcal{B} := \{B \in \Sigma(X_j) \mid g \in B \text{ and } g_0 \notin B, \text{ or } g \notin B \text{ and } g_0 \in B\}$$

is finite. Assuming the claim, the chain must stabilize, as each element of the chain (except for B_1) is an element of \mathcal{B} .

To prove the claim, first note that since X_j is H_j -almost invariant, we have

$$X_j g^{-1} \stackrel{H_j\text{-a}}{=} X_j g_0^{-1}.$$

Pick $g_1^j, \dots, g_{r_j}^j$ such that the symmetric difference of $X_j g^{-1}$ and $X_j g_0^{-1}$ is contained in $\coprod_{k=1}^{r_j} H_j (g_k^j)^{-1}$. We have:

$$\begin{aligned} g' X_j^{(*)} \in \mathcal{B} &\iff g' X_j^{(*)} \text{ separates } g \text{ and } g_0 \\ &\iff X_j^{(*)} \text{ separates } (g')^{-1}g \text{ and } (g')^{-1}g_0 \\ &\iff (g')^{-1} \text{ is in the symmetric difference of } X_j g^{-1} \text{ and } X_j g_0^{-1} \\ &\iff (g')^{-1} \in \coprod_{k=1}^{r_j} H_j (g_k^j)^{-1} \\ &\iff g' \in \coprod_{k=1}^{r_j} g_k^j H_j \end{aligned}$$

As H_j stabilizes X_j , and as there are only finitely many j , it follows that \mathcal{B} is finite.

This completes the proof that V_g satisfies DCC. \square

Lemma 8.6. *Let (Σ, \leq) be any pocset.*

1. Any two vertices (i.e. ultrafilters on (Σ, \leq)) satisfying DCC can be connected via a finite edge path.
2. If a vertex is connected to some vertex satisfying DCC, then the vertex satisfies DCC.

Proof. 1. Assume, for contradiction, that V and W satisfy DCC but differ on infinitely many (distinct) elements, say $A_1, A_2, \dots \in V$ and $A_1^*, A_2^*, \dots \in W$. As every element of Σ comes from one of finitely many splittings, after passing to a subsequence, all the A_k 's come from a single splitting, and hence are nested. Note that as both V and W are ultrafilters, we cannot have $A_k \leq A_l^*$ or $A_k^* \leq A_l$. Since the A_k 's are nested, this shows that given $k, l \in \mathbb{N}$, either $A_k \leq A_l$ or $A_l \leq A_k$. Thus, after reordering, we either get an ascending chain in the A_k 's and descending chain in the A_k^* 's, or vice-versa. Hence the chain stabilizes, a contradiction to the A_k 's being distinct.

2. Let V be any vertex satisfying DCC, and W any vertex connected to V . Then V and W differ by only finitely many complementary pairs (A, A^*) . Any descending chain in W must have all but finitely many of its elements in V , so must stabilize (since V satisfies DCC).

□

8.4 Minimal Cubings

In [18], Niblo, Sageev, Scott, and Swarup constructed another cubing L from a finite collection $\{X_j | j = 1, \dots, n\}$ of H_j -almost invariant subsets of a group G , where, G and all the H_j 's are finitely generated. As before, let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. The authors assumed that the subsets are already in good position, and use the partial order of almost inclusion. The vertices of L are characterized as ultrafilters on

(Σ, \leq) satisfying DCC. As in Sageev's cubing, define higher skeleta of L inductively: whenever you see the boundary of an n -cube, attach an n -cube. To check that L is a CAT(0) cubical complex, use the same arguments from [21] that prove that Sageev's cubing is a CAT(0) cubical complex. To prove that L is nonempty, the authors constructed basic vertex analogues.

The goal of this chapter is to prove that one can construct the same cubing L from [18], under a new set of assumptions. Namely, here we assume that the X_j 's come from splittings and that the splittings satisfying sandwiching, and do not assume that G or the H_j 's are finitely generated. The key fact needed is that the cubing is always nonempty (see Theorem 8.11).

Theorem 8.7. *Let G be any group with a finite collection $\{\sigma_j | j = 1, \dots, n\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j = 1, \dots, n\}$ satisfies sandwiching. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. Then there exists a CAT(0) cubical complex L , with a bijective correspondence between Σ and the set of oriented hyperplanes of L , such that two elements of Σ cross if, and only if, their corresponding hyperplanes cross in some square. Moreover, the hyperplane corresponding to any $A \in \Sigma$ determines a $\text{Stab}(A)$ -almost invariant set that is $\text{Stab}(A)$ -almost equal to A .*

As noted in [18], every ultrafilter on (Σ, \leq) is also an ultrafilter on (Σ, \subset) , and any ultrafilter satisfying DCC with respect to \leq also satisfies DCC with respect to \subset , so that every vertex in L is canonically a vertex in C . We will see that the embedding $L^0 \hookrightarrow C^0$ naturally extends to an embedding $L \hookrightarrow C$. However, in general, C contains many vertices that are not in L . For example, either all basic vertices are in L , or L contains no basic vertices:

Lemma 8.8. *Let G be any group with a finite collection $\{X_j | j = 1, \dots, n\}$ of H_j -almost invariant subsets. Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. Suppose almost inclusion \leq defines a partial order on Σ . Then V_g is an ultrafilter on (Σ, \leq) for all $g \in G$ if, and only if, V_g is an ultrafilter on (Σ, \leq) for some $g \in G$.*

Proof. We will use the arguments from the proof of Lemma 8.5 show that V_g satisfies DCC with respect to the partial order \leq , as follows (regardless of whether or not V_g actually is an ultrafilter on (Σ, \leq)). Suppose we have a descending chain $A_1 \geq A_2 \geq \dots$ in a basic vertex V_g . Since the A_i come from only finitely many splittings, after passing to a subchain, we have in fact $A_1 \supset A_2 \supset \dots$, which must stabilize, since it is a descending chain in (V_g, \subset) and (V_g, \subset) satisfies DCC.

Now, suppose there exists $g \in G$ such that V_g is not an ultrafilter on (Σ, \leq) . Clearly for all $A \in \Sigma$, either A or A^* is in V_g . Hence there must exist $A, B \in \Sigma$ with $A \in V_g$, $A \leq B$, and $B \notin V_g$. The pair (A, B) prevents V_g from being an ultrafilter. Now, for any $g' \in G$, the pair $(g'A, g'B)$ prevents $V_{g'}$ from being an ultrafilter. \square

Note that if the X_j 's are in very good position, then the minimal cubing will be the same as Sageev's cubing. Otherwise, Sageev's cubing will contain many more vertices than the minimal cubing.

In what follows, we prove that L is nonempty (assuming sandwiching, but not assuming any finite generation). To create an ultrafilter on (Σ, \leq) , we will start with a basic ultrafilter on $\Sigma(X_1) \subset \Sigma$, then extend. In general, suppose $\Sigma_0 \subset \Sigma$ and V_0 is an ultrafilter on Σ_0 . If we hope to extend V_0 to an ultrafilter on all of Σ , we must add to V_0 all elements B of Σ for which $A \leq B$ for some $A \in V_0$. We call this process taking the closure of V_0 . More formally:

Definition 8.9. Let (Σ, \leq) be any partially ordered set with complementation, and

let V_0 be an ultrafilter on some subset Σ_0 of Σ . The *closure* of V_0 is:

$$\overline{V_0} := V_0 \cup \{B \in \Sigma - \Sigma_0 \mid \text{there exists } A \in V_0 \text{ such that } A \leq B\}.$$

Lemma 8.10. *Let (Σ, \leq) be any partially ordered set with complementation, and let V_0 be an ultrafilter on some subset Σ_0 of Σ . Let $\overline{\Sigma_0}$ denote the set $\{A \in \Sigma \mid A \in \overline{V_0} \text{ or } A^* \in \overline{V_0}\}$. Then $\overline{V_0}$, the closure of V_0 in Σ , is an ultrafilter on $(\overline{\Sigma_0}, \leq)$.*

Proof. To prove that $\overline{V_0}$ is an ultrafilter on $(\overline{\Sigma_0}, \leq)$, we must show that conditions 1 and 2 of Definition 8.3 are satisfied.

1. Clearly for all $B \in \overline{\Sigma_0}$, at least one of B, B^* is in $\overline{V_0}$. We must show that if $B \in \overline{V_0}$, then $B^* \notin \overline{V_0}$. Suppose, for contradiction, that $B \in \overline{V_0}$ and $B^* \in \overline{V_0}$. By construction, either $B, B^* \in \Sigma_0$ or $B, B^* \in \overline{\Sigma} - \Sigma_0$, so that either $B, B^* \in V_0$ or $B, B^* \in \overline{V_0} - V_0$. As V_0 is an ultrafilter, it is impossible to have both B and B^* in V_0 . Hence we must have $B, B^* \in \overline{V_0} - V_0$. By the definition of “closure,” there exist $A, A' \in V_0$ such that $A \leq B$ and $A' \leq B^*$ (i.e. $B \leq A'^*$). Transitivity of \leq now implies $A \leq A'^*$. Since $A \in V_0$, this implies $A'^* \in V_0$, a contradiction to V_0 being an ultrafilter. This completes the proof that we cannot have simultaneously $B \in \overline{V_0}$ and $B^* \in \overline{V_0}$.
2. Assume $B, C \in \overline{\Sigma_0}$ with $B \in \overline{V_0}$ and $B \leq C$. We must show that $C \in \overline{V_0}$. By the construction of $\overline{V_0}$, there exists $A \in V_0$ with $A \leq B$. By transitivity of \leq , we have $A \leq C$ (equivalently, $C^* \leq A^*$). We break up the rest of the proof into two cases, depending on whether $C \in \Sigma_0$.
 - If $C \in \Sigma_0$, since V_0 is an ultrafilter on Σ_0 , either C or C^* must be in V_0 . If $C^* \in V_0$, then $C^* \leq A^*$ implies $A^* \in V_0$, which is impossible since $A \in V_0$. Hence we must have $C \in V_0$. Since V_0 is a subset of $\overline{V_0}$, this implies $C \in \overline{V_0}$.

- If $C \in \overline{\Sigma_0} - \Sigma_0$, then $A \leq C$ implies $C \in \overline{V_0}$.

Hence we must have $C \in \overline{V_0}$.

This completes the proof that the closure of an ultrafilter is an ultrafilter. \square

The following theorem (whose proof is lengthy) proves that the cubing L is nonempty.

Theorem 8.11. *Let G be any group with a finite collection $\{\sigma_j | j = 1, \dots, n\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j = 1, \dots, n\}$ satisfies sandwiching. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Let $\Sigma := \bigcup_{j=1}^n \Sigma(X_j)$. Then there exists an ultrafilter on (Σ, \leq) satisfying DCC.*

Proof. Fix $g \in G$. We will start with a basic ultrafilter V_1 on $\Sigma(X_1)$, take its closure in Σ , and inductively add in part of a basic ultrafilter on each $\Sigma(X_j)$ until we have defined an ultrafilter on all of Σ . The ultrafilters produced in all steps are, in order, $V_1 \subset \overline{V_1} \subset V_2 \subset \overline{V_2} \subset \dots \subset V_n =: V$. In Step ja, we add to the ultrafilter $\overline{V_{j-1}}$ all elements A of $\Sigma(X_j)$ such that $g \in A$ and neither A nor A^* was already in the ultrafilter, to get the ultrafilter V_j . In Step jb, we take the closure of the ultrafilter from Step ja, to get $\overline{V_j}$. In the end, we get an ultrafilter V on Σ .

- Step 1a: Define V_1 as a basic ultrafilter on $\Sigma(X_1)$:

$$V_1 := \{A \in \Sigma(X_1) | g \in A\}.$$

Let $\Sigma_1 := \Sigma(X_1)$. On $\Sigma(X_1)$, the inclusion relation is the same as \leq , so Lemma 8.5 proves that V_1 is an ultrafilter on (Σ_1, \leq) .

- Step 1b: Define $\overline{V_1}$ to be the closure of V_1 in Σ :

$$\overline{V_1} := V_1 \cup \{B \in \Sigma(X_2, X_3, \dots, X_n) | \text{there exists } A \in V_1 \text{ with } A \leq B\}.$$

Let $\overline{\Sigma}_1 := \{A \in \Sigma \mid A \in \overline{V}_1 \text{ or } A^* \in \overline{V}_1\}$. By Lemma 8.10, \overline{V}_1 is an ultrafilter on $(\overline{\Sigma}_1, \leq)$.

Perform the following two steps for $1 < j < n$.

- Step ja: Define V_j to be the union of \overline{V}_{j-1} and part of a basic ultrafilter on $\Sigma(X_j)$:

$$V_j := \overline{V}_{j-1} \cup \{A \in \Sigma(X_j) \mid g \in A \text{ and } A \notin \overline{\Sigma}_{j-1}\}.$$

Let $\Sigma_j := \{A \in \Sigma \mid A \in V_j \text{ or } A^* \in V_j\}$. To prove that V_j is an ultrafilter on (Σ_j, \leq) , we must show that conditions 1 and 2 of Definition 8.3 are satisfied.

1. Σ_j is defined to be the union of the elements of V_j and their complements.

We must show that if $A \in V_j$, then $A^* \notin V_j$. Suppose, for contradiction, that both A and A^* are elements of V_j . By construction, either $A, A^* \in \Sigma_j - \overline{\Sigma}_{j-1}$, or $A, A^* \in \overline{\Sigma}_{j-1}$, so that either $A, A^* \in V_j - \overline{V}_{j-1}$, or $A, A^* \in \overline{V}_{j-1}$. If $A, A^* \in \overline{V}_{j-1}$, this would contradict \overline{V}_{j-1} being an ultrafilter; hence it must be the case that $A, A^* \in V_j - \overline{V}_{j-1}$. This implies $g \in A$ and $g \in A^*$, also a contradiction. Hence it is impossible to have simultaneously $A \in V_j$ and $A^* \in V_j$.

2. Assume $B, C \in \Sigma_j$ with $B \in V_j$ and $B \leq C$. We must show that $C \in V_j$.

We break up the proof that $C \in V_j$ into two cases, depending on whether B was added to the ultrafilter in Step ja or a previous step.

- (a) Suppose B was added in Step ja, i.e. $B \in V_j - \overline{V}_{j-1}$. Note that $B \in \Sigma(X_j)$. Either $C \in \overline{\Sigma}_{j-1}$ or $C \in \Sigma_j - \overline{\Sigma}_{j-1}$. We want to show that $C \in V_j$. If $C \in \overline{\Sigma}_{j-1}$ and $C \notin V_j$, then we must have $C^* \in V_j$. Then $C^* \leq B^*$ would imply B^* was added to the ultrafilter by Step (j-1)b, so it would be impossible to have $B \in V_j - \overline{V}_{j-1}$. If instead $C \in \Sigma_j - \overline{\Sigma}_{j-1}$,

then $C \in \Sigma(X_j)$. As B and C are both elements of $\Sigma(X_j)$, having $B \leq C$ implies $B \subset C$. Hence $g \in B$ implies $g \in C$, so that $C \in V_j$.

(b) Suppose B was added in a previous step, i.e. $B \in \overline{V_{j-1}}$. Then by construction of $\overline{V_{j-1}}$, there exists $A \in V_{j-1}$ with $A \leq B$. By transitivity of \leq , we have $A \leq C$. It follows that $C \in \overline{V_{j-1}}$. As $\overline{V_{j-1}} \subset V_j$, we have $C \in V_j$.

In any case, we conclude $C \in V_j$.

Hence V_j is an ultrafilter on Σ_j .

- Step jb: Define $\overline{V_j}$ to be the closure of V_j in Σ :

$$\overline{V_j} := V_j \cup \{B \in \Sigma(X_{j+1}, X_{j+2}, \dots, X_n) \mid \text{there exists } A \in V_j \text{ with } A \leq B\}.$$

Let $\overline{\Sigma_j} := \{A \in \Sigma \mid A \in \overline{V_j} \text{ or } A^* \in \overline{V_j}\}$. By Lemma 8.10, $\overline{V_j}$ is an ultrafilter on $(\overline{\Sigma_j}, \leq)$.

Perform one last step to define an ultrafilter on all of Σ .

- Step na (this is just Step ja with $j = n$)

Note that $\Sigma_n = \Sigma$, so there is no need for Step nb. Let $V := V_n$.

We have successfully defined an ultrafilter V on (Σ, \leq) . Next we prove that V satisfies DCC. To make the proof less cumbersome, I will write *WLOG* to denote *without loss of generality*.

Suppose we have an (infinite) descending chain $B_1 \geq B_2 \geq \dots$ in V . We will obtain $A_k \leq B_k$, show that the A_k 's stabilize, and then show that the B_k 's stabilize.

1. Since each element of Σ comes from one of only finitely many splittings, WLOG (after passing to a subchain of $(B_k)_{1 \leq k}$) there exists a fixed $j \in \{1, 2, \dots, n\}$ such that $B_k \in \Sigma(X_j)$.

2. If there exists an infinite subchain of the B_k 's that were added to V in Stepja, then since g is in each element of the subchain, the proof of Lemma 8.5 shows that the subchain must stabilize, so that the original chain stabilizes, completing the proof of the theorem. Otherwise, WLOG (after passing to a subchain of $(B_k)_{1 \leq k}$), each of the B_k 's was added to V in a type "b" Step (before Stepja). Recall that j was fixed in the previous step, and $B_k \in \Sigma(X_j)$, for all k .
3. For each B_k , since B_k was added in a type "b" Step before Step ja, there exists $A_k \in V_{j-1}$ such that $A_k \leq B_k$.
4. WLOG (after possibly replacing A_k by something less than A_k), each A_k was added in a type "a" Step, so that $g \in A_k$.
5. WLOG (after passing to a subchain of $(A_k \leq B_k)_{1 \leq k}$), there exists a fixed j' (for some $1 \leq j' \leq j - 1$) such that $A_k \in \Sigma(X_{j'})$, for all k . In particular, since $X_{j'}$ arises from a splitting, all the A_k 's are nested.
6. We cannot have $A_k \subset A_l^*$ or $A_k^* \subset A_l$ (since all the A 's belong to the ultrafilter V), hence for all $k \neq l$, we must have $A_k \subset A_l$ or $A_k \supset A_l$.
7. In this step, we show that the A_k 's stabilize. If there is an infinite ascending subchain of the A_k 's such that each is contained in the next, then WLOG (after replacing each A_k in the subchain by the first one) all the A_k 's in that subchain are equal, so move on to the next step of the proof. Otherwise, WLOG (after passing to a subchain of $(A_k \leq B_k)_{1 \leq k}$) we have $A_k \supset A_l$, for all $k < l$. Since $g \in A_k$ for all k , and since all of the A_k 's are in $\Sigma(X_{j'})$, the A_k 's must stabilize after finitely many steps (by Lemma 8.5). So WLOG (after passing to a subchain of $(A_k \leq B_k)_{1 \leq k}$), all the A_k 's are identical.

8. Recall that $B_k \geq A_1$, for all k . We now have $B_1 \geq B_2 \geq \dots \geq A_1$. But since $\{\sigma_j | 1 \leq j \leq n\}$ satisfies sandwiching, this contradicts interval finiteness (see Proposition 7.4), unless the B_k 's stabilize. Hence the B_k 's stabilize.

This completes the proof that V satisfies DCC. In particular, we have shown there exists an ultrafilter on (Σ, \leq) satisfying DCC. \square

We are now ready to prove Theorem 8.7. As before, define the vertices of L to be the ultrafilters on (Σ, \leq) satisfying DCC, and define higher skeleta of L inductively: whenever you see the boundary of an n -cube, attach an n -cube. To check that L is a CAT(0) cubical complex, apply the arguments from [21] that prove that Sageev's cubing is a CAT(0) cubical complex. Theorem 8.11 shows that L is nonempty. Furthermore, the proof of Theorem 8.11 shows that each $A \in \Sigma$ is in some vertex of L : pick any $g \in A$, reorder the splittings (and their associated X_j) such that $A \in X_1$, and construct V as in the proof of Theorem 8.11. For each $A \in \Sigma$, since A is in some vertex of L and A^* is in some vertex of L , there exists a (nonempty) oriented hyperplane corresponding to the set of all edges that exit A . By Theorem 8.1, the hyperplane corresponding to A determines a $\text{Stab}(A)$ -almost invariant subset of G that is $\text{Stab}(A)$ -almost equal to A . This completes the proof of Theorem 8.7.

8.5 Putting the X_j 's in Very Good Position

Let G be any group. Take any finite collection $\{\sigma_1, \dots, \sigma_n\}$ of pairwise non-isomorphic splittings of G satisfying sandwiching. Let X_j be an H_j -almost invariant set arising from σ_j , and let $\Sigma = \{gX_j, gX_j^* | g \in G, j = 1, \dots, n\}$. We have shown that Σ is in good position, i.e. if two corners of (A, B) are small, then (at least) one is empty. This allowed us to define the partial order \leq on Σ . Now we show how to find $X'_j \stackrel{H_j\text{-a}}{=} X_j$ such that X'_j arises from the same splitting of G from which X_j

arises, and $\Sigma' = \{gX'_j, gX'^*_j | g \in G, j = 1, \dots, n\}$ is in very good position. This result was previously proved by Niblo–Sageev–Scott–Swarup [18] for a finite collection of almost invariant sets over finitely generated subgroups of a finitely generated group.

Theorem 8.12. *Let G be any group with a finite collection $\{\sigma_j | j = 1, \dots, n\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j = 1, \dots, n\}$ satisfies sandwiching. For each j , let X_j be an H_j -almost invariant set arising from σ_j . Then there exist $X'_j \stackrel{H_j\text{-a}}{=} X_j$, such that $\Sigma' := \bigcup_{j=1}^n \Sigma(X'_j)$ is in very good position.*

In preparation for proving the theorem, we take a look at how hyperplanes in the cubings C and L compare to each other. For each j , let \mathcal{H}_j be the hyperplane in C determined by the equivalence class of edges in C exiting X_j . (Or, equivalently, the class of edges equivalent to any given edge exiting X_j , with the equivalence relation generated by square-equivalence.) Define the halfspace \mathcal{H}_j^+ of C by:

$$\mathcal{H}_j^+ = \{V \in C^{(0)} | X_j \in V\}.$$

For any vertex $v \in C$, define $(X_j)_v$ by:

$$(X_j)_v := \{g \in G | gv \in \mathcal{H}_j^+\}.$$

Note that $(X_j)_{V_e} = X_j$, where V_e is the basic ultrafilter on (Σ, \subset) consisting of all elements of Σ containing the identity. Let \mathcal{K} be the hyperplane in L determined by the equivalence class of edges exiting X_j . Define the halfspace \mathcal{K}_j^+ of L by:

$$\mathcal{K}_j^+ := \{W \in L^{(0)} | X_j \in W\}.$$

Recall the canonical embedding $L_0 \hookrightarrow C_0$, in which we view any vertex in L , i.e. an ultrafilter $W \subset \Sigma$ on (Σ, \leq) satisfying DCC, as an ultrafilter on (Σ, \subset) . Now we will see how this extends to an embedding $L \hookrightarrow C$. If two edges in L are opposite

sides of a square in C then all four vertices of the square in C are in L , so that the other two edges of the square in C are also in L . Hence two edges in L are on the opposite sides of a square in L if, and only if, they are on opposite sides of a square in C . It follows that $\mathcal{H}_j \cap L = \mathcal{K}_j$ and $\mathcal{H}_j^+ \cap L = \mathcal{K}_j^+$.

Now we can use the cubing L to put the X_j 's in very good position.

Lemma 8.13. *Fix a vertex $w \in L \subset C$, and define $X'_j := \{g \in G \mid gw \in \mathcal{K}_j^+\}$. Then each X'_j is H_j -almost invariant, the collection $\Sigma' := \bigcup_{j=1}^n \Sigma(X'_j)$ is in very good position, and $X_j \mapsto X'_j$ induces a G -equivariant isomorphism from (Σ, \leq) to (Σ', \subset) .*

Proof. Viewing w as a vertex in C , we have $(X_j)_w = \{g \in G \mid gw \in \mathcal{H}_j^+\}$. As $L \hookrightarrow C$ is G -equivariant, it follows that $X'_j = (X_j)_w$. Now, applying Lemma 8.1, $X'_j = (X_j)_w$ is H_j -almost invariant and H_j -almost equal to $(X_j)_{v_e} = X_j$.

Next we show, as proved in [18, Lemma 4.1], that $X_j \mapsto X'_j$ induces a G -equivariant isomorphism from (Σ, \leq) to (Σ', \subset) . We are assuming that none of the X_j come from isomorphic splittings. Since w is an ultrafilter on (Σ, \leq) , it follows that for all $g \in G$, the translate gw is also an ultrafilter on (Σ, \leq) . If $A, B \in \Sigma$, let A', B' denote the images in Σ' of A and B by the map that sends $X_j \mapsto X'_j$.

Suppose $A \leq B$, i.e. $A \cap B^*$ is empty or the only small corner of the pair (A, B) . Note that since $X_j \stackrel{H_j\text{-a}}{=} X'_j$, we have $A \stackrel{\text{Stab}(A)\text{-a}}{=} A'$ and $B \stackrel{\text{Stab}(B)\text{-a}}{=} B'$. Hence a corner of the pair (A, B) is small if, and only if, the corresponding corner of the pair (A', B') is small. If $A \cap B^*$ is the only small corner of the pair (A, B) , then since A' and B' are nested, we must have $A' \subset B'$, as desired. If the pair (A, B) has two small corners, then $A \subset B$. To see that $A' \subset B'$, simply note that since $X'_j = \{g \in G \mid gw \in \mathcal{H}_j^+\}$, we have $A' = \{g \in G \mid A \in gw\}$, and $B' = \{g \in G \mid B \in gw\}$. In either case, we conclude that $A' \subset B'$.

Conversely, suppose that $A' \subset B'$. We need to show that $A \leq B$. Since Σ is in

good position, it follows that $A \leq B$ or $B \leq A$. The above paragraph shows that $A \leq B$. Hence $A \leq B$ if, and only if, $A' \subset B'$.

It follows that Σ' is in very good position. □

This completes the proof of Theorem 8.12.

CHAPTER 9

Existence of Algebraic Regular Neighborhoods

We start with a finite collection of splittings $\{\sigma_1, \dots, \sigma_n\}$ over subgroups H_j of a group G . Let X_j be an H_j -almost invariant set arising σ_j . Suppose $\{X_1, \dots, X_n\}$ satisfies sandwiching (see Definition 3.9). We also assume no two of the σ_j 's are isomorphic to each other. We will construct a bipartite G -tree $T(X_1, \dots, X_n)$, and show T is an algebraic regular neighborhood of $\{\sigma_1, \dots, \sigma_n\}$ (see Definition 3.13).

Let Σ denote the set of all translates of all the X_j and their complements. Since no two σ_j 's are isomorphic to each other, \leq defines a partial order on Σ (see Corollary 6.5). Construct the cubing L from Theorem 8.7. We will construct a bipartite tree from L .

Define a “cross connected” relation on Σ , generated by relations of the form “ A crosses B .” More precisely, A is cross connected to B if one of the following holds.

1. A is equal to B or B^* .
2. There exists some $m \geq 0$ and a sequence (A, B_1, \dots, B_m, B) such that A crosses B_1 , B_1 crosses B_2, \dots, B_{m-1} crosses B_m , and B_m crosses B .

This defines an equivalence relation. Call each equivalence class a **cross connected component** (*ccc*).

One can easily see *ccc*'s of Σ from looking at the cubing L . Each oriented hyperplane in the cubing corresponds to a unique element of Σ . Call a vertex of L a **cut vertex** if it separates L into more than one component. If we remove all cut vertices from the cubing, we are left with a disjoint collection of components, where each component is a subcubing with some vertices missing. Each *ccc* has all its hyperplanes contained in a single component. Moreover, since the components have no cut vertices, each component's hyperplanes come from only one *ccc*. This gives a bijective correspondence between the components of (L minus its cut vertices) and the *ccc*'s of Σ .

We introduce some basic notation. View the cubing as a disjoint union of subcubings which are glued together at cut vertices. Let CUT denote the set of cut vertices. Let SUB denote the set of (disjoint) subcubings. Note that we have a bijective correspondence between SUB and the *ccc*'s of Σ . For a given subcubing $\alpha \in SUB$, define the **corner vertices of α** to be the vertices of α that are glued to cut vertices. Let $CRN(\alpha)$ denote the set of corner vertices of α .

For each subcubing $\alpha \in SUB$, create a tree whose vertices are $CRN(\alpha)$ plus a central vertex, and an edge connecting each element of $CRN(\alpha)$ to the central vertex. Call this tree \star_α .

Glue the \star_α 's together by, for each element of CUT , identifying all corner vertices of all the \star_α 's which came from that element of CUT . Color the equivalence classes of corner vertices as V_1 -vertices, and color all the central vertices as V_0 -vertices.

The result is a bipartite G -tree. Let $T(X_1, \dots, X_n)$ denote this tree.

Theorem 9.1 (Existence of algebraic regular neighborhoods). *Let G be any group with a finite collection $\{\sigma_j | j = 1, \dots, n\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j = 1, \dots, n\}$ satisfies sandwiching. For each j , let X_j be an H_j -almost*

invariant set arising from σ_j . Then $T := T(X_1, \dots, X_n)$ is algebraic regular neighborhood of $\{\sigma_j | j = 1, \dots, n\}$.

To prove the theorem, we need to show that T satisfies the five conditions of Definition 3.13.

Lemma 9.2 (First condition). *Each σ_j is enclosed by some V_0 -vertex orbit in T , and each V_0 -vertex orbit encloses some σ_j .*

Proof. Fix j . We will use the original cubing L to construct a particular refinement of T . Recall that CUT denotes the cut vertex set of the original cubing, SUB denotes the set of subcubings, and $CRN(\alpha)$ denotes the set of corner vertices of $\alpha \in SUB$. For each subcubing $\alpha \in SUB$ not containing any X_j -hyperplanes, define \star_α as above.

For each subcubing $\alpha \in SUB$ that contains $\Sigma(X_j)$ -hyperplanes. Let $\#_\alpha$ denote the dual tree to the $\Sigma(X_j)$ -hyperplanes in α . For each element of $CRN(\alpha)$, make a vertex and attach it to $\#_\alpha$ with an edge. Specifically, attach the edge to the vertex of $\#_\alpha$ that corresponds to the component of $\alpha - (\Sigma(X_j)\text{-hyperplanes})$ containing the corner vertex. Call this tree \star'_α .

Glue the \star'_α 's (for ccc 's containing $\Sigma(X_j)$ -hyperplanes) and \star_α 's (for ccc 's not containing $\Sigma(X_j)$ -hyperplanes) together by, for each element of CUT , identifying all corner vertices which came from that element of CUT . Color the equivalence classes of corner vertices as V_1 -vertices, and color all the other vertices as V_0 -vertices.

This new tree maps naturally to T by collapsing the new edge orbit. On the other hand, the new tree maps to a tree for σ by collapsing all edges except for the new edge orbit. □

Lemma 9.3 (Second condition). *If σ is a splitting of G over H , where σ is sand-*

wiched by σ_j and $i(\sigma, \sigma_j) = 0$, for all $j \in J$, and $i(\sigma, \sigma_j) = 0$, then σ is enclosed by some V_1 -vertex orbit in T .

Proof. Let X be an H -almost-invariant set arising from σ , where X is sandwiched by X_j and $i(\sigma, \sigma_j) = 0$, for all $1 \leq j \leq n$.

Construct a new cubing from $\Sigma(X_1, \dots, X_n, X)$ (using the partial order \leq). Since X does not cross any element of $\bigcup_{j=1}^n \Sigma(X_j)$, each new subcubing α in SUB is simply an edge, and \star_α consists of two V_1 -vertices connected to a V_0 -vertex. Construct the new tree $T(X_1, \dots, X_n, X)$. The new tree projects naturally to a tree for σ , by collapsing each old \star_α to a point, and forgetting the new V_0 -vertices. On the other hand, the new tree naturally projects to the old tree, with each new \star_α being collapsed to a single V_1 -vertex. Hence σ is enclosed by a V_1 -vertex orbit. \square

Lemma 9.4 (Third condition). *T is a minimal G -tree.*

Proof. Let T_0 be the minimal sub- G -tree of T (or any fixed vertex, if G fixes a vertex of T).

If T_0 has no V_0 -vertices, then since T is bipartite, T_0 must consist of a single V_1 -vertex which is fixed by G . Let V denote a V_0 -vertex adjacent to the fixed V_1 -vertex. Since the orbit of V encloses σ_j for some j , and since $\Sigma(X_j)$ has infinite chains, V satisfying DCC implies that there exists a translate of V not adjacent to the fixed V_1 -vertex. This is impossible. Hence T_0 must contain a V_0 -vertex.

Take any V_0 -vertex V' in T_0 , and pick $j \in J$ such that σ_j is enclosed by the orbit of V' . If $T_0 \neq T$, then we can find a V_0 -vertex V'' in $T - T_0$. Pick $k \in J$ such that σ_k is enclosed by the orbit of V'' . As σ_j is sandwiched by σ_k , there exists a translate of V' that is not in T_0 . This is impossible, as T_0 is G -invariant. Hence we must have $T_0 = T$. This completes the proof that T is a minimal G -tree. \square

Lemma 9.5 (Fourth condition). *There exists a bijection*

$$f : \{j \in J \mid \sigma_j \text{ is isolated}\} \rightarrow G\text{-orbits of isolated } V_0\text{-vertices of } T$$

such that $f(j)$ encloses σ_j .

Proof. Each isolated V_0 -vertex corresponds to a subcubing $\alpha \in SUB$ consisting of exactly one edge, or equivalently, exactly one hyperplane. This hyperplane corresponds to a unique pair $\{A, A^*\} \subset \Sigma$. \square

Lemma 9.6 (Fifth condition). *Every non-isolated V_0 -vertex orbit in T encloses some non-isolated σ_j .*

Proof. Any non-isolated V_0 -vertex corresponds to a subcubing $\alpha \in SUB$ containing at least two hyperplanes that cross each other. \square

This completes the proof of Theorem 9.1.

CHAPTER 10

Uniqueness of Algebraic Regular Neighborhoods

We prove uniqueness of algebraic regular neighborhoods for an arbitrary collection of splittings of G satisfying sandwiching.

Theorem 10.1 (Uniqueness of algebraic regular neighborhoods). *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j \in J\}$ satisfies sandwiching, and that T_1 and T_2 are algebraic regular neighborhoods of $\{\sigma_j | j \in J\}$. Then there exists a G -equivariant, color preserving isomorphism from T_1 to T_2 .*

The proof of Theorem 10.1 is laid out in this chapter. We will use the same strategy that Scott and Swarup used to prove [25, Theorem 6.7]. Namely, insert an edge orbit in T_1 for each edge splitting of T_2 that is not already an edge splitting of T_1 , and vice-versa. Then we will show a contradiction if we actually had to insert any edge orbits. To “insert edge orbits” in T_1 or T_2 , we need to know that the edge splittings in T_1 and T_2 are compatible with the edge splittings to be inserted.

Lemma 10.2. *Suppose $\{\sigma_k | k \in K\}$ and $\{\sigma_l | l \in L\}$ are collections of splittings of G , such that their union satisfies sandwiching. Assume $\{\sigma_k | k \in K\}$ and $\{\sigma_l | l \in L\}$ are each compatible, and that $i(\sigma_k, \sigma_l) = 0$ for all $k \in K$ and $l \in L$. Then $\{\sigma_k | k \in K \cup L\}$ is compatible.*

Proof. For each $k \in K \cup L$, let X_k be an almost invariant set arising from σ_k . Without loss of generality, choose the X_k 's so that if σ_k is isomorphic to σ_l , then $X_k = X_l$ (as subsets of G). Let Σ denote the collection of all translates of the almost invariant sets and their complements:

$$\Sigma := \bigcup_{k \in K \cup L} \Sigma(X_k).$$

The proof of Theorem 7.5 directly carries through, provided we can prove interval finiteness. We need to show that for all $A, B \in \Sigma$, there are only finitely many $C \in \Sigma$ with $A \leq C \leq B$.

Fix $A, B \in \Sigma$, and recall that $\{\sigma_k | k \in K\} \cup \{\sigma_l | l \in L\}$ satisfies sandwiching. Pick $A_1, B_1 \in \{\sigma_k | k \in K\}$ such that $A_1 \leq A$ and $B \leq B_1$. Similarly, pick $A_2, B_2 \in \{\sigma_l | l \in L\}$ such that $A_2 \leq A$ and $B \leq B_2$. There are only finitely many $C \in \bigcup_{k \in K} \Sigma(X_k)$ with $A_1 \leq C \leq B_1$, and only finitely many $C \in \bigcup_{l \in L} \Sigma(X_l)$ with $A_2 \leq C \leq B_2$. Hence there are only finitely many $C \in \Sigma$ with $A \leq C \leq B$.

Apply Dunwoody's theorem to get a G -tree (see Chapter 3.11). For each edge orbit Ge , let n_e denote the number of splittings in $\{\sigma_k | k \in K \cup L\}$ that are isomorphic to the edge splitting for e , and subdivide each edge in Ge into an interval of n_e edges. The result is a compatibility tree for $\{\sigma_k | k \in K \cup L\}$. \square

To apply Lemma 10.2 to the proof of Theorem 10.1, we need to know that each edge splitting of an algebraic regular neighborhood is sandwiched by each σ_j .

Lemma 10.3. *Let G be any group with any collection $\{\sigma_j | j \in J\}$ of pairwise non-isomorphic splittings. Suppose $\{\sigma_j | j \in J\}$ satisfies sandwiching, and let T be an algebraic regular neighborhood of $\{\sigma_j | j \in J\}$. Then each edge splitting of T is sandwiched by $\{\sigma_j | j \in J\}$.*

Proof. Assume there exists some edge e of T and some $j \in J$ such that splitting from

e is not sandwiched by σ_j . Let σ denote the splitting from e .

Let V be some V_0 -vertex of T whose orbit encloses σ_j . The convex hull of all translates of V is a G -invariant subtree of T . The assumption that σ is not sandwiched by σ_j implies that all translates of V lie on one side of e , so that e is not in the convex hull of all translates of V . This implies that T is not a minimal G -tree, a contradiction to T being an algebraic regular neighborhood. \square

Now we present the proof of Theorem 10.1.

Proof. First we prove the theorem in the case that no σ_j is an isolated splitting, i.e. no σ_j has intersection number zero with every other splitting in the collection. We will show that the edge splittings of T_1 and T_2 are isomorphic. Then by a uniqueness result (see Corollary 6.7), T_1 and T_2 are G -isomorphic. Assume (for contradiction) that T_1 and T_2 have different edge splittings. By “different edge splittings,” we mean that T_1 (or T_2) has an edge splitting not isomorphic to any edge splitting in T_2 (or T_1), or that T_1 (or T_2) has strictly more edge orbits than T_2 (or T_1) yielding splittings in a given isomorphism class.

Let σ be some edge splitting of T_1 , call it σ , that is not in T_2 (in the sense described above). As each σ_j is enclosed by some V_0 -vertex of T_1 , each σ_j has intersection number zero with σ . Moreover, by Lemma 10.3, σ is sandwiched by each $\Sigma(X_j)$. By condition number 2 of the definition of algebraic regular neighborhood, σ is enclosed by some V_1 -vertex of T_2 , so that we can refine T_2 by adding one edge orbit which represents σ .

By Lemma 10.2, we can apply the above procedure simultaneously for all edge splittings of T_1 that are not in T_2 . Let the tree T_{21} denote a tree obtained from T_2 by splitting at V_1 -vertices for each edge splitting of T_1 that was not already in T_2 . When

splitting at a V_1 -vertex, color both endpoints of the new edge as V_1 -vertices. Define T_{12} similarly. T_{12} and T_{21} may have infinitely many edge orbits, but by Corollary 6.7, T_{12} and T_{21} are isomorphic G -trees.

If T_{21} has an edge e not in T_2 , then when adding in e , we would have split T_2 at a V_1 -vertex. Under the isomorphism from T_{21} to T_{12} , the edge e must map to an original edge of T_1 , so the isomorphism must identify a V_1 -vertex of T_{21} with an original V_0 -vertex of T_1 . Pick some splitting σ_k enclosed by that V_0 -vertex. By the isomorphism, σ_k is enclosed by a V_1 -vertex of T_2 , and hence has intersection number zero with every splitting in $\{\sigma_j | j \in J\}$, so that σ_k is an isolated splitting. This contradicts the assumption that none of the splittings in $\{\sigma_j | j \in J\}$ are isolated. Hence no edges were added to T_2 , i.e. $T_{21} = T_2$.

A similar argument shows that $T_{12} = T_1$. Hence the isomorphism between T_{21} to T_{12} is actually an isomorphism between T_2 and T_1 . If the isomorphism did not preserve color, then as in the above paragraph, the isomorphism would identify a V_0 -vertex of one tree with a V_1 vertex of another, and hence one of the σ_j 's would be isolated. This completes the proof of uniqueness of algebraic regular neighborhoods, in the case where no σ_j is isolated.

Second, we prove the theorem in the case where $\{\sigma_j | j \in J\}$ has some isolated splittings. For each V_0 -vertex in an orbit corresponding to an isolated σ_j in the definition of algebraic regular neighborhood, forget the vertex. This leaves an edge bounded by two V_1 -vertices and yielding a splitting isomorphic to σ_j . Let T'_1 denote the resulting tree. Define T'_2 similarly.

If all the σ_j 's are isolated, then no V_0 -vertices remain, so T'_1 and T'_2 are compatibility trees for $\{\sigma_j | j \in J\}$. Then Corollary 6.7 proves that T'_1 and T'_2 are G -equivariantly isomorphic.

If not all of the σ_j 's are isolated, consider each edge splitting in T'_1 that is not in T'_2 . Without loss of generality, we can take each such edge orbit to consist of edges where one endpoint is V_0 and the other is V_1 (as opposed to edges bounded by two V_1 -vertices, resulting from a forgotten V_0 -vertex). Now apply the above procedure to T'_2 and T'_1 to obtain T'_{21} and T'_{12} , and an isomorphism from T'_{21} to T'_{12} . If $T'_{21} \neq T'_2$ or $T'_{12} \neq T'_1$, then the isomorphism from T'_{21} to T'_{12} must identify a V_0 and a V_1 -vertex. This is impossible, as the non-forgotten V_0 -vertices are not isolated. Similarly, the isomorphism from T'_{21} to T'_{12} must preserve color. Hence we get a G -equivariant, color preserving isomorphism from T'_2 to T'_1 .

To get an isomorphism from T_2 to T_1 , add a V_0 -vertex in the middle of every edge bounded by two V_1 -vertices. This completes the proof of uniqueness of algebraic regular neighborhoods. \square

CHAPTER 11

Mixed Almost Invariant Sets when G is Finitely Generated

We can also construct minimal cubings for finite collections of “mixed” almost invariant sets, where some of the almost invariant sets have finitely generated stabilizers, and the rest arise from splittings.

Theorem 11.1. *Let G be a finitely generated group with any finite collection $\{X_i | i = 1, \dots, m\}$ of H_i -almost invariant subsets, where each H_i is finitely generated, each X_i is nontrivial, and $\bigcup_{i=1}^m \Sigma(X_i)$ is in very good position. Let $\{\sigma_j | j = 1, \dots, n\}$ be any finite collection of pairwise non-isomorphic splittings of G . For each j , let Y_j be a K_j -almost invariant set arising from σ_j , where each K_j is not finitely generated. Assume that $\{X_1, \dots, X_m, Y_1, \dots, Y_n\}$ satisfies sandwiching. Let $\Sigma := \bigcup_{i=1}^m \Sigma(X_i) \cup \bigcup_{j=1}^n \Sigma(Y_j)$. Then there exists a $CAT(0)$ cubical complex L , with a bijective correspondence between Σ and the set of oriented hyperplanes of L , such that two elements of Σ cross if, and only if, their corresponding hyperplanes cross.*

To prove Theorem 11.1, we will apply the proof of Theorem 8.7, with a few minor modifications.

Lemma 11.2. *Let Σ be as in Theorem 11.1. The relation \leq on Σ given by:*

$$A \leq B \iff A \subset B \text{ or } A \cap B^* \text{ is the only small corner of the pair } (A, B)$$

is well-defined, and partially orders Σ .

Proof. Since G is finitely generated, a corner of the pair (A, B) is $\text{Stab}(A)$ -finite if, and only if, the corner is $\text{Stab}(B)$ -finite of [23, Lemma 2.3]. Hence “smallness” of a corner is well-defined.

Next we claim that if a pair (A, B) has two small corners, then one is empty. Suppose two corners of (A, B) are small. By Lemma 6.1, $\text{Stab}(A)$ and $\text{Stab}(B)$ are commensurable. Since finite index subgroups of finitely generated groups are finite generated, and since we are assuming that the H_i 's are finitely generated and the K_j 's are not finitely generated, this implies either $A, B \in \bigcup_{i=1}^m \Sigma(X_i)$ or $A, B \in \bigcup_{i=1}^n \Sigma(Y_j)$. In the first case, since $\bigcup_{i=1}^m \Sigma(X_i)$ is in very good position, we must have an empty corner of the pair (A, B) . In the second case, since no two Y_j 's yield isomorphic splittings, we must have $A, B \in \Sigma(Y_j)$ for the same j , and hence one corner of (A, B) is empty.

Finally we show that \leq defines a partial order on Σ . Since $\bigcup_{i=1}^m \Sigma(X_i)$ is in very good position, The relation \leq , when restricted to $\bigcup_{i=1}^m \Sigma(X_i)$, is identical to inclusion. Now the proof of Proposition 6.5 shows that \leq defines a partial order on all of Σ . \square

Lemma 11.3. *Let Σ be as in Theorem 11.1. For all $A, B \in \Sigma$, there are only finitely many $C \in \Sigma$ such that $A \leq C \leq B$.*

Proof. Since G and all the H_i are finitely generated, [24, Lemma 1.15] shows that for all $A, B \in \Sigma(X_i)$, there are only finitely many $C \in \Sigma(X_i)$ such that $A \leq C \leq B$. Since each Y_j comes from a splitting, it follows that for all $A, B \in \Sigma(Y_j)$, there are only finitely many $C \in \Sigma(X_i)$ such that $A \leq C \leq B$.

By the proof of Proposition 7.4, for all $A, B \in \Sigma$, only finitely many $C \in \Sigma$ satisfy

$A \leq C \leq B$. □

To prove Theorem 11.1, apply the construction laid out in Chapter 8.4. The only other modification needed is to note that since $\bigcup_{i=1}^m \Sigma(X_i)$ is in very good position, each $\Sigma(X_i)$ is nested. This guarantees that the A_k 's – from part (5) of the proof that V in Theorem 8.11 satisfies DCC – are nested.

Corollary 11.4. *Let G be a finitely generated group with any finite collection $\{X_i | i = 1, \dots, m\}$ of H_i -almost invariant subsets, where each H_i is finitely generated, each X_i is nontrivial, and $\bigcup_{i=1}^m \Sigma(X_i)$ is in very good position. Let $\{\sigma_j | j = 1, \dots, n\}$ be any finite collection of pairwise non-isomorphic splittings of G . For each j , let Y_j be a K_j -almost invariant set arising from σ_j , where each K_j is not finitely generated. Assume that $\{X_1, \dots, X_m, Y_1, \dots, Y_n\}$ satisfies sandwiching. Let $\Sigma := \bigcup_{i=1}^m \Sigma(X_i) \cup \bigcup_{j=1}^n \Sigma(Y_j)$. Then we can put $\{X_1, \dots, X_m, Y_1, \dots, Y_n\}$ in very good position.*

I now briefly describe how to use Theorem 11.1 to construct an algebraic regular neighborhood of any finite collection $\{X_i | i = 1, \dots, m\}$ of H_i -almost invariant subsets in good position, where G and all the H_i 's are finitely generated. For the definition of an algebraic regular neighborhood, see [25, Definition 6.1]. First apply [18, Section 3] to construct a minimal cubing for the X_i 's, and to put them in very good position. Note that the existence of a minimal cubing for the X_i 's implies that the X_i 's satisfy sandwiching. Next, apply Theorem 11.1 to produce a minimal cubing for $\{X_i | i = 1, \dots, m\}$, and construct T from Chapter 9. To prove that T is an algebraic regular neighborhood of the X_i 's, note that, with the exception of the proof of Lemma 9.3, all the arguments from the proof of Theorem 9.1 still work. To prove the appropriate analogue of Lemma 9.3, let σ be any splitting of G having intersection number zero with each X_i , and let X be an H -almost invariant set arising

from σ . Then apply Theorem 11.1 to $\{X_i | i = 1, \dots, m\} \cup \{X\}$.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] B. H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.
- [2] D. E. Cohen. Ends and free products of groups. *Math. Z.*, 114:9–18, 1970.
- [3] M. J. Dunwoody. Accessibility and groups of cohomological dimension one. *Proc. London Math. Soc. (3)*, 38(2):193–215, 1979.
- [4] M. J. Dunwoody and M. E. Sageev. JSJ-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.
- [5] M. J. Dunwoody and E. L. Swenson. The algebraic torus theorem. *Invent. Math.*, 140(3):605–637, 2000.
- [6] H. Freudenthal. Über die Enden topologischer Räume und Gruppen. *Math. Z.*, 33(1):692–713, 1931.
- [7] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.
- [8] V. Guirardel. Cœur et nombre d’intersection pour les actions de groupes sur les arbres. *Ann. Sci. École Norm. Sup. (4)*, 38(6):847–888, 2005.
- [9] V. Guirardel and G. Levitt. JSJ decompositions: definitions, existence, uniqueness. I: The JSJ deformation space. arXiv:0911.3173, 2009.
- [10] V. Guirardel and G. Levitt. JSJ decompositions: definitions, existence, uniqueness. II: Compatibility and acylindricity. arXiv:1002.4564, 2009.
- [11] G. Higman, B. H. Neumann, and H. Neumann. Embedding theorems for groups. *J. London Math. Soc.*, 24:247–254, 1949.
- [12] H. Hopf. Enden offener Räume und unendliche diskontinuierliche Gruppen. *Comment. Math. Helv.*, 16:81–100, 1944.
- [13] C. H. Houghton. Ends of locally compact groups and their coset spaces. *J. Austral. Math. Soc.*, 17:274–284, 1974. Collection of articles dedicated to the memory of Hanna Neumann, VII.
- [14] W. H. Jaco and P. B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Amer. Math. Soc.*, 21(220):viii+192, 1979.
- [15] K. Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.
- [16] P. H. Kropholler. An analogue of the torus decomposition theorem for certain Poincaré duality groups. *Proc. London Math. Soc. (3)*, 60(3):503–529, 1990.
- [17] H. Müller. Decomposition theorems for group pairs. *Math. Z.*, 176(2):223–246, 1981.

- [18] G. Niblo, M. Sageev, P. Scott, and G. A. Swarup. Minimal cubings. *Internat. J. Algebra Comput.*, 15(2):343–366, 2005.
- [19] G. A. Niblo. A geometric proof of Stallings’ theorem on groups with more than one end. *Geom. Dedicata*, 105:61–76, 2004.
- [20] E. Rips and Z. Sela. Cyclic splittings of finitely presented groups and the canonical JSJ decomposition. *Ann. of Math. (2)*, 146(1):53–109, 1997.
- [21] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.
- [22] O. Schreier. Die untergruppen der freien gruppen. *Abh. Mat Sem. Univ. Hamburg*, 3:167–169, 1927.
- [23] P. Scott. The symmetry of intersection numbers in group theory. *Geom. Topol.*, 2:11–29 (electronic), 1998.
- [24] P. Scott and G. A. Swarup. Splittings of groups and intersection numbers. *Geom. Topol.*, 4:179–218 (electronic), 2000.
- [25] P. Scott and G. A. Swarup. Regular neighbourhoods and canonical decompositions for groups. *Astérisque*, (289):vi+233, 2003.
- [26] P. Scott and G. A. Swarup. Errata for ‘regular neighbourhoods and canonical decompositions for groups’. <http://www.math.lsa.umich.edu/~pscott/preprints.html>, October 2006.
- [27] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge, 1979.
- [28] Z. Sela. Structure and rigidity in (Gromov) hyperbolic groups and discrete groups in rank 1 Lie groups. II. *Geom. Funct. Anal.*, 7(3):561–593, 1997.
- [29] J.-P. Serre. *Arbres, amalgames, SL_2* . Société Mathématique de France, Paris, 1977. Avec un sommaire anglais, Rédigé avec la collaboration de Hyman Bass, Astérisque, No. 46.
- [30] J.-P. Serre. *Trees*. Springer-Verlag, Berlin, 1980. Translated from the French by John Stillwell.
- [31] E. Specker. Endenverbände von Räumen und Gruppen. *Math. Ann.*, 122:167–174, 1950.
- [32] J. R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.
- [33] J. R. Stallings. Group theory and 3-manifolds. In *Actes du Congrès International des Mathématiciens (Nice, 1970), Tome 2*, pages 165–167. Gauthier-Villars, Paris, 1971.
- [34] R. G. Swan. Groups of cohomological dimension one. *J. Algebra*, 12:585–610, 1969.
- [35] F. Waldhausen. On the determination of some bounded 3-manifolds by their fundamental groups alone. *Proc. Sympos. on Topology and its Applications, Beograd*, pages 331–332, 1969.