The coarse geometry of group splittings



Alexander Margolis
University College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2018

To my fiancée Shayna
for her unwavering love and support

Acknowledgements

First and foremost, I would like to thank my supervisor Panos Papasoglu for his help and advice. Some of the underlying themes and questions addressed in this thesis were first introduced to me in his *Geometric Group Theory* course that I had the pleasure of taking before starting my DPhil. Throughout my doctorate, he has suggested several interesting projects and ideas, sharing his expertise and knowledge. He has offered much invaluable advice, and made several helpful comments on my work. His mathematical influence on me will undoubtedly extend many years into the future.

I am grateful to have learnt from and worked with many wonderful mathematicians over the past four years, both at the Mathematical Institute in Oxford and in the wider geometric group theory community.

I would like to thank Bruce Kleiner for allowing me to use results contained in his unpublished manuscript [Kle02], which I give an account of in Section 3.2.2. I would also like to thank Jingyin Huang and Chris Cashen for reading a preliminary version of [Mar18a], the paper on which Chapter 5 is based, as well as Martin Dunwoody for a helpful discussion concerning the Kropholler conjecture.

I am extremely grateful to the EPSRC for funding my research. Given the worryingly large number of cuts in funding across the broad spectrum of higher education over the past eight years, I feel very fortunate to have been given the opportunity to undertake my DPhil.

Finally, I would like to thank all my friends and family for their encouragement and support, particularly my parents Andrew and Caroline and my fiancée Shayna.

Abstract

This thesis addresses Gromov's program of studying the geometry of finitely generated groups up to quasi-isometry. In particular, we investigate when we can characterise group splittings geometrically, and how this can be used to determine whether groups are quasi-isometric.

In Chapters 2 and 3 we develop several coarse topological techniques, building on work of Kapovich–Kleiner [KK05], that we expect to be of interest in their own right. We use these techniques to show that if G is a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is coarsely 3-separated by an essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ space W, then G splits over a subgroup H that is at finite Hausdorff distance from W. This can used to deduce that splittings of the form $G = A *_H B$, where G is of type $FP_{n+1}^{\mathbb{Z}_2}$ and H is a coarse $PD_n^{\mathbb{Z}_2}$ group such that both $|\mathrm{Comm}_A(H) : H|$ and $|\mathrm{Comm}_B(H) : H|$ are greater than two, are invariant under quasi-isometry. These results are contained in [Mar18b].

In Chapter 4 we extend the methods of Chapter 3 to understand the 2-separating case. Building on work of Scott–Swarup [SS03], we describe a *regular neighbourhood* of essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ spaces.

In Chapter 5 we show that the JSJ tree of cylinders over abelian subgroups of a RAAG is a quasi-isometry invariant. This provides a necessary condition for two RAAGs to be quasi-isometric. We then restrict to the class of RAAGs whose JSJ trees of cylinders over infinite cyclic subgroups have abelian rigid vertex stabilizers. We show that two such RAAGs are quasi-isometric if and only if they have equivalent JSJ trees of cylinders. The latter result is a special case of results we prove in [Mar18a].

Statement of Originality

This thesis is entirely my own work. Parts of this thesis have appeared in the paper [Mar18b] and the preprint [Mar18a]. All results due to mathematicians other than myself are acknowledged as such.

Alex Margolis
Oxford, 2018

Contents

1 Introduction **1**

2 Methods in coarse geometry and topology **15**

2.1 Coarse geometry 15

2.1.1 Coarse maps 16

2.1.2 Finitely generated groups as geometric objects 18

2.1.3 Coarse acyclicity 20

2.2 Metric complexes and coarse cohomology 22

2.2.1 Metric complexes 24

2.2.2 Uniformly acyclic complexes 29

2.2.3 Coarse cohomology 34

2.2.4 Roe's coarse cohomology 36

2.2.5 Group cohomology and finiteness properties of groups 38

2.2.6 Topology at infinity 38

2.3 Coarse Poincaré duality 40

2.3.1 Coarse Poincaré duality spaces 40

2.3.2 Poincaré duality groups 42

2.3.3 A technical lemma 43

3 Coarse detection of group splittings **49**

3.1 Coarse separation 49

3.1.1 Coarse complementary components 49

3.1.2 Coarse separation 53

3.1.3 Relative ends of groups 54

3.1.4 A quantitative coarse Mayer–Vietoris sequence 57

3.2 Constructing splittings 63

3.2.1 Essential components 63

3.2.2 Kleiner's mobility sets 69

3.2.3 Proof of Theorem 1.1 71
3.2.4 Applications 74

4 The geometry of JSJ decompositions 78
4.1 Groups acting on trees 81
4.2 A pretree structure on coarse complementary components 87
4.2.1 Groups with infinitely many ends 94
4.2.2 Bowditch's JSJ tree 95
4.2.3 The regular neighbourhood of Scott-Swarup 96
4.2.4 Coarse separation by quasi-lines 97
4.3 The regular neighbourhood 97
4.3.1 Construction of the regular neighbourhood 100
4.3.2 Quasi-isometric rigidity of the regular neighbourhood 108
4.3.3 A toolkit for computing the regular neighbourhood 111
4.4 More refined data from the JSJ decomposition 115
4.4.1 Relative quasi-isometries 116
4.4.2 Decorated trees 120

5 Quasi-isometry classification of certain RAAGs 126
5.1 Preliminaries on RAAGs 126
5.2 The JSJ tree of cylinders of a RAAG 128
5.3 Quasi-isometries of free products 135
5.4 Relative quasi-isometry type of cylinders 145
5.5 Constructing the quasi-isometry 148
5.6 The quasi-isometry problem 152
5.7 n-clique tree-graded RAAGs 154

Bibliography 157

Index 166

List of Figures

1.1 Essential and non-essential complementary components of a line . . . 4
1.2 Essential and non-essential complementary components of a plane . . 4
1.3 RAAGs that don't have equivalent JSJ trees of cylinders, hence aren't q.i. 11
1.4 A 3-clique tree-graded graph 12
1.5 RAAGs that do have equivalent JSJ trees of cylinders, yet aren't q.i. 13

2.1 A commutative diagram relating the topology at infinity of X and Y 44

3.1 The maps used to construct the coarse Mayer–Vietoris sequence . . . 61
3.2 A commutative diagram induced by f 64

5.1 Examples of visual graph of groups decompositions of a RAAG 129
5.2 The graph of groups associated to a separating clique 130
5.3 Schematic diagram of a free product of graphs 136
5.4 The graph of groups $\mathcal{G}$ 137
5.5 An example of a type I move 138
5.6 An example of a type II move 140
5.7 An example of a type III move 142
5.8 The decomposition of star(v) into coned-off blocks 145
5.9 The defining graphs of RAAGs that are not quasi-isometric 146

Chapter 1

Introduction

One of the fundamental questions of geometric group theory, originally posed by Gromov [Gro93], is the following: what properties of a finitely generated group are invariant under quasi-isometry? We say that a group G *splits* over a subgroup H if either G is an amalgamated free product $A *_H B$, where H is a proper subgroup of both A and B, or G is an HNN extension of the form $G = A*_H$.

Quasi-isometry invariance of splittings

In this thesis, we show that group splittings can often be detected from the large-scale geometry of a group. This can be used to determine when two groups are quasi-isometric. The most celebrated result along these lines is the following theorem of Stallings, providing a correspondence between the number of ends of a group — a large-scale geometric property, and whether a group splits over a finite subgroup — an algebraic property.

Theorem ([Sta68], [Sta71])**.** *A finitely generated group splits over a finite subgroup if and only if it has more than one end.*

Along with a theorem of Dunwoody [Dun85], Stallings' theorem allows us to decompose groups into 'smaller' pieces via graph of groups decompositions. A result of Papasoglu and Whyte [PW02] classifies finitely presented groups up to quasi-isometry in terms of their one-ended vertex groups. Papasoglu [Pap05] gives a geometric characterisation of splittings over two-ended groups, generalising Stallings' theorem. Generalising Stallings' theorem and Papasolgu's theorem to splittings over more complicated groups allows us to better understand the structure of groups up to quasi-isometry.

Consider the Cayley graph Γ of a group G with respect to some finite generating set. We say that $C \subseteq G$ is a *coarse complementary component* of $W \subseteq G$ if for

some $R \geq 0$, $C \backslash N_R(W)$ is the vertex set of a union of components of $\Gamma \backslash N_R(W)$. The motivation behind this definition is that quasi-isometries preserve unions of complementary components, but do not necessarily preserve a single complementary component. Thus the notion of coarse complementary components is a natural one when working with quasi-isometries.

A coarse complementary component is said to be *deep* if it is not contained in $N_R(W)$ for any $R \geq 0$; otherwise it is said to be *shallow*. A collection of deep coarse complementary components of W is said to be *coarse disjoint* if the intersection of any pair is shallow. We say that W *coarsely n-separates* G if there exist n deep, coarse disjoint, coarse complementary components of W. We say that W *coarsely separates* G if W coarsely 2-separates G. Coarse n-separation is a quasi-isometry invariant. The following proposition relates coarse separation to group splittings.

Proposition ([Pap12, Lemma 2.2])**.** *If a finitely generated group G splits over a finitely generated subgroup H, then H coarsely separates G.*

A group has more than one end if and only if it is coarsely separated by a point. Consequently, Stallings' theorem may be rephrased as follows:

Theorem. *A finitely generated group splits over a finite subgroup if and only if it is coarsely separated by a point.*

We prove a partial generalisation of Stallings' theorem, giving a large-scale geometric criterion that guarantees the existence of a group splitting. We use this to show that group splittings are often invariant under quasi-isometry. Triangle groups provide examples of groups which don't split, but have finite index subgroups that do split over two-ended subgroups; hence admitting a splitting over a two-ended group is not invariant under quasi-isometry. However, there is a theorem by Papasoglu which partially generalises Stallings' theorem. In the following theorem, a *line* is defined to be a coarsely embedded copy of $\mathbb{R}$.

Theorem ([Pap05])**.** *Let G be a finitely presented one-ended group that is not virtually a surface group. A line coarsely separates G if and only if G splits over a two-ended subgroup.*

This is known to be false if one drops the condition that G is finitely presented. In [Pap12], Papasoglu constructs a line that coarsely separates the lamplighter group — a finitely generated group that is not finitely presented and doesn't split over a two-ended subgroup.

To construct a splitting using the geometry of a group, we first construct a subgroup which coarsely separates the group, and then show that the group splits over a subgroup commensurable to it. Much work has already been done on the second step, for example see [DS00].

Before stating our results, we need to define a few terms. For a ring R, a group G is said to be of *type* FP_n^R if it admits a partial projective resolution

$$P_n \to P_{n-1} \to \cdots \to P_0 \to R \to 0$$

of the trivial RG-module R, such that each P_i is finitely generated as an RG-module. A group is of type F_n if it has a classifying space with finite n-skeleton. If G is of type F_n, it is of type FP_n^R for any R. These are examples of *finiteness properties*, generalising the notions of being finitely generated and finitely presented.

We will work with groups of type $FP_n^{\mathbb{Z}_2}$. The use of $\mathbb{Z}_2$ coefficients is fairly natural in the context of group splittings and coarse separation. For example, the number of ends can be detected using cohomology with $\mathbb{Z}_2$ coefficients [SW79] and Dunwoody's accessibility theorem holds for groups of type $FP_2^{\mathbb{Z}_2}$ [Dun85]. We remark that a group of type $FP_n^{\mathbb{Z}}$ (often simply denoted as FP_n) is necessarily of type $FP_n^{\mathbb{Z}_2}$.

Stallings' theorem holds for finitely generated groups, which are necessarily of type $FP_1^{\mathbb{Z}_2}$, whereas Papasoglu's theorem only holds for finitely presented groups, which are necessarily of type $FP_2^{\mathbb{Z}_2}$. This motivates the principle that when examining splittings over groups that are more complicated than finite and two-ended groups, we assume the ambient group has higher finiteness properties.

Coarse $PD_n^{\mathbb{Z}_2}$ spaces are defined in [KK05]. They are spaces that have the same large-scale homological properties as $\mathbb{R}^n$. For example, the universal cover of a closed aspherical n-dimensional manifold is a coarse $PD_n^{\mathbb{Z}_2}$ space. To construct a splitting, we make the assumption that a coarse $PD_n^{\mathbb{Z}_2}$ space coarsely separates a group of type $FP_{n+1}^{\mathbb{Z}_2}$.

When dealing with the case of a line coarsely separating a group as in Papasoglu's theorem, as well as similar results such as Bowditch's JSJ tree [Bow98], one needs a geometric criterion for recognizing virtually surface groups. In doing this, one uses the Tukia, Gabai and Casson-Jungreis theorem on convergence groups acting on the circle ([Tuk88], [Gab92] and [CJ94]). Since this theorem has no analogue in higher dimensions, we cannot rule out generalisations of triangle groups; we therefore make the assumption that a coarse $PD_n^{\mathbb{Z}_2}$ space coarsely 3-separates a group. Unfortunately, even this assumption is not sufficient for our purposes. Therefore, rather than working with

deep components, we use what we call *essential components*. This is a generalisation of the essential components found in [Pap05].

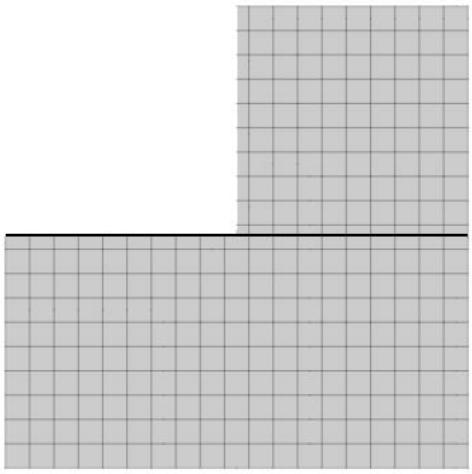


Figure 1.1: An example of essential and non-essential complementary components of a line

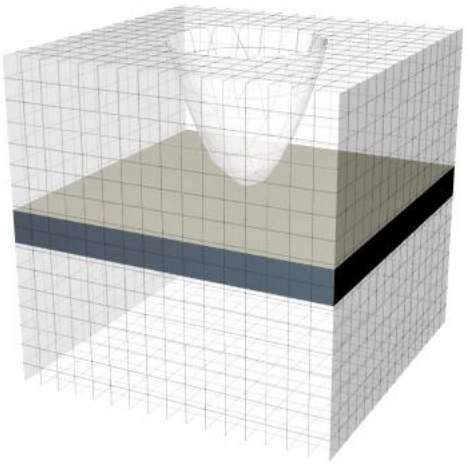


Figure 1.2: An example of essential and non-essential complementary components of a plane

Essential components are necessarily deep and are invariant under quasi-isometry. An example is that of a coarse $PD_n^{\mathbb{Z}_2}$ space which coarsely separates a coarse $PD_{n+1}^{\mathbb{Z}_2}$ space into two coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-spaces — see the coarse Jordan separation theorem of [KK05]. One can think of a copy of $\mathbb{R}^n$ coarsely separating $\mathbb{R}^{n+1}$ into two half-spaces. Each coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-space is essential. More generally, any coarse complementary component that contains such a half-space is necessarily essential. The definition of essential components is rather technical, so we do not provide it here. However, we do give some examples of essential and non-essential components.

In Figure 1.1, a line coarsely separates the space into two deep components. In Figure 1.2, a plane coarsely separates the space into two deep components. In both cases, the bottom component is essential and the top component is not. These examples illustrate the only ways in which a coarse complementary component fails to be essential. In Figure 1.1, the boundary of the top component is not the entire line, but only half the line. It will be shown in Proposition 3.30 that this cannot occur for essential components. In Figure 1.2, the boundary of the top component is the entire plane, so the component fails to be essential in a more subtle way. It is not essential because it contains a non-trivial 1-cycle at infinity — a hole that cannot be filled as we move away from the plane.

We say that $W \subseteq G$ is *essentially embedded* if every deep coarse complementary component of W is essential. We are now in a position to state our main theorem; all the hypotheses of the theorem are invariant under quasi-isometry. This is a special case of a slightly more general result — Theorem 3.42 — which we prove in Chapter 3.

Theorem 1.1. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and let $W \subseteq G$ be an essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ subspace that coarsely 3-separates G. Then there exists a subgroup $H \leq G$, contained in $N_R(W)$ for some $R \geq 0$, such that G splits over H.*

We now discuss some consequences of Theorem 1.1 to demonstrate its applicability to situations in which the essentially embedded condition is not a priori known to hold. In deciding whether a group splits over a certain class of subgroups, e.g. virtually $\mathbb{Z}^n$ subgroups, it simplifies the argument if we make the natural assumption that it does not split over a 'smaller' class of subgroups, e.g. virtually $\mathbb{Z}^r$ subgroups for $r < n$. For instance in Papasoglu's theorem, which determines if a group splits over a two-ended group, we assume that the ambient group is one-ended; therefore, Stallings' theorem says it cannot split over a finite subgroup.

We want to find a higher dimensional analogue of a space being one-ended that rules out splittings over certain classes of subgroups. The right generalisation of one-endedness is acyclicity at infinity over $\mathbb{Z}_2$, which will be defined in Section 2.2.6. If a group is of type $FP_{n+1}^{\mathbb{Z}_2}$ and is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$, then it cannot split over a virtually $\mathbb{Z}^r$ subgroup for any $r < n$.

A finitely generated group G has one end if and only if it is 0-acyclic at infinity over $\mathbb{Z}_2$. If G is the fundamental group of a closed aspherical n-manifold, or more generally is a coarse $PD_n^{\mathbb{Z}_2}$ space, then it is $(n-2)$-acyclic at infinity over $\mathbb{Z}_2$. A group is a *coarse $PD_n^{\mathbb{Z}_2}$ group* if, when equipped with the word metric with respect to some finite generating set, it is a coarse $PD_n^{\mathbb{Z}_2}$ space. We prove the following criterion which determines when a coarse $PD_n^{\mathbb{Z}_2}$ subgroup is essentially embedded.

Proposition 3.34. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$ and let $H \leq G$ be a coarse $PD_n^{\mathbb{Z}_2}$ group. Then H is essentially embedded if and only if no infinite index subgroup of H coarsely separates G.*

Combining this with Theorem 1.1 and the observation that the hypotheses of Theorem 1.1 are invariant under quasi-isometry, we deduce the following:

Theorem 1.2. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$. Suppose $H \leq G$ is a coarse $PD_n^{\mathbb{Z}_2}$ group that coarsely 3-separates G, and no infinite index subgroup of H coarsely separates G. Then for any quasi-isometry*

$f : G \to G'$ there is a subgroup $H' \leq G'$, at finite Hausdorff distance from $f(H)$, such that G' splits over H'.

It is shown in [BM00] that if G is the fundamental group of a finite graph of groups whose vertex and edge groups satisfy appropriate finiteness and acyclicity at infinity conditions, then so does G. We therefore deduce the following:

Corollary 1.3. *Suppose $G = A *_H B$ (or $G = A*_H$) is a splitting, where H is a coarse $PD_n^{\mathbb{Z}_2}$ group, A and B are of type $FP_{n+1}^{\mathbb{Z}_2}$ and are $(n-1)$-acyclic over infinity over $\mathbb{Z}_2$, H coarsely 3-separates G, and no infinite index subgroup of H coarsely separates G. Then for any quasi-isometry $f : G \to G'$ there is a subgroup $H' \leq G'$, at finite Hausdorff distance from $f(H)$, such that G' splits over H'.*

This corollary is particularly useful when combined with Theorem 8.7 of [Vav12], which gives an algebraic characterisation of when the 3-separating hypothesis of Corollary 1.3 holds.

There are several examples of groups that are acyclic at infinity. Say G is the extension of N by Q, where N and Q are groups of type F_{n+1} that are r and s-acyclic at infinity over $\mathbb{Z}_2$ respectively. Then Theorem 17.3.6 in [Geo08] tells us that G is $\min(n, s + r + 2)$-acyclic at infinity over $\mathbb{Z}_2$. For example, if N and Q are finitely presented one-ended groups, then G is 1-acyclic at infinity over $\mathbb{Z}_2$. Results from [BM01] and [DM02] give conditions for Coxeter and right-angled Artin groups to be acyclic at infinity. These results allow us to apply Theorem 1.2 and Corollary 1.3.

Theorem 1.2 can be simplified if H is a virtually polycyclic group and thus necessarily a coarse $PD_n^{\mathbb{Z}_2}$ group. In this case, one can drop the condition that no infinite index subgroup of H coarsely separates G since it is implied by G being $(n-1)$-acyclic at infinity. For ease of notation, we say a group is *VPC-n* if it is virtually polycyclic group of Hirsch length n.

Corollary 3.43. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$. Suppose $H \leq G$ is a VPC-n subgroup that coarsely 3-separates G. Then for any quasi-isometry $f : G \to G'$, there is a subgroup $H' \leq G'$, at finite Hausdorff distance from $f(H)$, such that G' splits over H'.*

Corollary 3.43 may be coupled with the quasi-isometric rigidity of virtually $\mathbb{Z}^n$ groups. If a group G is of type $FP_{n+1}^{\mathbb{Z}_2}$, is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$ and is coarsely 3-separated by a virtually $\mathbb{Z}^n$ group, then any group quasi-isometric to G splits over a virtually $\mathbb{Z}^n$ subgroup. A similar statement holds for virtually nilpotent groups.

Two subgroups $H, K \leq G$ are *commensurable* if $H \cap K$ has finite index in both H and K. The *commensurator* of H is the subgroup

$$\mathrm{Comm}_G(H) := \{g \in G \mid H \text{ and } g^{-1}Hg \text{ are commensurable}\}.$$

We deduce the following from Theorem 1.1:

Theorem 3.44. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is the fundamental group of a finite graph of groups $\mathcal{G}$. Suppose $\mathcal{G}$ contains an edge e with associated edge monomorphisms $i_0 : G_e \to G_v$ and $i_1 : G_e \to G_w$ such that the following holds: G_e is a coarse $PD_n^{\mathbb{Z}_2}$ group and $|\mathrm{Comm}_{G_v}(i_0(G_e)) : i_0(G_e)|$ and $|\mathrm{Comm}_{G_w}(i_1(G_e)) : i_1(G_e)|$ are both greater than one and not both equal to two. If $f : G \to G'$ is a quasi-isometry, then $f(G_e)$ has finite Hausdorff distance from some subgroup $H' \leq G'$, and G' splits over a subgroup commensurable to a subgroup of H'.*

Our work builds on results by Vavrichek [Vav12], Mosher–Sageev–Whyte in [MSW03] and [MSW11] and Papasoglu in [Pap07]. We remark that one of the conditions Vavrichek uses is the *non-crossing* condition, which is not known to be invariant under quasi-isometry. Under our hypotheses, we obtain the non-crossing condition automatically (see Lemma 3.35).

We now consider the situation where we remove the 3-separating hypothesis. It is shown in [DR93] that if $G = \mathrm{Comm}_G(H)$ and H coarsely separates G, then G splits over a subgroup commensurable to H. In the case where $G = \mathrm{Comm}_G(H)$, we obtain the following coarse geometric generalisation of this result.

Theorem 3.45. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and suppose G splits over a coarse $PD_n^{\mathbb{Z}_2}$ subgroup H and that $\mathrm{Comm}_G(H) = G$. Let $f : G \to G'$ be a quasi-isometry. Then one of the following occurs:*

(1) G' is a coarse $PD_{n+1}^{\mathbb{Z}_2}$ group;

(2) $f(H)$ has finite Hausdorff distance from some $H' \leq G'$, and G' splits over H'.

To deal with the general case, we recall the Algebraic Torus Theorem of Dunwoody–Swenson [DS00]. To do this, we need the notion of a *quadratically hanging vertex group* in a graph of groups.

Definition 1.4. A vertex group G_v in a graph of groups, whose incident edge groups are VPC-n, is said to be *quadratically hanging* if:

- there exists a VPC-$(n-1)$ normal subgroup $N \triangleleft G_v$ such that G_v/N is a Fuchsian group that is neither finite nor two-ended;
- incident edge groups $G_e \leq G_v$ contain N, and the quotients G_e/N correspond precisely to conjugates of peripheral subgroups of G_v/N.

Theorem ([DS00])**.** *Let G be a finitely generated group and $H \leq G$ be a VPC-n subgroup that coarsely separates G. Assume no subgroup $K \leq H$ of infinite index in H coarsely separates G. Then one of the following occurs:*

(1) G is VPC-$(n+1)$;

(2) G splits over a subgroup commensurable to H;

(3) G is the extension of a VPC-$(n-1)$ group by a cocompact Fuchsian group;

(4) G splits as a graph of groups with VPC-n edge groups and a quadratically hanging vertex group containing $H \leq G$.

Note that both the first and third cases of the Algebraic Torus Theorem are in some sense exceptional: in both cases the ambient group G is itself a coarse $PD^{\mathbb{Z}_2}_{n+1}$ group. We prove the following coarse analogue of the Algebraic Torus Theorem.

Corollary 4.4. *Let G be a group of type $FP^{\mathbb{Z}_2}_{n+1}$ that is coarsely 2-separated by an essentially embedded coarse $PD^{\mathbb{Z}_2}_n$ subspace W. Then either:*

(1) G is a coarse $(n+1)$-manifold group;

(2) G splits over a subgroup at finite Hausdorff distance from W;

(3) G acts non-trivially (i.e. not fixing a vertex) and minimally on a tree T, containing a vertex $v \in T$ enclosing W, where $G_v = \mathrm{Stab}([\alpha])$ for some non-trivial cohomology class $[\alpha] \in H^{n+1}(G, \mathbb{Z}_2 G)$.

We refer the reader to Chapter 4 for definitions of enclosing vertices and coarse $(n+1)$-manifold groups. Corollary 4.4 will be deduced from Theorem 4.1, which describes a tree that encloses any finite collection $\{W_1, \dots W_m\}$ of essentially embedded coarse $PD^{\mathbb{Z}_2}_n$ subspaces. We remark that the vertex v in case (3) generalises quadratically hanging vertices. The cohomology class $[\alpha]$ can be thought of as a coarse fundamental class, and the associated vertex space is in some sense "$(n+1)$-manifold like".

We remark that Corollary 4.4 gives new information even in the algebraic situation when W is a subgroup, e.g. where G is a hyperbolic group with simply connected boundary containing a codimension one surface subgroup $W \leq G$. Such a theorem may prove useful in developing a JSJ theory over non-slender subgroups.

Quasi-isometry classification of certain RAAGs

Given a finite simplicial graph Γ, the *right-angled Artin group* (RAAG) $A(\Gamma)$ is a group with the following presentation: generators correspond to vertices of Γ, and relations are commutators $[v, w]$ for vertices v and w that are joined by an edge. RAAGs have become interesting objects of study in recent years due to their rich subgroup structure, e.g. [BB97], [HW08]. The work of Agol and Wise uses RAAGs to understand 3-manifold groups [Wis09], [Ago13]. See the survey article [Cha07] for more information about RAAGs.

We are interested in classifying which RAAGs are quasi-isometric to one another. Much progress has been made on this problem and generalisations of it in recent years: see [BN08], [BKS08], [BJN10], [Hua17a], [Hua15], [HK18], [Hua18] and [HO17].

There are two very different approaches to solving the above problem:
Method 1: One identifies a class of RAAGs and shows that quasi-isometries amongst these RAAGs preserve maximal flats. One then shows that quasi-isometries of RAAGs in this class preserve a large amount of combinatorial data of the RAAG, and uses this data to deduce that RAAGs in this class are *quasi-isometrically rigid*, i.e. two RAAGs are quasi-isometric if and only if they are commensurable. This is the principle used in [BKS08], [Hua17a] and [Hua15]. It is loosely modelled on the quasi-isometric rigidity of higher rank symmetric spaces and mapping class groups.
Method 2: One identifies a class of RAAGs that have graph of group decompositions into smaller RAAGs whose quasi-isometry type is known, e.g. RAAGs that have graph of group decompositions with infinite cyclic edge groups and $\mathbb{Z}^2$ vertex groups. One shows that such a decomposition completely determines the quasi-isometry type of a RAAG. This is the principle used in [BN08] and [BJN10]. There are many RAAGs to which Method 2 can be applied that are quasi-isometric but are not commensurable.

The techniques required for Method 1 to work in [BKS08], [Hua17a] and [Hua15] completely break down if a RAAG contains a non-adjacent transvection; other approaches are needed to determine when RAAGs containing non-adjacent transvections are quasi-isometric. In this thesis, following Method 2 as above, we classify the quasi-isometry types of certain RAAGs via their JSJ trees of cylinders. This determines the quasi-isometry type of many RAAGs that cannot be determined by any other known method. This is a generalisation of the classification of tree RAAGs in [BN08].

It is easy to see that every RAAG can be built up from the trivial group via iterated amalgamated free products and HNN extensions. This motivates Method 2 as a natural step in the quasi-isometry classification of RAAGs. In this thesis, we examine the case where edge groups in the graph of groups decomposition are abelian.

The JSJ tree of cylinders of a finitely presented group G is a canonical tree T that encodes information about all splittings of G over a certain class of groups. Each vertex of T is one of three *types*: cylindrical, hanging or rigid. A *relative quasi-isometry* between vertex stabilisers of the JSJ tree of cylinders is a quasi-isometry that coarsely preserves the *peripheral structure* — the set of conjugates of incident edge stabilizers.

Definition. Let T and T' be two JSJ trees of cylinders of two finitely presented groups. We say that T and T' are *equivalent* if there exists a tree isomorphism $\chi : T \to T'$ such that the following hold:

(1) χ preserves vertex type (cylinder, hanging or rigid);

(2) for all $v \in VT$, there is a relative quasi-isometry between vertex stabilizers of v and $\chi(v)$, such that the induced map on the peripheral structure agrees with $\chi|_{\mathrm{lk}(v)}$.

Proposition 4.81 characterises when two JSJ trees of cylinders are equivalent in terms of their quotient graphs of groups. It is a special case of a more general statement of Cashen-Martin [CM17]. However, it is rather technical to state, so we defer doing so till Section 4.4.2. This characterisation is significantly easier to use in practice than directly constructing a map χ as above.

Theorem 1.5 (Theorem 5.17)**.** *Let $n > 0$. Suppose two RAAGs $A(\Gamma)$ and $A(\Lambda)$ don't split over any subgroup isomorphic to $\mathbb{Z}^k$ for all $k < n$. Then if $A(\Gamma)$ and $A(\Lambda)$ are quasi-isometric, they have equivalent JSJ trees of cylinders over subgroups isomorphic to $\mathbb{Z}^n$.*

By adapting the work of Clay [Cla14] and Groves–Hull [GH17], we give an explicit *visual* description of the JSJ tree of cylinders of a RAAG $A(\Gamma)$, i.e. a graph of groups that can be read off from the defining graph Γ. See Definition 5.12.

We note that in Theorem 1.5, the assumption that RAAGs don't split over $\mathbb{Z}^k$ for $k < n$ is not restrictive in practice. Firstly, a RAAG is one-ended if and only if its defining graph is connected and not a single vertex. Using the work of Papasoglu and Whyte [PW02], we note that if two RAAGs are not one-ended, they are quasi-isometric if and only if they have quasi-isometric one-ended factors in their Gruško decomposition. We thus reduce to RAAGs that are one-ended.

Having reduced to one-ended RAAGs, we apply Theorem 1.5 inductively. We first apply Theorem 1.5 with $n = 1$. This typically yields RAAG vertex groups that do not split over infinite cyclic subgroups, and we can thus apply Theorem 1.5 with $n = 2$

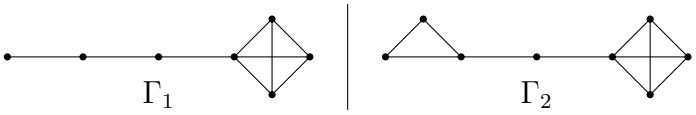


Figure 1.3: The defining graph of two RAAGs that don't have equivalent JSJ trees of cylinders, hence are not quasi-isometric. This example is also considered in [BHS17].

to determine if these vertex groups are quasi-isometric. We can continue in this way, inductively applying Theorem 1.5. This sort of argument gives a necessary condition for two RAAGs to be quasi-isometric.

Proposition 5.25 gives a criterion to determine whether cylindrical vertex stabilizers in the JSJ tree of cylinders of a RAAG are relatively quasi-isometric. Along with Proposition 4.81, this gives a powerful method that is frequently able to demonstrate that two RAAGs have JSJ trees of cylinders that are not equivalent, and hence the RAAGs are not quasi-isometric.

We are not aware of such a statement being used in the literature concerning quasi-isometries of RAAGs. For example, the RAAGs considered in the introduction to [BHS17], whose defining graphs are shown in Figure 1.3, can be shown not to be quasi-isometric using this method.

Indeed, consider the RAAGs with defining graphs shown in Figure 1.3. The JSJ tree of cylinders decomposition can be described using the graphs of groups $C(\Gamma_1)$ and $C(\Gamma_2)$, as in Definition 5.12. We note that $C(\Gamma_2)$ has a cylinder stabilizer isomorphic to $\mathbb{Z} \times (\mathbb{Z}^2 * \mathbb{Z})$. This can easily be seen not be quasi-isometric to any cylinder stabilizer of $C(\Gamma_1)$ — the details of this easily follow from Section 5.4. Thus the RAAGs $A(\Gamma_1)$ and $A(\Gamma_2)$ cannot have equivalent JSJ trees of cylinders, hence are not quasi-isometric.

We now restrict to the following class of RAAGs and give a necessary and sufficient condition for RAAGs in this class to be quasi-isometric.

Definition 1.6. We say a subgraph is *biconnected* if it is connected and has no cut vertex. A *block graph* is a finite connected simplicial graph such that every biconnected component is a clique. A RAAG whose defining graph is a block graph is called a *block graph RAAG*. Let $\mathcal{C}_{ab}$ denote the class of block graph RAAGs.

Theorem 5.28. *Two RAAGs in $\mathcal{C}_{\text{ab}}$ are quasi-isometric if and only if they have equivalent JSJ trees of cylinders.*

Using Theorem 5.28, we are also able to algorithmically decide if RAAGs in $\mathcal{C}_{\text{ab}}$ are quasi-isometric:

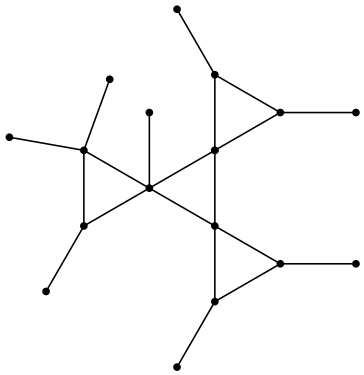


Figure 1.4: A 3-clique tree-graded graph

Theorem 5.34. *There is an algorithm that, when given as input finite simplicial graphs* Γ *and* Λ *such that* $A(\Gamma) \in \mathcal{C}_{\text{ab}}$*, determines whether* $A(\Gamma)$ *is quasi-isometric to* $A(\Lambda)$*.*

We note that the class $\mathcal{C}_{\text{ab}}$ includes the RAAGs whose defining graphs are shown in Figure 1.3, as well as tree RAAGs considered in [BN08]. As a sample application of Theorem 5.28, we give a classification for the following class of RAAGs.

Definition. Let Γ be a finite, connected, simplicial graph, and let C_0 denote the set of cut vertices of Γ. For $n \geq 2$, we say that Γ is an *n-clique tree-graded graph* if:

- every vertex of Γ is either contained in C_0 or has valence one;
- Γ has diameter at least three;
- each maximal biconnected subgraph is either an n-clique, all of whose vertices are in C_0, or a 2-clique containing exactly one vertex in C_0.

A simplicial graph is a 2-clique tree-graded graph if and only if it is a tree of diameter at least three. An example of a 3-clique tree-graded graph is shown in Figure 1.4. We say that $A(\Gamma)$ is an *n-clique tree-graded RAAG* whenever Γ is an n-clique tree-graded graph. The following theorem generalises results of Behrstock-Neumann [BN08], who prove the $n = 2$ case.

Theorem 5.36. *Fix some* $n \geq 2$*. All n-clique tree-graded RAAGs are quasi-isometric to one another, and every RAAG quasi-isometric to an n-clique tree-graded RAAG is itself an n-clique tree-graded RAAG.*

Though it will not be discussed in this thesis, we briefly mention that there exist RAAGs that are not quasi-isometric yet have JSJ trees of cylinders that are equivalent

in the sense considered above. The defining graphs of two such RAAGs are shown in Figure 1.5. By adapting the methods used in Chapter 5, we show in [Mar18a] that one can distinguish such RAAGs via their JSJ trees of cylinders by considering *stretch factors*. In [Mar18a], we generalise Theorem 5.28 to a much larger class of RAAGs by taking these stretch factors into account.

In contrast to the quasi-isometrically rigid classes of RAAGs considered in [BKS08], [Hua17a] and [Hua15], there are several RAAGs to which Theorem 5.28 can be applied that are quasi-isometric but not commensurable. In [CRKZ16], it is shown that there are infinitely many commensurability classes within the class of 2-clique tree-graded RAAGs, all of which lie in a single quasi-isometry class.

The quasi-isometry classification of certain right-angled Coxeter groups by Dani-Thomas [DT17] is very much in the spirit of Method 2. We remark that although every right-angled Artin group is a finite index subgroup of a right-angled Coxeter group [DJ00], the groups considered in [DT17] are one-ended and hyperbolic, so are not commensurable to any RAAG. Thus the results of [DT17] have no overlap with Theorem 5.28.

Some of the methods developed in Chapter 5 are of interest in their own right, and might be used in settings other than RAAGs. Studying the JSJ decompositions of groups is a natural approach to their quasi-isometry classification, and the techniques developed here may be used when the conditions of [CM17] are not known to hold. In particular, we do not assume rigid vertices are quasi-isometrically rigid, and we do not assume that cylinder stabilizers are two-ended.

Organisation of this thesis

In Chapter 2 we develop various tools in coarse geometry and topology. In Section 2.2, we use work of Kapovich and Kleiner [KK05] to define a notion of cohomology

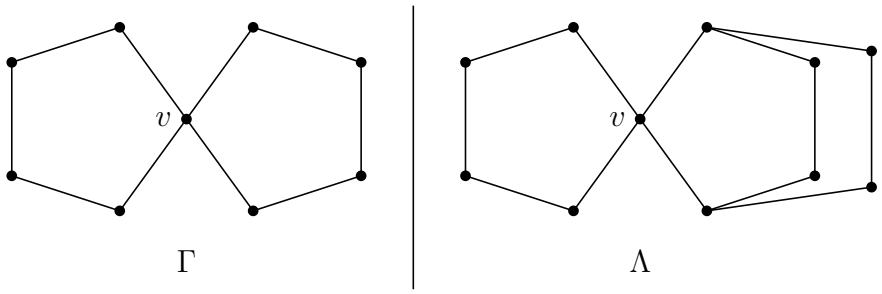


Figure 1.5: The defining graphs of RAAGs that are not quasi-isometric, yet have equivalent JSJ trees of cylinders

that is invariant under quasi-isometry. For finitely generated groups, this agrees with the cohomology with group ring coefficients. There is a more general notion of coarse cohomology due to Roe (see [Roe03]). However, our approach to coarse cohomology is more amenable to the quantitative methods essential in our work. Our approach to coarse cohomology makes use of the theory of *metric complexes* as defined in the appendix of [KK05]. In our exposition, metric complexes can be replaced with *bounded geometry CW complexes* as defined in [MSW03], with the caveat that this weakens our results from groups of type $FP_n^{\mathbb{Z}_2}$ to groups of type F_n. In Section 2.3, we discuss the notion of coarse $PD_n^{\mathbb{Z}_2}$ spaces as defined in [KK05]. We also prove a lemma about coarse $PD_n^{\mathbb{Z}_2}$ spaces that is needed to deduce Theorem 1.2.

In Chapter 3 we develop more coarse topological methods. In Section 3.1, we introduce coarse complementary components and describe the coarse Mayer–Vietoris sequence in coarse cohomology. In Section 3.2 we define essential components and prove Theorem 1.1 and its corollaries. We make use of the notion of a *mobility set* due to Kleiner, and contained in an unpublished manuscript [Kle02]. We give a self-contained exposition of the parts of [Kle02] that we use.

In Chapter 4 we switch our focus from trying to find a single splitting of a group, and instead try to understand all splittings of a group. In Section 4.2, we develop a general framework, using coarse complementary components, of encoding a collection of coarse complementary components into a combinatorial structure called a pretree. In Section 4.3, under suitable cohomological assumptions, we obtain a coarse analogue of the regular neighbourhood of Scott–Swarup [SS03].

In Chapter 5 we prove the aforementioned results concerning right-angled Artin groups. In Section 5.2 we give an explicit visual description the JSJ tree of cylinders of a RAAG and prove Theorem 5.17. In Section 5.5, we prove Theorem 5.28 giving a complete quasi-isometry classification of RAAGs in $\mathcal{C}_{\text{ab}}$.

Chapter 2

Methods in coarse geometry and topology

A common theme in recent years has been to view groups as geometric and topological objects. The geometric aspect of this was essentially started by pioneering work of Gromov, e.g. [Gro81] and [Gro93]. On the topological side, work of Hopf [Hop44] and Stallings [Sta68] describes aspects of the large scale topology of a finitely generated group. More recently, work of Farb–Schwartz [FS96] and Kapovich–Kleiner [KK05] develop methods of coarse algebraic topology. In this section and the next, we expound upon some of the methods and tools that are used to describe this large-scale geometry and topology.

2.1 Coarse geometry

Let (X, d) be a metric space. For $x \in X$ and $\emptyset \neq A \subseteq X$, we let $d(x, A) := \inf\{d(a, x) \mid a \in A\}$. We define $N_r^X(A) := \{x \in X \mid d(x, A) \leq r\}$. For each $x \in X$, we let $N_r^X(x) := N_r^X(\{x\})$, which is just the closed r ball around x. When unambiguous, we denote N_r^X by N_r. We say $Y \subseteq X$ is an (ϵ, δ)-*net* if $d(y, y') \geq \epsilon$ for all distinct $y, y' \in Y$, and if for every $x \in X$, there exists some $y \in Y$ such that $d(y, x) \leq \delta$.

If $A, B \subseteq X$ are non-empty, we define the *Hausdorff distance* to be

$$d_{\text{Haus}}(A, B) := \inf\{r \geq 0 \mid A \subseteq N_r(B) \text{ and } B \subseteq N_r(A)\}.$$

We are often only interested in whether $d_{\text{Haus}}(A, B)$ is finite. In such a case, we simply say that A and B are *coarsely equivalent*. Coarse equivalence is an equivalence relation among subsets of a given metric space. We let $[A]$ denote the coarse equivalence class containing A. Given $A, B \subseteq X$, we say that A is *coarsely contained* in B if for some $r \geq 0$, $A \subseteq N_r(B)$.

2.1.1 Coarse maps

The material here is fairly standard, see [Roe03] for more details. A function $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is *proper* if the inverse images of compact sets are compact.

Definition 2.1. Let (X, d_X) and (Y, d_Y) be metric spaces and $\eta, \phi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be proper non-decreasing functions. We say a map $f : X \to Y$ is an (η, ϕ)-*coarse embedding* if for all $x, y \in X$,

$$\eta(d_X(x, y)) \leq d_Y(f(x), f(y)) \leq \phi(d_X(x, y)).$$

We say that f is a *coarse embedding* if there exist proper non-decreasing functions η and ϕ such that f is an (η, ϕ)-coarse embedding. We say that η (resp. ϕ) is the *lower* (resp. *upper*) *distortion function* of f.

In the literature, coarse embeddings are also known as *uniform embeddings* or *uniformly proper embeddings*, sometimes with the additional assumption that the upper distortion function is affine.

Remark 2.2. For each proper non-decreasing function $\eta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, we define another proper non-decreasing function $\tilde{\eta} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ by $\tilde{\eta}(R) := \sup(\eta^{-1}([0, R]))$. We observe that whenever $\eta(S) \leq R$, then $S \leq \tilde{\eta}(R)$. Conversely, if $R < \eta(S)$, then $\tilde{\eta}(R) \leq S$.

Definition 2.3. A map $f : X \to Y$ is B-*dense* if $N_B^Y(f(X)) = Y$. We say that f is *coarsely surjective* if it is B-dense for some B. Two maps $f, g : X \to Y$ are r-*close* if $\sup_{x \in X} d_Y(f(x), g(x)) \leq r$. We say that f and g are *close* if they are r-close for some r.

Using Remark 2.2, we verify that a coarse embedding $f : X \to Y$ is coarsely surjective if and only if there exists a coarse embedding $g : Y \to X$ such that fg and gf are close to id_Y and id_X respectively. We then say that g is a *coarse inverse* to f and that f is a *coarse isometry*. We say $f : X \to Y$ is an (η, ϕ, B)-*coarse isometry* if it is an (η, ϕ)-coarse embedding and is B-dense.

If the distortion functions of a coarse embedding f are affine, then we say f is a *quasi-isometric embedding*. For constants $K \geq 1$ and $A \geq 0$, we say that $f : X \to Y$ is a (K, A)-*quasi-isometric embedding* if for every $x, x' \in X$,

$$\frac{1}{K} d_X(x, x') - A \leq d_Y(f(x), f(x')) \leq K d_X(x, x') + A.$$

A coarsely surjective quasi-isometric embedding is known as a *quasi-isometry*. If a (K, A)-quasi-isometric embedding is A-dense, we say that it is a (K, A)-quasi-isometry. Among the class of metric spaces, being quasi-isometric is an equivalence relation. We let $[[X]]$ denote the equivalence class consisting of all metric spaces quasi-isometric to X. We will see examples of coarse isometries and quasi-isometries in the context of group theory in Section 2.1.2.

For $V \subseteq X$, we say that $f, g : V \to Y$ are *A-close* if $\sup_{v \in V}\{d_Y(g(v), f(v))\} \leq A$. We say that f and g are *close* if they are A-close for some $A \geq 0$. We remark that if $V, V' \subseteq X$ are coarsely equivalent and $f : X \to Y$ is a quasi-isometry, then $f(V)$ and $f(V')$ are coarsely equivalent. Thus a quasi-isometry sends coarse equivalence classes to coarse equivalence classes.

We say that X and Y are *bi-Lipschitz equivalent* if there exists a bijective $(K, 0)$-quasi-isometric embedding from X to Y for some $K \geq 1$. Any quasi-isometry $f : X \to Y$ can be factored through a bi-Lipschitz equivalence by passing to and from suitable nets. More precisely, for any quasi-isometry $f : X \to Y$, $f|_N$ is close to a bi-Lipschitz equivalence $g : N \to M$ for suitably chosen nets $N \subseteq X$ and $M \subseteq Y$.

A metric space is said to be *coarse geodesic* (resp. *quasi-geodesic*) if it is coarsely isometric (resp. quasi-isometric) to a geodesic metric space. If $f : X \to Y$ is a coarse embedding and X is a quasi-geodesic metric space, then the upper distortion function of f can always be taken to be affine. Consequently, a coarse isometry between quasi-geodesic metric spaces is a quasi-isometry.

We now show that if a space is coarse geodesic, one can approximate it by a simplicial complex known as the *Rips complex*.

Definition 2.4. Let (X, d) be a metric space. For each $r \geq 0$, we define the *Rips complex* $P_r(X)$ to be the simplicial complex with vertex set X, where $\{x_0, \ldots, x_n\} \subseteq X$ spans an n-simplex if for all i, j

$$d(x_i, x_j) \leq r.$$

We define the *Rips graph* $P_r^1(X)$ to be the 1-skeleton of $P_r(X)$.

If $P_r^1(X)$ is connected, it can be endowed with the path metric in which edges have length 1. The following proposition shows that if X is coarse geodesic and r is sufficiently large, X and $P_r^1(X)$ are coarsely isometric. A *t-chain* of length n from x to y consists of a sequence $x = x_0, x_1, \ldots, x_n = y$ such that $d(x_i, x_{i+1}) \leq t$.

Proposition 2.5. *Let (X, d) be a metric space. The following are equivalent:*

(1) X is coarse geodesic;

(2) there exists a $t > 0$ and a proper non-decreasing function η such that for all $x, y \in X$, there is a t-chain from x to y of length at most $\eta(d(x, y))$;

(3) there exists a $t > 0$ such that for all $r \geq t$, the Rips graph $P_r^1(X)$ is connected and the inclusion $X \to P_r^1(X)$ is a coarse isometry.

Proof. (1) $\implies$ (2): There exists a geodesic metric space X' and an (η, ϕ, B)-coarse isometry $f : X' \to X$. For all $x, y \in X$, there are $x', y' \in X$ such that

$$d_X(f(x'), x), d_X(f(y'), y) \leq B.$$

Letting $R := d_{X'}(x', y')$, there is a geodesic $p : [0, R] \to X'$ from x' to y'. We choose $n \in \mathbb{N}$ such that $n-1 < R \leq n$. Let $x_{-1} = x$, $x_i = f(p(i))$ for $0 \leq i < n$, $x_n = f(y')$ and $x_{n+1} = y$. Letting $t := \max(B, \phi(1))$, we see $x_{-1}, \ldots, x_{n+1}$ is a t-chain of length $n + 2$ from x to y. By Remark 2.2, we see that $R + 2 \leq d_{X'}(x', y') + 3 \leq \tilde{\eta}(d(x, y) + 2B) + 3$.

(2) $\implies$ (3): Suppose there exist η and t such that (2) holds, and let $r \geq t$. Any two points of X can be joined by a t-chain; hence $P_r^1(X)$ is connected. Let d_r be the induced path metric on $P_r^1(X)$. Any $x, y \in X$ can be joined by a t-chain of length n with $n \leq \eta(d(x, y))$. Such a t-chain corresponds to an edge path of length n in $P_r^1(X)$, so $d_r(x, y) \leq n$. As $(P_r^1(X), d_r)$ is a geodesic metric space, x and y can be joined by an r-chain of length $d_r(x, y)$. By the triangle inequality, we see that $d(x, y) \leq d_r(x, y)r$. This implies that

$$\frac{d(x, y)}{r} \leq d_r(x, y) \leq \eta(d(x, y)),$$

verifying that the inclusion $X \to P_r^1(X)$ is indeed a coarse isometry.

(3) $\implies$ (1): This is clear, since $P_r^1(X)$ is a geodesic metric space. □

2.1.2 Finitely generated groups as geometric objects

Definition 2.6. Let (X, d) be a metric space. For $x, y \in X$, a *1-geodesic* between x and y is a 1-chain $x = x_0, x_1, \ldots, x_n = y$ such that $d(x_i, x_j) = |i - j|$ for each i, j. We say that X is *1-geodesic* if every pair of points can be joined by a 1-geodesic.

Let G be a group with finite generating set S. We equip G with the word metric d_S with respect to S. This metric is unique up to quasi-isometry — if S' is another finite generating set then the identity map $id_G : (G, d_S) \to (G, d_{S'})$ is a quasi-isometry. Unless otherwise stated, we always assume that finitely generated groups are equipped

with the word metric with respect to some finite generating set. All finitely generated groups equipped with the word metric are 1-geodesic.

If $H \leq G$ is a finitely generated subgroup equipped with the word metric, then the inclusion $\iota : H \to G$ is a coarse embedding, but not necessarily a quasi-isometric embedding (see Corollary 1.19 and Remark 1.20 of [Roe03]). Thus the *intrinsic* geometry of H (H equipped with the word metric) and the *extrinsic* geometry of H (H considered as a subspace of G, where G is itself equipped with the word metric) are the same up to coarse isometry.

In the case where H is a *distorted* subgroup, i.e. ι is not a quasi-isometric embedding, then ι is a coarse isometry onto its image, but not a quasi-isometry onto its image. This explains why coarse isometries are necessary in this thesis, and the notion of quasi-isometry is not sufficient.

Definition 2.7. Let G be a group and $H \leq G$ be a subgroup. We say $A \subseteq G$ is *H-finite* if there is a finite $T \subseteq G$ such that $A \subseteq HT$.

Remark 2.8. If $H \leq K \leq G$, then K is H-finite if and only if H has finite index in K.

Lemma 2.9. *Let G be a finitely generated group and H be a subgroup. Then $K \subseteq G$ is H-finite if and only if it is contained in $N_r(H)$ for some $r \geq 0$. In particular, if $H \leq K \leq G$, then H has finite index in K if and only if $K \subseteq N_r(H)$ for some $r \geq 0$.*

Proof. If $X \subseteq N_r(H)$, then every $x \in X$ can be written as ht, where $h \in H$ and $t \in G$ is a word of length at most r. Hence $X \subseteq HT$, where T contains all the finitely many words in G of length at most r. Conversely, if $X \subseteq HT$ for some finite $T \subseteq G$, we let $r := \max(\{\text{length}(t) \mid t \in T\})$. Then for all $x \in X$, we can write $x = ht$ for some $h \in H$ and $t \in T$, showing that $x \in N_r(H)$. □

Lemma 2.10 (See Lemma 2.2 of [MSW11])**.** *Let G be a finitely generated group and let $H, K \leq G$ be subgroups. Then for every $r, s \geq 0$, $K \cap H$ has finite Hausdorff distance from $N_r(K) \cap N_s(H)$.*

Proof. For every $g \in N_r(K) \cap N_s(H)$, we choose $k_g \in K$ and $h_g \in H$ such that $g \in N_r(k_g) \cap N_s(h_g)$; therefore $d(e, k_g^{-1} h_g) \leq r + s$. As the set $\Lambda := \{k_g^{-1} h_g \mid g \in N_r(K) \cap N_s(H)\}$ is finite, we choose $g_1, \ldots, g_n \in N_r(K) \cap N_s(H)$ such that $\Lambda := \{k_{g_1}^{-1} h_{g_1}, \ldots, k_{g_n}^{-1} h_{g_n}\}$.

We pick $R \geq 0$ large enough so that $\bigcup_{i=1}^{n} (N_r(k_{g_i}) \cap N_s(h_{g_i})) \subseteq N_R(e)$. We now claim that $K \cap H \subseteq N_r(K) \cap N_s(H) = N_R(K \cap H)$. Indeed, if $g \in N_r(K) \cap N_s(H)$,

then $k_g^{-1}h_g = k_{g_i}^{-1}h_{g_i}$ for some $1 \leq i \leq n$. Therefore, as $k_g k_{g_i}^{-1} = h_g h_{g_i}^{-1} \in K \cap H$, we see that

$$g \in N_r(k_g) \cap N_s(h_g) = k_g k_{g_i}^{-1}(N_r(k_{g_i}) \cap N_s(h_{g_i})) \subseteq N_R(k_g k_{g_i}^{-1}) \subseteq N_R(K \cap H).$$

□

Definition 2.11. Two subgroups $K, H \leq G$ are said to be *commensurable* if $K \cap H$ has finite index in both K and H.

Proposition 2.12. *Let G be a finitely generated group and $H, K \leq G$ be subgroups. Then:*

(1) $H \subseteq N_r(K)$ for some $r \geq 0$ if and only if H is commensurable to a subgroup of K;

(2) $d_{\text{Haus}}(K, H)$ is finite if and only if H and K are commensurable.

Proof. We observe that H is commensurable to a subgroup of K if and only if $H \cap K$ has finite index in H. Suppose there exists an $r \geq 0$ such that $H \subseteq N_r(K)$. By Lemma 2.10, $H = H \cap N_r(K)$ has finite Hausdorff distance from $H \cap K$; hence by Lemma 2.9, $H \cap K$ has finite index in H. Conversely, Lemma 2.9 says that if $H \cap K$ has finite index in H, then there exists an $r \geq 0$ such that $H \subseteq N_r(H \cap K) \subseteq N_r(K)$. This proves (1); (2) follows from (1). □

2.1.3 Coarse acyclicity

Let R be a commutative ring with unity. The following definition does not explicitly appear in the literature, but generalises related definitions such as Gromov's *uniform k-connectedness on the large scale* [Gro93, Section 1.D].

Definition 2.13. Let $\lambda : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ and $\mu : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be functions such that $\lambda(i) \geq i$ and $\mu(i, r) \geq r$ for all $i, r \in \mathbb{R}_{\geq 0}$. We say that a metric space X is (λ, μ)-*coarsely uniformly n-acyclic over R* if for every $k \leq n$, $x \in X$ and $i, r \in \mathbb{R}_{\geq 0}$, the map

$$\tilde{H}_k(P_i(N_r(x)); R) \to \tilde{H}_k(P_{\lambda(i)}(N_{\mu(i,r)}(x)); R),$$

induced by inclusion, is zero. We say that X is *coarsely uniformly n-acyclic over R* if it is (λ, μ)-coarsely uniformly n-acyclic over R for some suitable λ and μ. If X is coarsely uniformly n-acyclic over R for every n, then we say it is *coarsely uniformly acyclic over R*.

Example 2.14. If G is a hyperbolic group, then $P_i(G)$ is contractible for i sufficiently large. Thus for all $n \in \mathbb{N}$ there exists a μ such that for all $r \geq 0$, $\tilde{H}_k(P_i(N_r(e)); R) \to \tilde{H}_k(P_i(N_{\mu(i,r)}(e)); R)$ is zero for all $k \leq n$. As G acts cocompactly on $P_i(G)$, we thus see that G is coarsely uniformly n-acyclic over R. Hence G is coarsely uniformly acyclic over R.

Proposition 2.15. *If Y is (λ, μ)-coarsely uniformly n-acyclic over R and $f : X \to Y$ is an (η, ϕ, B)-coarse isometry, then X is (λ', μ')-coarsely uniformly n-acyclic over R, where λ' and μ' depend only on λ, μ, η, ϕ and B.*

Proof. We choose g, a coarse inverse to f, such that gf is A-close to id_X and both f and g have upper distortion function ψ. One can choose ψ and A depending only on η, ϕ and B. For $i, r \in \mathbb{R}_{\geq 0}$ and $x \in X$, we let $i_1 := \psi(i)$ and $r_1 := \psi(r)$. Then f induces a map $f_\# : C_\bullet(P_i(N_r^X(x))) \to C_\bullet(P_{i_1}(N_{r_1}^Y(f(x))))$ given by $[x_0, \ldots, x_n] \mapsto [f(x_0), \ldots, f(x_n)]$ on each oriented n-simplex of $P_i(N_r^X(x))$.

We define $i_2 := \lambda(i_1)$, $r_2 := \mu(i_1, r_1)$ and let

$$\tau_\# : C_\bullet(P_{i_1}(N_{r_1}^Y(f(x)))) \to C_\bullet(P_{i_2}(N_{r_2}^Y(f(x))))$$

be the inclusion. Then for $k \leq n$, the map $\tilde{H}_k(P_{i_1}(N_{r_1}^Y(f(x)))) \xrightarrow{\tau_*} \tilde{H}_k(P_{i_2}(N_{r_2}^Y(f(x))))$, induced by $\tau_\#$, is zero. Letting $i_3 := \psi(i_2) + A$ and $r_3 := \psi(r_2) + A$, we see that g induces the chain map $C_\bullet(P_{i_2}(N_{r_2}^Y(f(x)))) \xrightarrow{g_\#} C_\bullet(P_{i_3}(N_{r_3}^X(x)))$ given by $[y_0, \ldots, y_n] \mapsto [g(y_0), \ldots, g(y_n)]$. We observe that

$$g_\# \tau_\# f_\#([x_0, \ldots x_n]) = [gf(x_0), \ldots gf(x_n)]$$

for every oriented n-simplex of $P_i(N_r^X(x))$.

We define the chain homotopy $h_\# : C_\bullet(P_i(N_r^X(x))) \to C_{\bullet+1}(P_{i_3}(N_{r_3}^X(x)))$ by

$$[x_0, \ldots, x_n] \to \sum_{i=0}^{n} (-1)^n [x_0, \ldots, x_i, gf(x_i), \ldots gf(x_n)]$$

on each oriented n-simplex. We note that $\partial h_\# + h_\# \partial = g_\# \tau_\# f_\# - \iota_\#$, where $\iota_\#$ is the inclusion $C_\bullet(P_i(N_r^X(x))) \to C_\bullet(P_{i_3}(N_{r_3}^X(x)))$. Thus $\iota_* = g_* \tau_* f_* = 0$ on the level of reduced homology for $k \leq n$. □

Definition 2.16. A metric space (X, d) has *bounded geometry* if for any $r > 0$, there is a $K_r \in \mathbb{N}$ such that $|\{y \in X \mid d(x, y) \leq r\}| \leq K_r$ for all $x \in X$.

Example 2.17. A non-trivial finitely generated group G, equipped with the word metric with respect to some finite generating set S, is a bounded geometry metric space. However, the Cayley graph of G with the respect to S is *not* a bounded geometry metric space.

In many situations, it is preferable to work with bounded geometry metric spaces, which can often be done by passing to a net. This is because the Rips complex of a bounded geometry metric is a (uniformly) locally finite simplicial complex, thus allowing us to define cohomology with compact supports.

Proposition 2.18. *A bounded geometry metric space is coarsely uniformly* 0*-acyclic over* R *if and only if it is coarse geodesic.*

Proof. Let X be a bounded geometry metric space which is (λ, μ)-coarsely uniformly 0-acyclic over R. Without loss of generality, we can choose μ such that $r \mapsto \mu(0, r)$ is non-decreasing. For all $x \in X$ and $r \geq 0$, the map $\tilde{H}_0(P_0(N_r(x)); R) \to \tilde{H}_0(P_{\lambda(0)}(N_{\mu(0,r)}(x)); R)$, induced by inclusion, is zero; thus $N_r(x)$ is contained in a single connected component of $P_{\lambda(0)}(N_{\mu(0,r)}(x))$. Therefore, for all $x, y \in X$, there exists a $\lambda(0)$-chain in $N_{\mu(0,d(x,y))}(x)$ from x to y of length $n_{x,y}$. We may assume $n_{x,y}$ is minimal, so that no element of the $\lambda(0)$-chain is repeated, hence $n_{x,y}$ is bounded by the size of $N_{\mu(0,d(x,y))}(x)$. Since X has bounded geometry, the function $\eta(r) := \sup_{x \in X} \left| N_{\mu(0,r)}(x) \right|$ is proper and non-decreasing. Thus $n_{x,y} \leq \eta(d(x, y))$, so Proposition 2.5 tells us that X is coarse geodesic.

For the converse, we observe that whenever X is a geodesic metric space, every $N_r(x)$ is connected; therefore, X is coarsely uniformly 0-acyclic. By Proposition 2.15, we see that every coarse geodesic space is also coarsely uniformly 0-acyclic. □

2.2 Metric complexes and coarse cohomology

In this section, we define a notion of cohomology that is invariant under coarse isometry. The cohomology we use here is defined only for bounded geometry metric spaces that are coarsely uniformly $(n-1)$-acyclic, and only then in dimensions at most n. Our coarse cohomology is a special case of a more general notion due to Roe — see Section 2.2.4 and [Roe03] for more details. However, our approach is more amenable to quantitative methods than the approach taken by Roe.

We make heavy use of technology from [KK05] in the next two sections. In particular, we use metric complexes as defined in the appendix of [KK05]. Before diving straight into the theory of metric complexes, we first give a brief explanation of how metric complexes naturally arise when working with homological finiteness properties of groups.

Definition 2.19. A group is said to be of *type* F_n if it acts freely, cocompactly and cellularly on an $(n-1)$-connected CW complex. For a ring R, a group G is said to be of *type* FP_n^R if it admits a partial projective resolution

$$P_n \to P_{n-1} \to \cdots \to P_0 \to R \to 0$$

of the trivial RG-module R, such that each P_i is finitely generated as an RG-module.

It is easily seen that a group of type F_n is of type FP_n^R for any R. A group of type $FP_n^{\mathbb{Z}}$ is of type FP_n^R for any commutative ring R. It is shown in [BB97] that for $n \geq 2$, there exist groups of type FP_n^R that are not of type F_n.

In [DK18], Drutu and Kapovich define a condition known as *coarse n-connectedness*. Under the presence of a cocompact group action, this is analogous to our definition of coarse uniform n-acyclicity, using homotopy groups instead of reduced homology groups. The following theorem characterises groups of type F_n and FP_n^R in terms of their coarse geometry.

Theorem 2.20 ([KK05],[DK18])**.** *If G is a finitely generated group, then:*

(1) G is of type F_n if and only if it is coarsely $(n-1)$-connected;

(2) G is of type FP_n^R if and only if it is coarsely uniformly $(n-1)$-acyclic over R.

We give an outline of how to construct an $(n-1)$-connected CW complex needed to prove (1), see [DK18] for details. We proceed by induction. Suppose G is coarsely $(n-1)$-connected and one has already constructed an $(n-1)$-dimensional, $(n-2)$-connected CW complex X, admitting a free, cellular and cocompact G-action. The coarse $(n-1)$-connected condition ensures that one need only attach finitely many orbits of n-cells to X to obtain an $(n-1)$-connected CW complex X'. In doing this, we use the Hurewicz theorem — as X is $(n-2)$-connected, $H_{n-1}(X) \cong \pi_{n-1}(X)$ can be generated by finitely many orbits of spherical $(n-1)$-cycles. We equivariantly attach n-cells to these spherical cycles to kill off $\pi_{n-1}(X)$ and obtain a suitable $(n-1)$-connected complex admitting a free cocompact G-action.

To prove (2), we would like to replicate the above argument, replacing homotopy groups with reduced homology groups and thus building a CW complex that is $(n-1)$-acyclic over R and admits a free, cellular, cocompact G-action — this would certainly show that G is of type FP_n^R. However, if one tries to replicate the above inductive argument, one runs into problems.

Indeed, suppose a group is coarsely uniformly $(n-1)$-acyclic and that one has constructed an appropriate $(n-1)$-dimensional, $(n-2)$-acyclic CW complex X. Since

X is not necessarily $(n-2)$-connected, we cannot assume that $H_{n-1}(X)$ is generated by spherical cycles. Thus we cannot necessarily kill $H_{n-1}(X)$ by attaching n-cells. However, there are only finitely many orbits of (possibly non-spherical) cycles that generate $H_{n-1}(X)$. Thus we can equivariantly attach finitely many orbits of what we temporarily call *pseudo-cells* and obtain some "complex" X' that is $(n-1)$-acyclic and admits a free cocompact G-action. This complex is an example of a *metric complex* introduced by Kapovich and Kleiner in [KK05].

The complex X' isn't necessarily a CW complex since the pseudo-cells defined above may not have spherical boundary. Consequently, it may not have a topological realisation. However, it does have a well-defined $(n-1)$-acyclic chain complex

$$F \xrightarrow{\partial} C_{n-1}(X) \to \dots C_0(X) \to 0$$

which extends the cellular chain complex of X. Here, F is a free R-module with R-basis the pseudo-cells described above. The boundary map $F \xrightarrow{\partial} C_{n-1}(X)$ takes each pseudo-cell to the $(n-1)$-cycle it is attached to. This shows us that G is of type FP_n^R.

We now proceed to formally define metric complexes. For those uncomfortable with metric complexes, we remark that much of this paper still works if one replaces metric complexes with CW complexes. More specifically, one needs to use bounded geometry CW complexes as defined in [MSW03]. Doing this weakens the hypothesis of our results from groups of type $FP_{n+1}^{\mathbb{Z}_2}$ to groups of type F_{n+1}, but the argument is otherwise virtually unchanged.

2.2.1 Metric complexes

Let X be a bounded geometry metric space. We fix a commutative ring R with unity. A *free R-module over X* consists of a triple (M, Σ, p), where M is a free R-module, Σ is a basis of M and p is a map $p : \Sigma \to X$. We say that X is the *control space*, Σ is the *standard basis* and p is the *projection map*.

For convenience, we will often denote a free module over X simply as M, where the choice of Σ and p are implicit. We say that a free module over X has *finite type* if $|p^{-1}(x)|$ is uniformly bounded. For each $\sigma = \sum_{b\in\Sigma} r_b b \in M$, we define the *support* of σ to be $\mathrm{supp}(\sigma) := \{p(b) \mid r_b \neq 0\} \subseteq X$. We define $\mathrm{diam}(\sigma)$ to be the diameter of $\mathrm{supp}(\sigma)$.

Say X' is also a bounded geometry metric space and that (M, Σ, p) and (M', Σ', p') are free R-modules over X and X' respectively. If there exists an $r \geq 0$, an R-module

homomorphism $\hat{f} : M \to M'$ and a map $f : X \to X'$ such that

$$\operatorname{supp}(\hat{f}(\sigma)) \subseteq N_r(f(p(\sigma)))$$

for every $\sigma \in \Sigma$, then we say that $\hat{f}$ has *displacement at most r over f*. If there exists such an r, then we say that $\hat{f}$ has *finite displacement over f*. When $X = X'$, we say that $\hat{f}$ has *finite displacement* when it has finite displacement over id_X.

Lemma 2.21. *Let (M, Σ, p), (M', Σ', p') and (M'', Σ'', p'') be free modules over X, X' and X'' respectively. Say $f : X \to X'$ is an arbitrary map and $g : X' \to X''$ is an (η, ϕ)-coarse embedding. If $\hat{f} : M \to M'$ has displacement at most r over f and $\hat{g} : M' \to M''$ has displacement at most s over g, then $\hat{g}\hat{f}$ has displacement at most $s + \phi(r)$ over gf.*

Proof. For every $\sigma \in \Sigma$, we can write $\hat{f}(\sigma) = \sum_{b \in \Sigma'} r_b b$. If $r_b \neq 0$, then $p'(b) \in \operatorname{supp}(\hat{f}(\sigma)) \subseteq N_r(f(p(\sigma)))$. Therefore $d_{X''}(g(p'(b)), (g \circ f)(p(\sigma))) \leq \phi(r)$, so that

$$\operatorname{supp}(\hat{g}(b)) \subseteq N_s(g(p'(b))) \subseteq N_{s+\phi(r)}((g \circ f)(p(\sigma))).$$

Thus $\operatorname{supp}(\hat{g}\hat{f}(\sigma)) \subseteq N_{s+\phi(r)}((g \circ f)(p(\sigma)))$. □

Definition 2.22. An *R-metric complex* consists of the tuple $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$, where X is a bounded geometry metric space and $C_\bullet$ is a chain complex such that:

(1) each (C_i, Σ_i, p_i) is a free R-module over X of finite type and each boundary map ∂_i has finite displacement;

(2) the composition $\varepsilon \circ \partial_1 : C_1 \to R$ is zero, where $\varepsilon : C_0 \to R$ is the standard augmentation given by $\sigma \mapsto 1_R$ for each $\sigma \in \Sigma_0$;

(3) the projection map $p_0 : \Sigma_0 \to X$ is onto.

When unambiguous, we will denote an R-metric complex $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ by $(X, C_\bullet)$, or even just by $C_\bullet$. We will also often refer to R-metric complexes simply as metric complexes, where the choice of ring R is implicit.

Example 2.23. If X is a bounded geometry metric space, then every $(X, C_\bullet(P_i(X); R))$ is an R-metric complex, where $C_\bullet(P_i(X); R)$ is the simplicial chain complex of the Rips complex $P_i(X)$. The standard basis Σ_k of $C_k(P_i(X); R)$ is the set of k-simplicies of $P_i(X)$ and the projection map can be given by any map $p_k : \Sigma_k \to X$ such that $p_k(\sigma) \in \sigma$ for each $\sigma \in \Sigma_k$.

Example 2.24. A CW (or simplicial) complex with bounded geometry, as defined in [MSW03], [KK05] and [DK18], is a special case of a metric complex.

Definition 2.25. We define the *dimension* of $(X, C_\bullet)$ to be $\sup\{n \in \mathbb{N} \mid C_n \neq 0\}$, which may be finite or infinite. Let $[C_\bullet]_n$ denote the n-truncation of $C_\bullet$, i.e. the chain complex

$$\cdots \to 0 \to C_n \to C_{n-1} \to \cdots \to C_0 \to 0.$$

We call $[C_\bullet]_n$ the *n-skeleton* of $C_\bullet$.

We say that a metric complex $(X, C_\bullet)$ has *n-displacement at most r* if for each $i \leq n$, the boundary map ∂_i has displacement at most r. Given metric complexes $(X, C_\bullet)$, $(Y, D_\bullet)$ and a map $f : X \to Y$, we say that a chain map $f_\# : C_\bullet \to D_\bullet$ (resp. chain homotopy $h_\# : C_\bullet \to D_{\bullet+1}$) has *$n$-displacement at most r over f* if for each $i \leq n$, f_i (resp. h_i) has displacement at most r over f.

Given a topological space (or CW complex) X and a subspace (or subcomplex) K, one can define the subchain complex $C_\bullet(K) \subseteq C_\bullet(X)$. Things aren't so simple with metric complexes, since the displacement of the boundary maps means that it is possible that $p_n(\sigma) \in K$, but $\operatorname{supp}(\partial\sigma)$ is not contained in K. This motivates the following definition:

Definition 2.26. Let $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ be a metric complex. We say a metric complex $(K, C'_\bullet, \Sigma'_\bullet, p'_\bullet)$ is a *subcomplex* of $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ if $K \subseteq X$, $C'_\bullet$ is subchain complex of $C_\bullet$ and for each i, $\Sigma'_i \subseteq \Sigma_i$ and $p'_i = p_i|_{\Sigma'_i}$. For any $K \subseteq X$, we define the subcomplex generated by K to be the largest metric complex $(K, C_\bullet[K], \Sigma'_\bullet, p'_\bullet)$ that is a subcomplex of $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$. We let $C_\bullet[K]_n$ denote the n-skeleton of $C_\bullet[K]$.

Remark 2.27. The subcomplex $(K, C_\bullet[K], \Sigma'_\bullet, p'_\bullet)$ can be described explicitly as follows. We let $\Sigma'_0 = p^{-1}(K)$, with $C_0[K]$ the free module generated by Σ'_0. We inductively define

$$\Sigma'_{k+1} := \{\sigma \in \Sigma_{k+1} \mid \partial\sigma \in C_k[K] \text{ and } p_{k+1}(\sigma) \in K\}$$

and let $C_{k+1}[K]$ be the free module generated by Σ'_{k+1}.

Lemma 2.28. *Let $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ be a metric complex and suppose $\operatorname{supp}(\rho) \subseteq K$ for some $\rho \in C_n$. If $C_\bullet$ has n-displacement at most r, then $\rho \in C_n[N_{nr}(K)]$.*

Proof. We prove this by induction. The case $n = 0$ is clear, since each $\sigma \in \Sigma_0$ lies in $C_\bullet[K]$ precisely when $p_0(\sigma) \in K$. Let $k < n$. We assume for all $L \subseteq X$ and $\rho \in \Sigma_k$, that if $p_k(\rho) \in L$, then $\rho \in C_\bullet[N_{kr}(L)]$. Suppose $p_{k+1}(\sigma) \in K$ for

some $\sigma \in \Sigma_{k+1}$. Since ∂_{k+1} has displacement at most r, $\mathrm{supp}(\partial_{k+1}\sigma) \subseteq N_r(K)$. Therefore $\partial_{k+1}\sigma$ is contained in $C_\bullet[N_{kr}(N_r(K))] \subseteq C_\bullet[N_{(k+1)r}(K)]$. By Remark 2.27, we see that $\sigma \in C_\bullet[N_{(k+1)r}(K)]$. Hence if $\mathrm{supp}(\rho) \subseteq K$ for some $\rho \in C_{k+1}$, then $\rho \in C_{k+1}[N_{(k+1)r}(K)]$. □

Corollary 2.29. *Suppose $(X, C_\bullet)$ and $(Y, D_\bullet)$ are metric complexes and $f_\# : C_\bullet \to D_\bullet$ has finite n-displacement over $f : X \to Y$. Then there exists an $r \geq 0$, depending only on the n-displacement of $f_\#$ and $D_\bullet$, such that $f_\#([C_\bullet]_n) \subseteq D_\bullet[N_r^Y(f(X))]$.*

Remark 2.30. Using the metric complex structure of the Rips complex $P_i(X)$ from Example 2.23, we see that if $\sigma \in C_k(P_i(X); R)$ and $\mathrm{supp}(\sigma) \subseteq K$, then $\sigma \in C_\bullet(P_i(N_i^X(K)); R)$.

Definition 2.31. We define the cochain complex $C^\bullet[K] := \mathrm{Hom}_R(C_\bullet[K], R)$ and let $\delta^k : C^k[K] \to C^{k+1}[K]$ denote the coboundary map dual to the boundary map ∂_{k+1}. For every $\alpha \in C^k[K]$, we define its *support* to be $\mathrm{supp}(\alpha) := \{p_k(\sigma) \mid \sigma \in \Sigma_k \text{ and } \alpha(\sigma) \neq 0\}$.

As X has bounded geometry and each ∂_k has finite displacement, we see that δ^k preserves cochains of finite support. We thus define $C_c^\bullet[K]$ to be the subcochain complex of $C^\bullet[K]$ consisting of cochains with finite support.

We let $H_k(C_\bullet[K])$, $H^k(C_\bullet[K])$ and $H_c^k(C_\bullet[K])$ denote the k^{th} homology or cohomology of $C_\bullet[K]$, $C^\bullet[K]$ and $C_c^\bullet[K]$ respectively. We can also take the reduced homology $\tilde{H}_k(C_\bullet[K])$, which is the k^{th} homology of the augmented chain complex $\cdots \to C_1[K] \to C_0[K] \xrightarrow{\varepsilon} R$.

Definition 2.32. Given a metric complex $(X, C_\bullet)$ and $K \subseteq X$, we define

$$\hat{H}_c^k(C_\bullet[K]) := \ker(H_c^k(C_\bullet[K]) \xrightarrow{\iota^*} H^k(C_\bullet[K])),$$

where the map ι^* is induced by the inclusion $C_c^\bullet[K] \xrightarrow{\iota} C^\bullet[K]$. We call $\hat{H}_c^k$ the *modified cohomology with compact supports*.

It will often be convenient to denote $H_k(C_\bullet[K])$, $H^k(C_\bullet[K])$, $H_c^k(C_\bullet[K])$ and $\hat{H}_c^k(C_\bullet[K])$ by $H_k[K]$, $H^k[K]$, $H_c^k[K]$ and $\hat{H}_c^k[K]$ respectively. This notation implicitly assumes the choice of some metric complex $(X, C_\bullet)$, but can be used when there is no ambiguity.

For $K \subseteq X$, we define $C_\bullet[X, K] = C_\bullet[X]/C_\bullet[K]$. Then there is a long exact sequence

$$\cdots \to H_1[K] \to H_1[X] \to H_1[X, K] \to H_0[K] \to H_0[X] \to H_0[X, K] \to 0$$

and similar such sequences for cohomology and cohomology with compact supports.

Definition 2.33. Let $(X, C_\bullet)$ and $(Y, D_\bullet)$ be metric complexes. We say a chain map $f_\# : C_\bullet \to D_\bullet$ (resp. chain homotopy $h_\# : C_\bullet \to D_{\bullet+1}$) is *proper* if there is a coarse embedding $f : X \to Y$ such that for every k, $f_\#$ (resp. $h_\#$) has finite k-displacement over f.

The following proposition shows that proper chain maps between metric complexes induce maps in cohomology with compact supports.

Proposition 2.34. *Let $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ and $(Y, D_\bullet)$ be metric complexes.*

(1) A proper chain map $f_\# : [C_\bullet]_n \to [D_\bullet]_n$ induces maps $\hat{f}^* : \hat{H}^k_c(D_\bullet) \to \hat{H}^k_c(C_\bullet)$ for $k \leq n$.

(2) If proper chain maps $f_\#, g_\# : [C_\bullet]_n \to [D_\bullet]_n$ are chain homotopic via a proper chain homotopy $h_\# : [C_\bullet]_{n-1} \to [D_{\bullet+1}]_{n-1}$, then $\hat{f}^* = \hat{g}^*$ for $k \leq n$.

Proof. (1): The chain map $f_\#$ induces a dual map $f^\#$ given by $\alpha \mapsto \alpha f_\#$. As $f_\#$ is proper, it has n-displacement at most r over some coarse embedding $f : X \to Y$. We claim that for $k \leq n$ and $\alpha \in D^k_c$, $f(\mathrm{supp}(f^\#\alpha)) \subseteq N^Y_r(\mathrm{supp}(\alpha))$. Indeed, suppose $x \in \mathrm{supp}(f^\#\alpha)$; then there exists a $\sigma \in p_k^{-1}(x)$ such that $\alpha(f_\#\sigma) = (f^\#\alpha)(\sigma) \neq 0$. Thus $\mathrm{supp}(\alpha)$ intersects $\mathrm{supp}(f_\#\sigma) \subseteq N^Y_r(f(x))$, so $f(x) \in N^Y_r(\mathrm{supp}(\alpha))$, proving the claim.

Since f is a coarse embedding, the above claim demonstrates that $f^\#$ preserves cochains of finite support. Suppose $k \leq n$ and $\alpha \in D^k_c$ is a cocycle such that $\alpha = \delta\beta$ for some $\beta \in D^{k-1}$; then $f^\#\alpha = \delta f^\#\beta$ where $f^\#\beta \in C^{k-1}$. It is straightforward to verify that $[\alpha] \mapsto [f^\#\alpha]$ is a well-defined map in modified cohomology with compact supports.

(2): A proper chain homotopy $h_\# : [C_\bullet]_{n-1} \to [D_{\bullet+1}]_{n-1}$ induces a dual map $h^\#$ such that $\delta h^\# + h^\#\delta = g^\# - f^\#$ in dimensions at most $n-1$. The above argument shows that $h^\#$ preserves cochains of finite support. Suppose $k \leq n$ and $\alpha \in D^k_c$ is a cocycle such that $\alpha = \delta\beta$ for some $\beta \in D^{k-1}$. Then $(g^\# - f^\#)\alpha = \delta(g^\# - f^\#)\beta = \delta(\delta h^\# + h^\#\delta)\beta = \delta h^\#\alpha$. Thus $\hat{f}^* = \hat{g}^*$ for $k \leq n$. □

Remark 2.35. Using ordinary (not modified) cohomology with compact supports, an analogue of Proposition 2.34 holds for $k < n$. This is our motivation for introducing modified cohomology with compact supports, as it gives us an invariant one dimension higher than we would otherwise have.

Definition 2.36. A *group action* of G on $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ consists of a pair $(\rho, \hat{\rho})$, where $\rho : G \curvearrowright X$ and $\hat{\rho} : G \curvearrowright C_\bullet$ are group actions by isometries and chain automorphisms respectively, such that for each i and $g \in G$, $\hat{\rho}(g)(\Sigma_i) = \Sigma_i$ and p_i is G-equivariant. This action is *free* (resp. *cocompact*) if the action $\rho : G \curvearrowright X$ is free (resp. cocompact).

Whenever a metric complex $(G, C_\bullet)$ admits a G-action, it will always be assumed that $\rho : G \curvearrowright G$ is the action by left multiplication.

Example 2.37. Our conditions for a metric complex to admit a G-action are reasonably restrictive. For instance, if a group G has torsion, then for sufficiently large r, G cannot act on $(G, P_r(G))$ despite the fact that G has a natural simplicial action on $P_r(G)$. This is because some $g \in G \backslash \{e\}$ fixes a simplex of $P_r(G)$, so the projections maps can never be G-equivariant. If G is torsion-free, then projection maps can be chosen so that G acts freely on $(G, P_r(G))$.

2.2.2 Uniformly acyclic complexes

The results in this subsection form the technical toolkit needed to define and work with coarse cohomology.

Definition 2.38 (cf. Definition 2.13)**.** Let $\mu : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a function such that $\mu(r) \geq r$ for each r. An R-metric complex $(X, C_\bullet)$ is said to be *μ-uniformly n-acyclic* if for every $x \in X$, $r \in \mathbb{R}_{\geq 0}$ and $k \leq n$, the map

$$\tilde{H}_k[B_r(x)] \to \tilde{H}_k[B_{\mu(r)}(x)],$$

induced by inclusion, is zero. We say $(X, C_\bullet)$ is *uniformly n-acyclic* if it is μ-uniformly n-acyclic for some μ. We say that $(X, C_\bullet)$ is *uniformly acyclic* if it is uniformly n-acyclic for all n.

As in Example 2.14, if G is a hyperbolic group, then $C_\bullet(P_i(G))$ is uniformly acyclic for i sufficiently large. Unfortunately, for a general coarsely uniformly $(n-1)$-acyclic metric space X, $C_\bullet(P_i(X))$ is not necessarily uniformly $(n-1)$-acyclic for large i.

Lemma 2.39. *Suppose X and Y are bounded geometry metric spaces, $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ is an R-metric complex with n-displacement at most d and Y is (λ, μ)-coarsely uniformly $(n-1)$-acyclic over R.*

(1) There exists an $i = i(\lambda)$ such that any (η, ϕ)-coarse embedding $f : X \to Y$ induces a chain map

$$f_\# : [C_\bullet]_n \to C_\bullet(P_i(Y); R)$$

of n-displacement at most $M = M(\lambda, \mu, \phi, d)$ over f.

(2) *For $i \geq 0$, suppose $f_\#, g_\# : [C_\bullet]_n \to C_\bullet(P_i(Y); R)$ are chain maps of n-displacement at most r over an (η, ϕ)-coarse embedding $f : X \to Y$. Then there exists a $j = j(i, \lambda) \geq i$ and a chain homotopy*

$$h_\# : [C_\bullet]_{n-1} \to C_{\bullet+1}(P_j(Y); R),$$

of $(n-1)$-displacement at most $N = N(i, \lambda, \mu, \phi, r, d)$ over f, between $\iota_\# f_\#$ and $\iota_\# g_\#$, where $\iota_\# : C_\bullet(P_i(Y); R) \to C_\bullet(P_j(Y); R)$ is the inclusion.

Proof. We define $y_\Delta := f(p_i(\Delta)) \in Y$ for every $\Delta \in \Sigma_i$.

(1): We proceed by induction. For the base case, we extend the map $f \circ p_0 : \Sigma_0 \to Y$ to an augmentation preserving map $f_0 : C_0 \to C_0(P_0(Y); R)$, which has displacement zero over f.

Let $k < n$. We assume that there exists some $i_k = i_k(\lambda)$ and a chain map $f_\# : [C_\bullet]_k \to C_\bullet(P_{i_k}(Y); R)$ of k-displacement at most $M_k = M_k(\lambda, \mu, \phi, d)$ over f. By Lemma 2.21, we see that $\mathrm{supp}(f_k \partial_{k+1} \Delta) \subseteq N^Y_{\phi(d)+M_k}(y_\Delta)$ for every $\Delta \in \Sigma_{k+1}$. Furthermore, by Remark 2.30, we see that $f_k \partial_{k+1} \Delta \in C_k(P_{i_k}(N^Y_{\phi(d)+M_k+i_k}(y_\Delta)); R)$. Letting $i_{k+1} := \lambda(i_k)$ and $M_{k+1} := \mu(i_k, \phi(d) + M_k + i_k)$, there exists an $\omega_\Delta \in C_{k+1}(P_{i_{k+1}}(N^Y_{M_{k+1}}(y_\Delta)); R)$ such that $\partial_{k+1} \omega_\Delta = f_k \partial_{k+1} \Delta$. We define $f_{k+1}(\Delta) = \omega_\Delta$ for each $\Delta \in \Sigma_{k+1}$, thus defining a chain map

$$f_\# : [C_\bullet]_{k+1} \to C_\bullet(P_{i_{k+1}}(Y); R)$$

of $(k+1)$-displacement at most M_{k+1} over f.

(2): We also use induction. For ease of notation, we omit the inclusion maps, simply identifying $C_\bullet(P_i(Y); R)$ with a subcomplex of $C_\bullet(P_j(Y); R)$. If $\Delta \in \Sigma_0$, then $\mathrm{supp}(f_0(\Delta) - g_0(\Delta)) \subseteq N^Y_r(y_\Delta)$; thus by Remark 2.30 we see that $f_0(\Delta) - g_0(\Delta) \in C_0(P_i(N^Y_{r+i}(y_\Delta)); R)$. Since $f_\#$ and $g_\#$ are both augmentation preserving, $f_0(\Delta) - g_0(\Delta)$ is a reduced 0-cycle. Thus there is an $\omega_\Delta \in C_1(P_{\lambda(i)}(N^Y_{\mu(i,r+i)}(y_\Delta)))$ such that $\partial \omega_\Delta = f_0(\Delta) - g_0(\Delta)$. Letting $h_0(\Delta) = \omega_\Delta$ for each $\Delta \in \Sigma_0$, we define a chain homotopy $h_\# : [C_\bullet]_0 \to C_1(P_{\lambda(i)}(Y); R)$ that has 0-displacement $\mu(i, r+i)$ over f such that $\partial h_\# + h_\# \partial = g_\# - f_\#$.

Let $k < n-1$. We assume there is a $j_k = j_k(i, \lambda) \geq i$ and a proper chain homotopy $h_\# : [C_\bullet]_k \to C_{\bullet+1}(P_{j_k}(Y); R)$ such that $\partial h_\# + h_\# \partial = g_\# - f_\#$ and $h_\#$ has k-displacement at most $N_k = N_k(i, \lambda, \mu, \phi, d, r)$ over f. Lemma 2.21 and Remark 2.30 tell us that for each $\Delta \in \Sigma_{k+1}$,

$$h_\#(\partial \Delta) - g_\#(\Delta) + f_\#(\Delta) \in C_{k+1}(P_{j_k}(N^Y_{R_k}(y_\Delta)); R),$$

where $R_k := \max(\phi(d) + N_k + j_k, r + j_k)$. We deduce that there exists an $\omega_\Delta \in C_{k+2}(P_{\lambda(j_k)}(N^Y_{\mu(j_k,R_k)}(y_\Delta)); R)$ such that $\partial_{k+2}\omega_\Delta = h_\#(\partial\Delta) - g_\#(\Delta) + f_\#(\Delta)$. We therefore define $h_{k+1}(\Delta) = \omega_\Delta$ for each $\Delta \in \Sigma_{k+1}$. □

Coupled with Proposition 2.42, the following lemma is essential in allowing us to define coarse cohomology. It is a metric complex version of Proposition 6.47 and Corollary 6.49 of [DK18]. It is similar to Lemma 2.39, except that the target space is now a uniformly acyclic metric complex rather than the Rips complex of a coarsely uniformly acyclic space.

Lemma 2.40. *Suppose $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ and $(Y, D_\bullet, \Sigma'_\bullet, p'_\bullet)$ are R-metric complexes, $(Y, D_\bullet)$ is μ-uniformly $(n-1)$-acyclic and $C_\bullet$ and $D_\bullet$ have n-displacement at most d_1 and d_2 respectively. Then:*

(1) every (η, ϕ)-coarse embedding $f : X \to Y$ induces a chain map $f_\# : [C_\bullet]_n \to D_\bullet$ of n-displacement at most $M = M(\mu, \phi, d_1, d_2)$ over f;

(2) for every (η, ϕ)-coarse embedding $f : X \to Y$ and every pair of chain maps $f_\#, g_\# : [C_\bullet]_n \to D_\bullet$ of n-displacement at most r over f, there exists a chain homotopy $h_\# : [C_\bullet]_{n-1} \to D_{\bullet+1}$ between $f_\#$ and $g_\#$ which has $(n-1)$-displacement at most $N = N(\mu, \phi, d_1, d_2, r)$ over f.

Proof. The proof is very similar to the proof of Lemma 2.39, so we will only outline where the proofs differ. For the base case of (1), we use the fact that p'_0 is surjective to find a $\Delta' \in \Sigma'_0$ such that $p'_0(\Delta') = f(p_0(\Delta))$ for each $\Delta \in \Sigma_0$. We thus define $f_0(\Delta) = \Delta'$ for each $\Delta \in \Sigma_0$ and then extend linearly.

For the inductive step, we use the displacement of $[D_\bullet]_n$. Defining y_Δ as in the proof of Lemma 2.39, Lemma 2.21 says that for each $\Delta \in \Sigma_{k+1}$, $\operatorname{supp}(f_k\partial_{k+1}\Delta) \subseteq N^Y_{\phi(d_1)+M_k}(y_\Delta)$. Thus by applying Lemma 2.28, we see that

$$f_k\partial_{k+1}\Delta \in D_k[N^Y_{\phi(d_1)+M_k+kd_2}(y_\Delta)].$$

We extend $f_\#$ by setting $f_{k+1}(\Delta) = \omega_\Delta$ for some $\omega_\Delta \in D_{k+1}[N^Y_{\mu(\phi(d_1)+M_k+kd_2)}(y_\Delta)]$ such that $\partial\omega_\Delta = f_k\partial_{k+1}\Delta$. Part (2) is proved analogously, also making use of Lemma 2.28. □

Making minor modifications to the above proofs, we deduce the following:

Lemma 2.41. *Let $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ and $(Y, D_\bullet)$ be R-metric complexes such that $(Y, D_\bullet)$ is μ-uniformly $(n-1)$-acyclic and $C_\bullet$ and $D_\bullet$ have n-displacement at most d_1 and d_2 respectively. Suppose also there is a subcomplex $(X, C'_\bullet, \Sigma'_\bullet, p'_\bullet)$ of $(X, C_\bullet)$, and a chain map $f'_\# : [C'_\bullet]_n \to D_\bullet$ of n-displacement at most M' over an (η, ϕ)-coarse embedding $f : X \to Y$. Then $f'_\#$ extends to a chain map $f_\# : [C_\bullet]_n \to D_\bullet$ of n-displacement at most $M = M(\mu, \phi, d_1, d_2, M')$ over f.*

Proof. For $i \leq n$ and $\sigma \in \Sigma'_i \subseteq \Sigma_i$, we define $f_\#(\sigma) := f'_\#(\sigma)$. For $\sigma \in \Sigma_i \backslash \Sigma'_i$, we proceed as in Lemma 2.40. □

The following proposition allows us to construct a uniformly acyclic metric complex for any coarsely uniformly acyclic metric space, formalizing the construction discussed in the introduction to this section.

Proposition 2.42 (See Lemma 5.10 and Proposition 11.4 of [KK05])**.** *Let X be a bounded geometry metric space.*

(1) If X is (λ, μ)-coarsely uniformly $(n-1)$-acyclic over R, then it is the control space of a μ'-uniformly $(n-1)$-acyclic R-metric complex $(X, C_\bullet)$ of n-displacement at most d, where d and μ' depend only on λ and μ.

(2) Suppose that for each n, X is (λ_n, μ_n)-coarsely uniformly $(n-1)$-acyclic over R. Then X is the control space of a uniformly acyclic R-metric complex $(X, C_\bullet)$ such that for each n, $(X, C_\bullet)$ is μ'_n-uniformly $(n-1)$-acyclic and has n-displacement at most d_n, where d_n and μ'_n depend only on λ_i and μ_i for $i \leq n$.

(3) Suppose a group G acts freely on X. Then (1) and (2) hold, and the resulting metric complex $(X, C_\bullet)$ can be chosen so that it admits a free G-action.

Proof. (1): We proceed by induction on n. For each $x \in X$ and $r \geq 0$, we see that the map $\tilde{H}_0(P_0(N_r(x)); R) \to \tilde{H}_0(P_{\lambda(0)}(N_{\mu(0,r)}(x)); R)$, induced by inclusion, is zero. This means $N_r(x)$ is contained in a single connected component of $P_{\lambda(0)}(N_{\mu(0,r)}(x))$, so $\tilde{H}_0[P_{\lambda(0)}(N_r(x))] \to \tilde{H}_0[P_{\lambda(0)}(N_{\mu(0,r)}(x))]$, also induced by inclusion, is zero. Let $\mu_1(r) := \mu(0, r)$. Therefore, $C_\bullet(P_{\lambda(0)}(X); R)$ is μ_1-uniformly 0-acyclic and has displacement at most $\lambda(0)$.

Let $0 < k < n$. For our inductive hypothesis, we assume there is a k-dimensional R-metric complex $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ that is μ_k-uniformly $(k-1)$-acyclic with k-displacement at most d_k, where μ_k and d_k depend only on λ and μ. By Lemma 2.39, there is an $i_k = i_k(\lambda)$ such that the identity $\mathrm{id}_X : X \to X$ induces the chain map $f_\# :$

$C_\bullet \to C_\bullet(P_{i_k}(X); R)$ of k-displacement at most $M_k = M_k(\lambda, \mu)$ over id_X. Let $\iota_\# : C_\bullet(P_{i_k}(X); R) \to C_\bullet(P_{\lambda(i_k)}(X); R)$ be the inclusion.

We now consider the algebraic mapping cylinder $D_\bullet$ of $\iota_\# f_\#$, defined by $D_j := C_j \oplus C_{j-1} \oplus C_j(P_{\lambda(i_k)}(X))$ with boundary maps $\tilde{\partial}(a, b, c) = (\partial a + b, -\partial b, \partial c - \iota_\# f_\# b)$. Each D_j is the direct sum of finite type free R-modules over X, so inherits the structure of a finite type free R-module over X. Since $C_\bullet$, $C_\bullet(P_{\lambda(i_k)}(X))$ and $\iota_\# f_\#$ have $(k+1)$-displacements at most d_k, $\lambda(i_k)$ and M_k respectively, we see that $(X, D_\bullet)$ is an R-metric complex of $(k+1)$-displacement at most $d_{k+1} := \max(d_k, \lambda(i_k), M_k)$.

We note that $(X, C_\bullet)$ can be considered a subcomplex of $(X, D_\bullet)$ via the inclusion map $\tau_\#$ given by $a \mapsto (a, 0, 0)$. Thus Lemma 2.41 allows us to extend the identity $\mathrm{id} : [C_\bullet]_k \to [C_\bullet]_k$ to a chain map $r_\# : [D_\bullet]_k \to C_\bullet$ of k-displacement at most $P_k = P_k(\mu_k, d_{k+1})$ over id_X such that $r_\# \tau_\# = \mathrm{id}_{[C_\bullet]_k}$.

We define the boundary map $\partial : D_{k+1} \to C_k$ to be the composition $r_\# \circ \tilde{\partial}$. We now consider the R-metric complex

$$D_{k+1} \xrightarrow{\partial} C_k \to \cdots \to C_0 \to 0, \tag{2.1}$$

which has $(k+1)$-displacement at most $d_{k+1} + P_k$ over the identity. For $x \in X$, let $\sigma \in C_k[N_r(x)]$ be a reduced cycle. As $f_\#$ has k-displacement at most M_k over id_X, $\mathrm{supp}(f_\# \sigma) \subseteq N_{r+M_k}(x)$; thus by Remark 2.30, $f_\# \sigma \in C_k(P_{i_k}(N_{r+M_k+i_k}(x)); R)$. Hence, there is an $\omega_\sigma \in C_{k+1}(P_{\lambda(i_k)}(N_{\mu(i_k, r+M_k+i_k)}(x)); R)$ such that $\partial \omega_\sigma = \iota_\# f_\# \sigma$. Letting $\lambda := (0, \sigma, \omega_\sigma) \in D_{k+1}$, we see $\tilde{\partial}\lambda = (\sigma, 0, 0) = \tau_\# \sigma$. Therefore $\partial \lambda = r_\# \tau_\# \sigma = \sigma$. Since $\lambda \in D_{k+1}[N_{\mu(i_k, r+M_k+i_k)}(x)]$, we see that (2.1) is μ'-uniformly k-acyclic, where $\mu'(r) := \mu(i_k, r + M_k + i_k)$.

(2): We apply the method used to prove (1) in each dimension.

(3): We inductively modify the chain complex produced in the proof of (1) so that it admits a G-action. The base case is trivial since G acts freely on the zero skeleton of every Rips complex. For the inductive hypothesis, we assume the metric complex $C_k \to \cdots \to C_0 \to 0$ admits a free G-action, is μ_k-uniformly $(k-1)$-acyclic and has k-displacement at most d_k. As in the proof of (1), we then construct a finite type free R-module $(D_{k+1}, \Sigma_{k+1}, p_{k+1})$ over X such that $D_{k+1} \xrightarrow{\partial} C_k \to \cdots \to C_0 \to 0$ is a μ'-uniformly k-acyclic R-metric complex with $(k+1)$-displacement at most $d_{k+1} + P_k$, where d_{k+1} and P_k are as defined above. Let $S := \mu'(d_{k+1} + P_k + k d_k)$.

We choose $T \subseteq X$ such that each G-orbit in X contains precisely one element of T. Let $\Sigma'_{k+1} = \{(g, t, \rho) \in G \times T \times \Sigma_{k+1} \mid \rho \in D_{k+1}[N_S(t)]\}$ and let D'_{k+1} be the free R-module with R-basis Σ'_{k+1}. Then $(D'_{k+1}, \Sigma'_{k+1}, p'_{k+1})$ is a finite type free R-module over X, where the projection map is defined by $p'_{k+1}(g, t, \rho) := gt$. We

define the boundary map $\partial' : D'_{k+1} \to C_k$ by $(g,t,\rho) \mapsto g\partial\rho$. Using the definition of D'_{k+1} and the displacement of $\partial : D_{k+1} \to D_k$, we see that ∂' has displacement at most $d_{k+1} + P_k + S$. As D'_{k+1} admits a G-action given by $g'(g,t,\rho) \mapsto (g'g,t,\rho)$, we see that

$$D'_{k+1} \xrightarrow{\partial'} C_k \to \cdots \to C_0 \to 0 \tag{2.2}$$

is a metric complex that admits a G-action and has $(k+1)$-displacement at most $d_{k+1} + P_k + S$.

We claim that for every $\Delta \in \Sigma_{k+1}$, there exists a $\hat{\Delta} \in D'_{k+1}$ such that $\partial'\hat{\Delta} = \partial\Delta$. Indeed, since ∂ has displacement at most $d_{k+1} + P_k$ over the identity, $\operatorname{supp}(\partial\Delta) \subseteq N_{d_{k+1}+P_k}(p_{k+1}(\Delta))$. Lemma 2.28 then tells us that $\partial\Delta \in C_k[N_{d_{k+1}+P_k+kd_k}(p_{k+1}(\Delta))]$. There exists a unique $g \in G$ and $t \in T$ such that $p_{k+1}(\Delta) = gt$; thus $g^{-1}\partial\Delta \in C_k[N_{d_{k+1}+P_k+kd_k}(t)]$. As $g^{-1}\partial\Delta$ is a reduced k-cycle, there exists a $\gamma_\Delta \in D_{k+1}[N_S(t)]$ such that $\partial\gamma_\Delta = g^{-1}\partial\Delta$. We can write $\gamma_\Delta = \sum_{i=1}^n r_i\gamma_i$, where $r_i \in R$ and $\gamma_i \in \Sigma_{k+1} \cap D_{k+1}[N_S(t)]$ for each i. Let $\hat{\Delta} := \sum_{i=1}^n r_i(g,t,\gamma_i) \in D'_{k+1}$. Then $\partial'\hat{\Delta} = \sum_{i=1}^n r_i g\partial\gamma_i = g\partial\sum_{i=1}^n r_i\gamma_i = g\partial\gamma_\Delta = \partial\Delta$. As $\operatorname{supp}(\hat{\Delta}) = \{p_{k+1}(\Delta)\}$ and $\partial\Delta \in C_k[N_{d_{k+1}+P_k+kd_k}(p_{k+1}(\Delta))] \subseteq C_k[N_S(p_{k+1}(\Delta))]$, it follows from Remark 2.27 that $\hat{\Delta} \in D'_{k+1}[N_S(p_{k+1}(\Delta))]$.

Let $\sigma \in C_k[N_r(x)]$ be a reduced k-cycle. Then there exists an $\omega \in D_{k+1}[N_{\mu'(r)}(x)]$ such that $\partial\omega = \sigma$. We can write $\omega = \sum_{i=1}^m r'_i\omega_i$, where $\omega_i \in \Sigma_{k+1}$ and $r'_i \in R$ for each i. Thus $\partial\sum_{i=1}^m r'_i\hat{\omega}_i = \partial'\sum_{i=1}^m r'_i\omega_i = \partial\omega = \sigma$ and $\sum_{i=1}^m r'_i\hat{\omega}_i \in D'_{k+1}[N_{\mu'(r)+S}(x)]$. Letting $\mu''(r) = \mu'(r) + S$, we thus see that (2.2) is μ''-uniformly k-acyclic. □

2.2.3 Coarse cohomology

Suppose $(X, C_\bullet)$ and $(X, C'_\bullet)$ are two uniformly $(n-1)$-acyclic R-metric complexes. Applying Lemma 2.40 to id_X, we define proper chain maps $f_\# : [C_\bullet]_n \to [C'_\bullet]_n$ and $g_\# : [C'_\bullet]_n \to [C_\bullet]_n$ of finite n-displacement over the identity. By Lemma 2.21, both $g_\# f_\#$ and $f_\# g_\#$ have finite n-displacement over the identity, thus Lemma 2.40 says they are chain homotopic to $\mathrm{id}_{[C_\bullet]_{n-1}}$ and $\mathrm{id}_{[C'_\bullet]_{n-1}}$ respectively. By Proposition 2.34, $f_\#$ and $g_\#$ induce maps $\hat{f}^* : \hat{H}^k_c(C'_\bullet) \to \hat{H}^k_c(C_\bullet)$ and $\hat{g}^* : \hat{H}^k_c(C_\bullet) \to \hat{H}^k_c(C'_\bullet)$ for $k \leq n$. Proposition 2.34 also tells us that each of $\hat{f}^*\hat{g}^*$ and $\hat{g}^*\hat{f}^*$ is the identity, thus $\hat{f}^*$ is an isomorphism.

Moreover, any two chain maps $f_\#, f'_\# : [C_\bullet]_n \to [C'_\bullet]_n$ of finite displacement over id_X are properly chain homotopic by Lemma 2.40. Thus by Proposition 2.34, $\hat{f}^*$ doesn't depend on the choice of $f_\#$. Consequently, for $k \leq n$ the isomorphism $\hat{f}^* : \hat{H}^k_c(C'_\bullet) \to \hat{H}^k_c(C_\bullet)$ is canonical, so $\hat{H}^k_c(C_\bullet)$ is effectively an invariant of X,

independent of the choice of $C_\bullet$. This discussion demonstrates that the following is well-defined:

Definition 2.43. Let X be a bounded geometry metric space that is coarsely uniformly $(n-1)$-acyclic over R. Let $(X, C_\bullet)$ be a uniformly $(n-1)$-acyclic R-metric complex, whose existence is guaranteed by Proposition 2.42. For $k \leq n$, we define the *coarse cohomology* of X to be $H^k_{\text{coarse}}(X; R) := \hat{H}^k_c(C_\bullet)$.

Remark 2.44. It follows easily from the definition of modified cohomology with compact supports that $\hat{H}^0_c(C_\bullet) = 0$ for every metric complex $(X, C_\bullet)$. Thus for every bounded geometry metric space, $H^0_{\text{coarse}}(X; R) = 0$.

The subsequent lemma demonstrates that it is only necessary to use modified cohomology with compact supports to define $H^n_{\text{coarse}}(X; R)$ when X is coarsely uniformly $(n-1)$-acyclic over R and not coarsely uniformly n-acyclic over R. Otherwise, we can just take the ordinary (not modified) cohomology with compact supports of a suitable metric complex.

Lemma 2.45. *Suppose X is an infinite metric space, $n > 0$ and $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ is a uniformly $(n-1)$-acyclic R-metric complex. Then for $k < n$, $H^k_c(C_\bullet) = \hat{H}^k_c(C_\bullet)$.*

Proof. Since $C_\bullet$ is $(n-1)$-acyclic, $H^k(C_\bullet) = 0$ for $0 < k < n$ and $H^0(C_\bullet) \cong R$. This can be seen using the universal coefficients theorem for cohomology, or can be calculated directly. Therefore, for $0 < k < n$, $H^k_c(C_\bullet) = \ker(H^k_c(C_\bullet) \xrightarrow{\iota^*} H^k(C_\bullet)) = \hat{H}^k_c(C_\bullet)$.

Suppose $\alpha \in C^0_c$ is a cocycle and $\Delta \in \Sigma_0$. For any $\Delta' \in \Sigma_0$, $\Delta - \Delta'$ is a reduced 0-cycle, hence there is a 1-chain $\rho \in C_1$ such that $\partial\rho = \Delta - \Delta'$. Therefore $\alpha(\Delta) - \alpha(\Delta') = \alpha(\partial\rho) = \delta\alpha(\rho) = 0$. Since X is infinite and α has finite support, we see that $\alpha(\Delta) = 0$. As Δ was arbitrary, $\alpha = 0$; hence $H^0_c(C_\bullet) = 0$. Therefore $\hat{H}^0_c(C_\bullet) = H^0_c(C_\bullet) = 0$. □

Proposition 2.46. *Let X and Y be bounded geometry metric spaces that are coarsely uniformly $(n-1)$-acyclic over R. Then a coarse embedding $f : X \to Y$ induces a homomorphism $f^* : H^k_{\text{coarse}}(Y; R) \to H^k_{\text{coarse}}(X; R)$ for all $k \leq n$. Moreover, for any bounded geometry metric spaces X, Y and Z that are coarsely uniformly $(n-1)$-acyclic over R, then:*

(1) $(\mathrm{id}_X)^* = \mathrm{id}_{H^k_{\text{coarse}}(X;R)}$;

(2) if $f, g : X \to Y$ are close coarse embeddings, then $f^* = g^*$;

(3) if $f : X \to Y$ and $g : Y \to Z$ are coarse embeddings, then $(gf)^* = f^*g^*$.

In particular, if $f : X \to Y$ is a coarse isometry, then f^* is an isomorphism.

Proof. Let $(X, C_\bullet)$, $(X, C'_\bullet)$, $(Y, D_\bullet)$ and $(Y, D'_\bullet)$ be uniformly $(n-1)$-acyclic R-metric complexes, which necessarily exist by Proposition 2.42.

$$\begin{array}{ccc} \hat{H}^k_c(D'_\bullet) & \xrightarrow{\hat{f}'^*} & \hat{H}^k_c(C'_\bullet) \\ \downarrow{\scriptstyle (\mathrm{id}_Y)^*} & & \downarrow{\scriptstyle (\mathrm{id}_X)^*}. \\ \hat{H}^k_c(D_\bullet) & \xrightarrow{\hat{f}^*} & \hat{H}^k_c(C_\bullet) \end{array}$$

We now use Proposition 2.34 and Lemma 2.40 to construct the commutative diagram shown above. In view of this commutative diagram and the discussion preceding Definition 2.43, we see f induces a map $f^* : H^k_{\mathrm{coarse}}(Y; R) \to H^k_{\mathrm{coarse}}(X; R)$, independent of the choice of $C_\bullet$ and $D_\bullet$. Parts (1)–(3) follow from Lemmas 2.21 and 2.40. □

Remark 2.47. Let X and Y be bounded geometry metric spaces that are (λ, μ)-coarsely uniformly $(n-1)$-acyclic over R, and suppose $f : X \to Y$ is an (η, ϕ)-coarse embedding. Using Proposition 2.42, we can construct μ'-uniformly $(n-1)$-acyclic R-metric complexes $(X, C_\bullet)$ and $(Y, D_\bullet)$ of n-displacement at most d, where μ' and d depend only on λ and μ. Therefore, by Lemma 2.40, we can construct a map $f_\# : [C_\bullet]_n \to D_\bullet$ of n-displacement at most $D = D(\lambda, \mu, \eta, \phi)$ over f. Thus there is some quantitative information associated to the induced map $f^* : H^k_{\mathrm{coarse}}(Y; R) \to H^k_{\mathrm{coarse}}(X; R)$. This idea is crucial in the results of Section 3.1.4 and their consequences.

2.2.4 Roe's coarse cohomology

There is a more general notion of coarse cohomology due to Roe that doesn't require coarse uniform acyclicity (see [Roe03]). When defined, H^*_{coarse} is naturally isomorphic to Roe's coarse cohomology. However, working with anti-Čech approximations à la Roe, we lose the quantitative information (see Remark 2.47) that we have when working with metric complexes. As Kapovich and Kleiner say in [KK05], "one inevitably loses quantitative information which is essential in many applications of coarse topology to quasi-isometries and geometric group theory".

Definition 2.48. Let (Y, d) be a bounded geometry metric space. An open cover $\mathcal{U}$ is called *good* if it is locally finite and the closure of each $U \in \mathcal{U}$ is compact. A sequence $\mathcal{U}_1, \mathcal{U}_2, \ldots$ of good covers of Y is said to be an *anti-Čech approximation* if there exists a sequence of real numbers $R_n \to \infty$ such that for all n:

(1) every $U \in \mathcal{U}_n$ has diameter at most R_n;

(2) $\mathcal{U}_{n+1}$ has Lebesgue number at least R_n, i.e. every set of diameter at most R_n is contained in some $U \in \mathcal{U}_{n+1}$.

The *nerve* $|\mathcal{U}|$ of an open cover $\mathcal{U}$ of Y is a simplicial complex with vertex set $\mathcal{U}$ such that $\{U_0, \ldots, U_n\} \subseteq \mathcal{U}$ spans an n-simplex if and only if $\cap_{i=0}^{n} U_i \neq \emptyset$. If $\mathcal{U}_1, \mathcal{U}_2, \ldots$ is an anti-Čech approximation, then for every n, $\mathcal{U}_n$ is a refinement of $\mathcal{U}_{n+1}$, i.e. every $U \in \mathcal{U}_n$ is contained in some element of $\mathcal{U}_{n+1}$; this shows the existence of simplicial inclusion maps $|\mathcal{U}_n| \to |\mathcal{U}_m|$ for every $n \leq m$.

We thus see that $\{C_c^\bullet(|\mathcal{U}_n|; R)\}_{n \in \mathbb{N}}$ is an inverse system whose bonds are induced by the inclusion. Let $\hat{C}^\bullet$ be the inverse limit of this inverse system. Roe defines $HX^*(Y; R)$ to be the cohomology of this inverse limit. Roe shows that $HX^*(Y; R)$ is independent of the choice of anti-Čech approximation, and that a coarse embedding $f : X \to Y$ induces a map $f^* : HX^*(Y; R) \to HX^*(X; R)$. The full details in this are found in [Roe93] and [Roe03].

We can now pick an anti-Čech approximation where for each n, $\mathcal{U}_n := \{B_n(y) \mid y \in Y\}$. The nerve $|\mathcal{U}_n|$ is simply the Rips complex $P_n(Y)$. Thus $HX^*(Y; R)$ is the cohomology of the inverse limit $\hat{C}^\bullet = \varprojlim C_c^\bullet(P_n(Y); R)$. By Theorem 3.5.8 of [Wei94], there is a short exact sequence

$$0 \to \varprojlim{}^1 H_c^{k-1}(P_n(Y); R) \to HX^k(Y; R) \to \varprojlim H_c^k(P_n(Y); R) \to 0, \tag{2.3}$$

where the $\varprojlim^1$ term is known as the *derived limit* and will not be discussed here. This makes it hard to calculate coarse cohomology in general.

However, the situation is simpler if Y is coarsely uniformly acyclic over R. If this is the case, the $\varprojlim^1 H_c^{k-1}(P_n(Y); R)$ term vanishes and $\varprojlim H_c^k(P_n(Y); R)$ is isomorphic to the image of the map $H_c^k(P_j(Y); R) \to H_c^k(P_i(Y); R)$, induced by inclusion, for some $j \gg i \gg 0$. We will not show this now, but it will be shown in the proof of Lemma 3.20. Moreover, Lemma 3.20 will show that the image of the map $H_c^k(P_j(Y); R) \to H_c^k(P_i(Y); R)$ is naturally isomorphic to $H^k_{\text{coarse}}(Y)$. This is also evident from Theorem 5.28 of [Roe03]. Thus the condition that Y is coarsely uniformly acyclic means that we can calculate $HX^k(Y; R)$ without using inverse limits.

By not passing to the inverse limit and simply using either metric complexes, or Rips complexes with sufficiently large parameters, we preserve quantitative information such as the diameter of cocycles and the displacement of maps. We use the notation $H^k_{\text{coarse}}(Y; R)$ rather than $HX^k(Y; R)$ to emphasise this point. We remark that if Y is only coarsely uniformly $(n-1)$-acyclic over R, the preceding discussion still holds in dimensions at most n, taking modified cohomology with compact supports in the top dimension.

2.2.5 Group cohomology and finiteness properties of groups

We recall that a finitely generated group G can be thought of as a bounded geometry metric space by endowing it with the word metric with respect to some finite generating set. Consider a free RG-module M with RG-basis $\mathcal{B}$, so that $G \cdot \mathcal{B}$ is an R-basis of M. We can define a projection $p : G \cdot \mathcal{B} \to G$ so that for each $gb \in G \cdot \mathcal{B}$, $p(gb) = g$. Then $(M, G \cdot \mathcal{B}, p)$ has the structure of a free R-module over G. Moreover, M is finitely generated as an RG-module if and only if it is of finite type. If $f : M \to N$ is an RG-module homomorphism between free RG-modules and M is finitely generated as an RG-module, then f has finite displacement over id_G.

We recall a group is of *type* FP_n^R if the trivial RG-module R has a partial projective resolution of length n consisting of finitely generated RG-modules. The preceding discussion demonstrates that such a partial projective resolution is in fact an n-dimensional, uniformly $(n-1)$-acyclic metric complex with control space G.

Conversely, suppose one has an n-dimensional uniformly $(n-1)$-acyclic metric complex $(G, C_\bullet)$ with control space G, admitting a G-action. This is a partial projective resolution of the trivial RG-module R of length n, consisting of finitely generated RG-modules. By applying Brown's lemma [Bro87], it is straightforward to see that a group is of type FP_n^R if and only if it is coarsely uniformly $(n-1)$-acyclic over R, proving part (2) of Theorem 2.20.

Proposition 2.49 ([Bro82],[GM86],[Roe03])**.** *Let G be a group of type FP_n^R. Then for $i \leq n$, $H^i(G, RG) \cong H^i_{\mathrm{coarse}}(G; R)$ as right RG-modules.*

This follows easily by applying [Bro82, VIII Lemma 7.4] which states that if M is an RG-module, there is a natural isomorphism $\mathrm{Hom}_{RG}(M, RG) \cong \mathrm{Hom}_c(M, R)$. Here, $\mathrm{Hom}_c(M, R)$ consists of all R-module homomorphisms $f : M \to R$ such that for every $m \in M$, $f(gm) = 0$ for all but finitely many $g \in G$. A little more care is required in the $i = n$ case, in which one has to use modified cohomology with compact supports; this is done using an argument identical to one found in [GM86]. Alternatively, in light of the discussion in Section 2.2.4, one can use [Roe03, Example 5.21] which shows that Roe's coarse cohomology is isomorphic to group cohomology with group ring coefficients.

2.2.6 Topology at infinity

To discuss the topology at infinity of spaces, we assume familiarity with inverse and direct limits; the necessary background can be found in [KK05]. We set $R = \mathbb{Z}_2$

and assume all homology and cohomology is taken with coefficients in $\mathbb{Z}_2$. We give a characterisation of coarse cohomology in terms of topology at infinity. We use an argument from [GM85], although it is considerably simpler when working over $\mathbb{Z}_2$.

Let X be a bounded geometry space which is coarsely uniformly n-acyclic over $\mathbb{Z}_2$ and let $(X, C_\bullet)$ be a uniformly n-acyclic $\mathbb{Z}_2$-metric complex. A *finite filtration* of X is a nested sequence $K_1 \subseteq K_2 \subseteq K_3 \subseteq \dots$ of finite subsets of X such that $\cup_i K_i = X$. The k^{th} *reduced homology at infinity* of X is defined to be the inverse limit $\varprojlim_i \tilde{H}_k[X\backslash K_i]$, where the bonds of this inverse system are the maps induced by inclusion. We say that X is *r-acyclic at infinity over $\mathbb{Z}_2$* if $\varprojlim_i \tilde{H}_k[X\backslash K_i]$ vanishes for $k \leq r$. These notions are independent of the choice of finite filtration and metric complex.

Proposition 2.50 ([GM85],[Geo08])**.** *Let $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ be a uniformly n-acyclic $\mathbb{Z}_2$-metric complex, where $n \geq 0$ and X is infinite. For $k \leq n$ and $d \in \mathbb{N}$, $\varprojlim_i \tilde{H}_k[X\backslash K_i]$ is d-dimensional precisely when $H^{k+1}_{\text{coarse}}(X; \mathbb{Z}_2)$ is. In particular for $r \leq n$, X is r-acyclic at infinity over $\mathbb{Z}_2$ if and only if $H^k_{\text{coarse}}(X; \mathbb{Z}_2) = 0$ for $k \leq r+1$.*

Proof. We first show that $H_k[X\backslash K_i]$ is finitely generated for every $i \in \mathbb{N}$ and every $k \leq n$. Let $j > i$ be large enough so that for $k \leq n$, every element of Σ_k lies in either $C_\bullet[K_j]_n$ or $C_\bullet[X\backslash K_i]_n$; Lemma 2.28 demonstrates that such a j exists. A standard argument produces the Mayer–Vietoris sequence

$$\dots \to H_k[K_j\backslash K_i] \to H_k[K_j] \oplus H_k[X\backslash K_i] \to H_k[X] \to \dots .$$

As K_j and $K_j\backslash K_i$ are finite, $H_k[K_j\backslash K_i]$ and $H_k[K_j]$ are finitely dimensional. Since $H_k[X]$ is zero, we see $H_k[X\backslash K_i]$ is a finite dimensional vector space. Let d_i^k be the dimension of $H_k[X\backslash K_i]$.

For each K_i we have a long exact sequence

$$\tilde{H}^0[X] \to \tilde{H}^0[X\backslash K_i] \to H^1[X, X\backslash K_i] \to H^1[X] \to \dots \to H^{n+1}[X].$$

It is well-known that there is an isomorphism $\varinjlim C^\bullet[X, X\backslash K_i] \cong C^\bullet_c[X]$, and that cohomology commutes with direct limits. Since $(X, C_\bullet)$ is uniformly n-acyclic, we obtain the isomorphisms $\varinjlim \tilde{H}^k[X\backslash K_i] \cong \hat{H}^{k+1}_c[X] = \ker(H^{k+1}_c[X] \to H^{k+1}[X])$ for $k \leq n$, showing that $\varinjlim \tilde{H}^k[X\backslash K_i] \cong H^{k+1}_{\text{coarse}}(X; \mathbb{Z}_2)$ for $k \leq n$.

Since we are working over $\mathbb{Z}_2$ and each $\tilde{H}_k[X\backslash K_i]$ is finite dimensional, there are natural isomorphisms $\tilde{H}^k[X\backslash K_i] \cong \text{Hom}_{\mathbb{Z}_2}(\tilde{H}_k[X\backslash K_i], \mathbb{Z}_2)$ and $\tilde{H}_k[X\backslash K_i] \cong$

$\mathrm{Hom}_{\mathbb{Z}_2}(\tilde{H}^k[X\backslash K_i], \mathbb{Z}_2)$. Thus either $\varprojlim_i \tilde{H}_k[X\backslash K_i]$ and $\varinjlim_i \tilde{H}^k[X\backslash K_i]$ are both finite dimensional, in which case there is an isomorphism

$$\varinjlim_i \tilde{H}^k[X\backslash K_i] \cong \mathrm{Hom}_{\mathbb{Z}_2}(\varprojlim_i \tilde{H}_k[X\backslash K_i], \mathbb{Z}_2),$$

or both are infinite dimensional. The result now follows from Remark 2.44, which says that $H^0_{\mathrm{coarse}}(X; \mathbb{Z}_2) = 0$. □

If X is an unbounded coarse geodesic metric space, then Proposition 2.5 says that there is a j such that the inclusion $X \to P_j(X)$ is a coarse isometry. We define the number of ends of X to be equal to the number of ends of $P_j(X)$. Then the above argument shows that the number of ends of X is equal to $\dim(H^1_{\mathrm{coarse}}(X; \mathbb{Z}_2)) + 1$. In particular, X is one-ended precisely when $H^1_{\mathrm{coarse}}(X; \mathbb{Z}_2) = 0$. We think of r-acyclicity at infinity over $\mathbb{Z}_2$ as a higher dimensional analogue of being one-ended.

These ideas are well illustrated in the case of a hyperbolic group G, where the 'topology at infinity' can be interpreted as the topology of the Gromov boundary ∂G. Bestvina and Mess show that when G is hyperbolic, there is an isomorphism $H^k(G, RG) \cong \check{H}^{k-1}(\partial G)$, where $\check{H}^{k-1}(\partial G)$ is the reduced Čech cohomology of the boundary [BM91]. Since $H^k(G, RG) \cong H^k_{\mathrm{coarse}}(G; R)$, this gives a very concrete interpretation of coarse cohomology. See also [Bes96] for a generalized notion of boundary defined in terms of group cohomology.

2.3 Coarse Poincaré duality

2.3.1 Coarse Poincaré duality spaces

We introduce the notion of coarse PD^R_n spaces, which roughly speaking, are spaces which have the same coarse cohomological properties as $\mathbb{R}^n$. The definition we use is found in the appendix of [KK05].

Let X be a bounded geometry metric space. An *R-chain complex over X* is a chain complex $C_\bullet$, such that each C_i is a free R-module over X and each boundary map $\partial : C_i \to C_{i-1}$ has finite displacement. We note that $C_\bullet$ is an R-chain complex over X for every R-metric complex $(X, C_\bullet)$. We extend the notion of finite displacement chain maps and homotopies from R-metric complexes to R-chain complexes over a metric space.

If $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ is a metric complex, then each C^k_c is a finite type free module over X with standard basis Ω_k dual to Σ_k. Each coboundary map $C^k_c \to C^{k+1}_c$ has finite displacement. Thus for any n, $C_c^{n-\bullet}$ is an R-chain complex over X.

We recall that an R-metric complex $(X, C_\bullet)$ admits an augmentation $\varepsilon : C_0 \to R$ such that $\varepsilon(\Delta) = 1_R$ for each $\Delta \in \Sigma_0$. We thus define reduced homology groups $\tilde{H}_k(C_\bullet)$ of this augmented chain complex. Suppose we are given a homomorphism $\alpha : C_c^n \to R$ that is zero on all coboundaries, which we also call an *augmentation of* $C_c^{n-\bullet}$. We then define an augmented cochain complex $\cdots \to C_c^0 \to \cdots \to C_c^n \xrightarrow{\alpha} R \to 0$ and calculate the *reduced cohomology* of this cochain complex, which we denote by $\tilde{H}_c^k(C_\bullet)$.

There is an ambiguity here, since $\tilde{H}_c^k(C_\bullet)$ may also refer to the cohomology of the cochain complex obtained by dualizing the augmented chain complex $\cdots \to C_0 \xrightarrow{\varepsilon} R$ and restricting to cochains of finite support. We do not use this form of reduced cohomology in this section — we always assume that reduced cohomology refers to the former notion, where the choice of a suitable augmentation $\alpha : C_c^n \to R$ is implicit.

Definition 2.51. A *coarse* PD_n^R *complex* consists of a uniformly acyclic R-metric complex $(X, C_\bullet)$, equipped with finite displacement chain maps

$$C_c^{n-\bullet} \xrightarrow{P} C_\bullet \text{ and } C_\bullet \xrightarrow{\bar{P}} C_c^{n-\bullet}$$

over id_X and finite displacement chain homotopies $\bar{P} \circ P \overset{\Phi}{\simeq} \mathrm{id}_{C_c^{n-\bullet}}$ and $P \circ \bar{P} \overset{\bar{\Phi}}{\simeq} \mathrm{id}_{C_\bullet}$ over id_X. We call P and $\bar{P}$ the *duality maps*. We say that X is a *coarse* PD_n^R *space* if it is the control space of some coarse PD_n^R complex.

We say that a group G *acts on a coarse* PD_n^R *complex* if it acts on the underlying metric complex $(X, C_\bullet)$, and the maps P, $\overline{P}$, Φ and $\overline{\Phi}$ are all G-equivariant.

Remark 2.52. If $(X, C_\bullet)$ is a coarse PD_n^R complex, then the cochain complex admits an augmentation $\alpha : C_c^n \to R$ given by the composition $C_c^n \xrightarrow{P} C_0 \xrightarrow{\varepsilon} R$. With respect to this augmentation, the maps P and $\bar{P}$ are *augmentation preserving*, i.e. $\varepsilon(P\sigma) = \alpha(\sigma)$ for all $\sigma \in C_c^n$ and $\varepsilon(\gamma) = \alpha(\bar{P}\gamma)$ for all $\gamma \in C_0$. We note that α induces an isomorphism $H_c^n(C_\bullet) \cong R$.

Proposition 2.53. *Suppose X is a coarse PD_n^R space, $(Y, D_\bullet)$ is a uniformly acyclic R-metric complex and $f : X \to Y$ is a coarse isometry. Then $(Y, D_\bullet)$ is a coarse PD_n^R complex.*

Proof. By definition, there exists a coarse PD_n^R complex $(X, C_\bullet)$ endowed with chain maps and homotopies P, $\overline{P}$, Φ and $\overline{\Phi}$ as above . Since f is a coarse isometry, it has a coarse inverse g. Using Lemma 2.40, we can construct finite displacement chain maps $f_\# : C_\bullet \to D_\bullet$ over f, and $g_\# : D_\bullet \to C_\bullet$ over g, as well as finite displacement chain homotopies $g_\# f_\# \simeq \mathrm{id}_{C_\bullet}$ and $f_\# g_\# \simeq \mathrm{id}_{D_\bullet}$ over id_X and id_Y respectively. By the

proof of Proposition 2.34, $f_\#$ and $g_\#$ induce maps on cochains with compact supports. Lemma 2.21 now tells us that

$$D_c^{n-\bullet} \xrightarrow{f_\# \circ P \circ f^\#} D_\bullet \text{ and } D_\bullet \xrightarrow{g^\# \circ \bar{P} \circ g_\#} D_c^{n-\bullet}$$

are finite displacement chain maps over id_Y. Furthermore, there are finite displacement chain homotopies

$$f_\# \circ P \circ f^\# \circ g^\# \circ \bar{P} \circ g_\# \simeq \mathrm{id}_{D_\bullet} \text{ and } g^\# \circ \bar{P} \circ g_\# \circ f_\# \circ P \circ f^\# \simeq \mathrm{id}_{D_c^{n-\bullet}}$$

over id_Y, which can be explicitly written down in terms of existing homotopies and chain maps. □

Remark 2.54. If X is a coarse PD_n^R space and $f : X \to Y$ is a coarse isometry, then by Propositions 2.15 and 2.42, there necessarily exists a uniformly acyclic $(Y, D_\bullet)$. Thus Proposition 2.53 tells us that being a coarse PD_n^R space is invariant under coarse isometries.

Definition 2.55. A *coarse* PD_n^R *group* is a finitely generated group that, when equipped with a word metric with respect to a finite generating set, is a coarse PD_n^R space.

2.3.2 Poincaré duality groups

For the convenience of the reader, we briefly define PD_n^R groups and relate them to coarse PD_n^R groups. Although knowledge of PD_n^R groups are not needed for our results, PD_n^R groups (and virtually PD_n^R groups) provide a source of examples of coarse PD_n^R groups.

Definition 2.56. A group G is a PD_n^R *group* if:

(1) G is of *type* FP^R, i.e. it has a finite length projective resolution of the trivial RG-module R in which each module is finitely generated;

(2) $H^i(G, RG) = 0$ for $i \neq n$, and $H^n(G, RG) = R$.

A $PD_n^{\mathbb{Z}}$ group G is called *orientable* if the action of G on $H^n(G, \mathbb{Z}G) \cong \mathbb{Z}$ is trivial. For $PD_n^{\mathbb{Z}_2}$ groups, we no longer have to worry about orientation. A $PD_n^{\mathbb{Z}}$ group is automatically a PD_n^R group for any commutative ring R with unity, as shown in Proposition V.3.7 of [DD89].

Proposition 2.57 ([KK05, Section 11.2])**.** *A group G acts freely and cocompactly on a coarse PD_n^R complex if and only if G is a PD_n^R group.*

Example 2.58 ([Bro82, VIII Example 10.1]). The fundamental group of a closed aspherical n-manifold is a $PD_n^{\mathbb{Z}}$ group.

Example 2.59 (Page 166, Example 1 in [Bie81]). If G is a torsion-free polycyclic group of Hirsch length n, then G is a $PD_n^{\mathbb{Z}}$ group.

Every virtually polycyclic group has a finite index subgroup that is a torsion-free polycyclic group. Hence every virtually polycyclic group is a coarse $PD_n^{\mathbb{Z}}$ group. The following example shows that there are coarse PD_n^R groups that are not PD_n^R groups.

Example 2.60. The infinite cyclic group $\mathbb{Z}$ is easily seen to be a $PD_1^{\mathbb{Z}}$ group. It is quasi-isometric to $H := \mathbb{Z}_2 * \mathbb{Z}_2$, which is known not to be a $PD_1^{\mathbb{Z}}$ group since it has torsion so is not of type FP (see [Bro82]). Proposition 2.15 and Proposition 2.42 allow us to construct a uniformly acyclic metric complex $(H, D_\bullet)$ that admits a free H-action. By Proposition 2.53, $(H, D_\bullet)$ is a coarse $PD_1^{\mathbb{Z}}$ complex and hence H is a coarse $PD_1^{\mathbb{Z}}$ group. However, even though H acts freely and cocompactly on $(H, D_\bullet)$ *as a metric complex*, Proposition 2.57 tells us it cannot act freely and cocompactly on $(H, D_\bullet)$ *as a $PD_1^{\mathbb{Z}}$ complex*. In other words, the duality maps and homotopies cannot be H-equivariant.

2.3.3 A technical lemma

The rest of this section will be devoted to a proof of the following Lemma, which plays an important role in the proof of Theorem 1.2.

Lemma 2.61. *Let Y be a bounded geometry metric space that is coarsely uniformly n-acyclic over $\mathbb{Z}_2$ for some $n \geq 0$. Let $X \subseteq Y$ be a subspace that is a $PD_n^{\mathbb{Z}_2}$ space, and let $f : X \to Y$ be the inclusion map. Suppose G is a finitely generated group acting freely on Y such that $GX = X$ and the action of G restricted to X is cocompact. If for $k \leq n$, the maps $f^* : H^k_{\text{coarse}}(Y; \mathbb{Z}_2) \to H^k_{\text{coarse}}(X; \mathbb{Z}_2)$, induced by f, are isomorphisms, then $Y = N_r(X)$ for some $r < \infty$.*

Proof. We first restrict to the case where Y is 1-geodesic. Since Y is coarsely uniformly 0-acyclic, Proposition 2.18 tells us it is coarse geodesic. Hence Proposition 2.5 says that for some s large enough, $P_s^1(Y)$ is connected and the inclusion $Y \to P_s^1(Y)$ is a coarse isometry. We thus remetrise Y by identifying it with the 0-skeleton of $P_s^1(Y)$, endowed with the subspace metric. This procedure doesn't alter our hypotheses.

We use Proposition 2.42 to construct $\mathbb{Z}_2$-metric complexes $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$ and $(Y, D_\bullet, \Sigma'_\bullet, p'_\bullet)$ that are uniformly acyclic and uniformly n-acyclic respectively and admit a G-action. By Proposition 2.53, we observe that $(X, C_\bullet)$ is a $PD_n^{\mathbb{Z}_2}$ complex with (not necessarily G-equivariant) duality maps P and $\bar{P}$ and homotopies Φ and $\bar{\Phi}$.

We may assume that $C_\bullet$ and $D_\bullet$ are $(n+1)$-dimensional by restricting to their $(n+1)$-skelata. By Lemma 2.40, there exists a chain map $f_\# : C_\bullet \to D_\bullet$ that has finite $(n+1)$-displacement over f. This induces a cochain map $f^\# : D_c^{n-\bullet} \to C_c^{n-\bullet}$.

We fix some $\tilde{x} \in X$ and define $K_i^X := N_i^X(\tilde{x})$ and $K_i^Y := N_i^Y(\tilde{x})$ for each $i \in \mathbb{N}$. For every $j \geq i$, we let $\iota_\#^{i,j} : C_\bullet[X \backslash K_j^X] \to C_\bullet[X \backslash K_i^X]$ and $\iota_{i,j}^\# : D_c^\bullet[Y, K_j^Y] \to D_c^\bullet[Y, K_i^Y]$ be the maps induced by inclusion. For each i, there exists a short exact sequence

$$1 \to D_c^{n-\bullet}[Y, K_i^Y] \xrightarrow{q_i^\#} D_c^{n-\bullet}[Y] \xrightarrow{p_i^\#} D_c^{n-\bullet}[K_i^Y] \to 1, \tag{2.4}$$

where $p_i^\#$ is the map induced by inclusion. We note that $q_i^\# \iota_{i,j}^\# = q_j^\#$ for all $i \leq j$.

We let ε and α be the augmentations defined in Remark 2.52. We define an augmentation β by $D_c^n \xrightarrow{f^\#} C_c^n \xrightarrow{\alpha} \mathbb{Z}_2$, and observe that β induces an isomorphism $H_c^n(D_\bullet) \cong \mathbb{Z}_2$. For each $i \in \mathbb{N}$, $\beta \circ q_i^\# : D_c^n[Y, K_i^Y] \to \mathbb{Z}_2$ defines an augmentation, allowing us to define the reduced cohomology $\tilde{H}_c^k(D_\bullet[Y, K_R^Y])$.

By Lemma 2.62, there is an augmentation preserving chain map $g_\# : C_\bullet \to D_c^{n-\bullet}$ of finite displacement over f. By Lemma 2.63, there is a $D \in \mathbb{N}$ such that for every $R \in \mathbb{N}$, $g_\#$ and $h_\# := Pf^\#$ induce augmentation preserving maps

$$g_\#^{R+D} : C_\bullet[X \backslash K_{R+D}^X] \to D_c^{n-\bullet}[Y, K_R^Y] \tag{2.5}$$

$$h_\#^{R+D} : D_c^{n-\bullet}[Y, K_{R+D}^Y] \to C_\bullet[X \backslash K_R^X]. \tag{2.6}$$

Moreover, Lemma 2.63 also tells us that the maps induced on reduced homology and cohomology by (2.5) and (2.6) give the commutative diagram shown in Figure 2.1, whose horizontal maps are induced by inclusion.

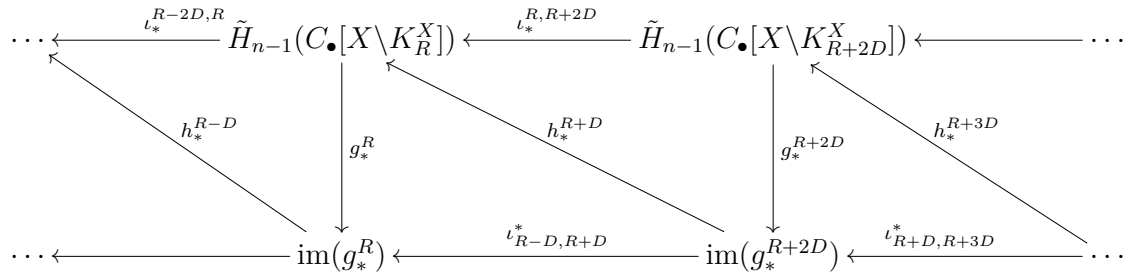


Figure 2.1: A commutative diagram relating the topology at infinity of X and Y

Since $D_\bullet$ is uniformly 0-acyclic, there is an R_0 such that for all $y \in Y$, the map $\tilde{H}_0(D_\bullet[N_1^Y(y)]) \to \tilde{H}_0(D_\bullet[N_{R_0}^Y(y)])$, induced by inclusion, is zero. By Corollary 2.29, there is some R_1 such that $g_\#(C_\bullet[X]_n) \subseteq D_c^{n-\bullet}[N_{R_1}^Y(X)]$. As $H_c^n(C_\bullet) \cong \mathbb{Z}_2$, Proposition 2.50 tells us $(\tilde{H}_{n-1}(C_\bullet[X \backslash K_i^X]))_{i=0}^\infty$ has inverse limit $\mathbb{Z}_2$. Hence there is an $i_0 \geq 0$ such that for all $i \geq i_0$, there is a non-trivial reduced cycle $\lambda_i \in C_{n-1}[X \backslash K_{i+D}^X]$ such that for all $j \geq i$, $\iota_*^{i+D,j+D}[\lambda_j] = [\lambda_i]$.

We let $\sigma_i = g_\#^{i+D}\lambda_i \in D_c^1[Y, K_i^Y]$. It follows from the way $g_\#^{i+D}$ is defined in Lemma 2.63, that $\mathrm{supp}(q_i^\# \sigma_i) \subseteq N_{R_1}^Y(X) \backslash K_i^Y$. We deduce from Figure 2.1 that $\iota_{i,j}^*[\sigma_j] = [\sigma_i]$ for all $j \geq i \geq i_0$. Let $i_1 := i_0 + 2D$. For $i \geq i_1$, as

$$[\iota_\#^{i_0+D,i-D} h_\#^i g_\#^{i+D} \lambda_i] = [\iota_\#^{i_0+D,i+D} \lambda_i] = [\lambda_{i_0}] \neq 0,$$

we see that $[\sigma_i] = [g_\#^{i+D}\lambda_i] \neq 0$.

We choose $i \geq i_1$ large enough so $GK_i^X = X$, and let $j = i + 2R_0 + R_1 + 1$. There is an $\omega \in C_c^0[Y]$ such that $\delta\omega = q_j^\# \sigma_j$. This is because $\tilde{H}_c^1(D_\bullet) = 0$ and, since $g_\#^{j+D}$ is augmentation preserving, $q_j^\# \sigma_j$ is a reduced cocycle. We let $R_2 := \mathrm{diam}(\omega)$.

We claim that $\omega(\Delta) = 1$ for every $\Delta \in {p_0'}^{-1}(K_i^Y)$. If $p_i^\# \omega = 0$, then by the exactness of (2.4), $\omega = q_i^\#(\omega_i)$ for some $\omega_i \in D_c^0[Y, K_i^Y]$. Thus $\iota_{i,j}^\# \sigma_j = \delta\omega_i$, which cannot happen as $[\iota_{i,j}^\# \sigma_j] = [\sigma_i]$ is non-trivial. Thus $p_i^\# \omega \neq 0$, so we can pick $\Delta_0 \in {p_0'}^{-1}(K_i^Y)$ such that $\omega(\Delta_0) = 1$.

Let $\Delta \in {p_0'}^{-1}(K_i^Y)$. Then there exists a 1-chain $p_0'(\Delta_0) = x_0, x_1, \ldots, x_n = p_0'(\Delta)$ in K_i^Y. This is because Y is 1-geodesic and $K_i^Y = N_i^Y(\tilde{x})$. For $1 \leq k < n$, we pick some $\Delta_k \in (p_0')^{-1}(x_k)$ and let $\Delta_n = \Delta$. Each $\Delta_{k+1} - \Delta_k$ is a reduced 0-cycle in $D_\bullet[N_1^Y(x_k)]$, hence there is some $\rho_k \in D_1[N_{R_0}^Y(x_k)] \subseteq D_1[K_j^Y]$ such that $\partial\rho_k = \Delta_{k+1} - \Delta_k$. Letting $\rho := \sum_{k=0}^{n-1} \rho_k \in C_1[K_j^Y]$, we see $\partial\rho = \Delta - \Delta_0$, and so

$$\omega(\Delta) - \omega(\Delta_0) = \omega(\partial\rho) = (q_j^\# \sigma_j)(\rho) = 0.$$

Thus $\omega(\Delta) = \omega(\Delta_0) = 1$, proving the claim.

We define a function $Q : X \to \mathbb{R} \cup \infty$ that measures the distance from a point of X to the furthest point of Y that isn't closer to any other point of X. More precisely, for each $y \in Y$, we define $\mathrm{height}(y) := d_Y(y, X) = \inf\{d_Y(x, y) \mid x \in X\}$. For $x \in X$, we define

$$Q(x) := \sup_{y \in Y}\{\mathrm{height}(y) \mid \mathrm{height}(y) = d_Y(x, y)\}.$$

Suppose $\sup_{x \in X} Q(x) \leq r < \infty$ and $y \in Y$. Then there exists an $x \in X$ such that $\mathrm{height}(y) = d_Y(x, y)$. Indeed, for any $x' \in X$, the set $\{x \in X \mid d_Y(y, x) \leq d_Y(y, x')\}$ is non-empty and finite. Hence there exists an $x \in X$ such that $d_Y(x, y) \leq d_Y(x'', y)$

for all $x'' \in X$, so height$(y) = d_Y(x, y)$. Since $d_Y(x, y) \leq Q(x) \leq r$, we see that $y \in N_r^Y(X)$. This holds for each $y \in Y$, so $Y = N_r^Y(X)$. As G acts on Y by isometries, Q is G-equivariant. Since $GK_i^X = X$, it is thus sufficient to show that Q is bounded on K_i^X.

Let $x \in K_i^X$; we claim that $Q(x) \leq R_2$. Suppose for contradiction that there is some $y \in Y$ such that height$(y) = d_Y(y, x) > R_2$. We choose a 1-geodesic $x = w_0, w_1, \ldots, w_n = y$, and pick $\Lambda_k \in p_0'^{-1}(w_k)$. For each $0 \leq k < n$, there exists a $\mu_k \in D_1[N_{R_0}^Y(w_k)]$ such that $\partial\mu_k = \Lambda_{k+1} - \Lambda_k$.

As $x \in K_i^X \subseteq K_i^Y$, we know that $\omega(\Lambda_0) = 1$. Since $d(x, y) > R_2 = \text{diam}(\omega)$, we see that $\omega(\Lambda_n) = 0$. Therefore, $(q_j^\# \sigma_j)(\sum_{k=0}^{n-1} \mu_k) = \omega(\Lambda_n) - \omega(\Lambda_0) = 1$. Hence there exists some $0 \leq t < n$ such that $(q_j^\# \sigma_j)(\mu_t) = 1$. Since $\text{supp}(q_j^\# \sigma_j) \subseteq N_{R_1}^Y(X) \backslash K_j^Y$ and $\mu_t \in D_1[N_{R_0}^Y(w_t)]$, we deduce that $w_t \notin K_{j-R_0}^Y$ and that there is some $z \in X$ such that $d(z, w_t) \leq R_0 + R_1$. As $w_0 = x \in K_i^X$, it follows that $t = d(w_0, w_t) \geq (j - R_0) - i = R_0 + R_1 + 1$. Hence

$$d(z, y) \leq d(z, w_t) + d(w_t, w_n) \leq (R_0 + R_1) + (n - t) \leq n - 1 < d(x, y),$$

which contradicts height$(y) = d(x, y)$. Thus $Q(x) \leq R_2$ for every $x \in K_i^X$. □

The following two lemmas freely use notation from the preceding proof.

Lemma 2.62. *There is an augmentation preserving chain map $g_\# : C_\bullet \to D_c^{n-\bullet}$ of finite displacement over f, such that $Pf^\# g_\#$ is chain homotopic to the identity via a finite displacement chain homotopy Λ.*

Proof. Since G acts cocompactly on $(X, C_\bullet, \Sigma_\bullet, p_\bullet)$, Σ_k has only finitely many G-orbits; we thus let $x_1^k, \ldots, x_{t_k}^k$ be a set of representatives for the G-orbits of Σ_k. We define $g_\#$ on each x_i^k and extend equivariantly. The finite displacement of $g_\#$ then follows readily from G-equivariance. Indeed, one can choose D_k large enough so that $\text{supp}(g_\#(x_i^k)) \subseteq N_{D_k}^Y(f(p_k(x_i^k)))$ for each x_i^k. Thus for each $h \in G$ and x_i^k, we see

$$\text{supp}(g_\#(hx_i^k)) \subseteq hN_{D_k}^Y(f(p_k(x_i^k))) = N_{D_k}^Y(f(p_k(hx_i^k))), \tag{2.7}$$

showing that $g_\#$ does indeed have finite displacement over f.

As the augmentation $\beta : D_c^n \to \mathbb{Z}_2$ induces an isomorphism $H_c^n(D_\bullet) \to \mathbb{Z}_2$, we may choose an n-cycle $\sigma \in D_c^n$ such that $\beta(\sigma) = 1$. Each $h \in G$ preserves cocycles and coboundaries, so preserves the unique non-trivial cohomology class in $H_c^n(D_\bullet)$; thus $[h\sigma] = [\sigma]$, so $\beta(h\sigma) = \beta(\sigma) = 1$.

For each x_i^0, we define $g_\#(x_i^0) = \sigma$ and extend equivariantly. We note that $g_\#$ is augmentation preserving, since for each $h \in G$,

$$\beta(g_\#(hx_i^0)) = \beta(hg_\#(x_i^0)) = \beta(h\sigma) = 1 = \varepsilon(hx_i^0).$$

For each x_i^1, we see that $\delta g_\#(\partial x_i^1) = 0$ and $\beta(g_\#\partial x_i^1) = \varepsilon(\partial x_i^1) = 0$. Therefore, as β induces an isomorphism $H_c^n(D_\bullet) \cong \mathbb{Z}_2$, there exists an $\omega_i^1 \in D_c^{n-1}$ such that $\delta\omega_i^1 = g_\#(\partial x_i^1)$. We define $g_\#(x_i^1) := \omega_i^1$ for each x_i^1 and extend equivariantly. We continue similarly for higher dimensions, using the fact $H_c^k(D_\bullet) = 0$ for $k < n$.

We now claim that $f^\# g_\# : C_\bullet \to C_c^{n-\bullet}$ is a chain map of finite displacement over id_X. If the claim is true, then by Lemma 2.21, $Pf^\# g_\#$ also has finite displacement over id_X. By Lemma 2.40, we thus see that $Pf^\# g_\#$ is chain homotopic to the identity chain map via a finite displacement chain homotopy Λ.

To prove the claim, we suppose $x \in \mathrm{supp}(f^\# g_\# h x_i^k)$. Then there exists a $\Delta \in \Sigma_{n-k}$ such that $p_{n-k}(\Delta) = x$ and $f^\# g_\# h x_i^k(\Delta) \neq 0$; thus $g_\#(hx_i^k)(f_\#\Delta) \neq 0$. We see that $\mathrm{supp}(g_\#(hx_i^k)) \subseteq N_{D_k}^Y(f(p_k(hx_i^k)))$ by (2.7). As $f_\#$ has n-displacement at most R over f for some $R \geq 0$, we see that $\mathrm{supp}(f_\#\Delta) \subseteq N_R^Y(p_{n-k}(\Delta))$. Therefore, as f is the inclusion map

$$d_X(x, p_k(hx_i^k)) = d_Y(f(p_{n-k}(\Delta)), f(p_k(hx_i^k))) \leq R + D_k.$$

Hence $\mathrm{supp}(f^\# g_\# hx_i^k) \subseteq N_{R+D_k}^X(p_k(hx_i^k))$, so $f^\# g_\#$ has finite displacement over id_X. □

Lemma 2.63. *Let $h_\# := Pf^\# : D_c^{n-\bullet} \to C_\bullet$. There is a $D \in \mathbb{N}$, such that for every $R \in \mathbb{N}$, $g_\#$ and $h_\#$ induce maps $g_\#^{R+D}$ and $h_\#^{R+D}$ as in (2.5) and (2.6). These maps then induce the commutative diagram as shown in Figure 2.1.*

Proof. We use similar methods to those used in Section 6 of [KK05]. We pick $D_0 \in \mathbb{N}$ such that $g_\#$ and $f_\#$ have n-displacement at most D_0 over f, while P, Λ and the boundary maps of $C_\bullet$ and $D_\bullet$ all have n-displacement at most D_0 over the identity; we now define $D = 2(n+1)D_0$. For each $i \in \mathbb{N}$ we define $q_i^\#$ and $p_i^\#$ as in (2.4), and let $r_\#^i : C_\bullet[X\backslash K_i^X] \to C_\bullet$ be the inclusion map.

Let $k \leq n$. If $\sigma \in C_k[X\backslash K_{R+D}^X]$, then $\mathrm{supp}(r_\#^{R+D}\sigma) \subseteq X\backslash K_{R+D}^X$, so

$$\mathrm{supp}(g_\# r_\#^{R+D}\sigma) \subseteq N_{D_0}^Y(X\backslash K_{R+D}^X) \subseteq Y\backslash K_{R+D-D_0}^Y \subseteq Y\backslash K_R^Y.$$

Therefore $p_R^\# g_\# r_\#^{R+D}\sigma = 0$, so by the exactness of (2.4) there exists an $\omega_\sigma \in D_c^{n-k}[Y, K_R^Y]$ such that $q_R^\#\omega_\sigma = g_\# r_\#^{R+D}\sigma$. As $q_R^\#$ is injective, we can define an

augmentation preserving chain map $g_{\#}^{R+D} : C_\bullet[X\backslash K_{R+D}^X]_n \to D_c^{n-\bullet}[Y, K_R^Y]$ such that $q_R^{\#} g_{\#}^{R+D} = g_{\#} r_{\#}^{R+D}$.

Similarly, we define $h_{\#}^{R+D} : D_c^{n-\bullet}[Y, K_{R+D}^Y]_n \to C_\bullet[X\backslash K_R^X]$ so that $r_{\#}^R h_{\#}^{R+D} = Pf^{\#} q_{R+D}^{\#}$. Indeed, suppose $\rho \in D_c^k[Y, K_{R+D}^Y]$ for $k \leq n$. If $\Delta \in (p_k')^{-1}(K_{R+D-nD_0}^Y)$, then by Lemma 2.28, $\Delta \in D_k[K_{R+D-nD_0+kD_0}^Y] \subseteq D_k[K_{R+D}^Y]$, so that $q_{R+D}^{\#}(\rho)(\Delta) = 0$. Therefore, $\mathrm{supp}(q_{R+D}^{\#}\rho) \subseteq Y\backslash K_{R+D-nD_0}^Y$ and so

$$\mathrm{supp}(Pf^{\#} q_{R+D}^{\#}\rho) \subseteq X\backslash K_{R+D-(n+2)D_0}^X = X\backslash K_{R+nD_0}^X.$$

By another application of Lemma 2.28, we see that $Pf^{\#} q_{R+D}^{\#}\rho = r_{\#}^R \gamma_\rho$ for some $\gamma_\rho \in C_\bullet[X\backslash K_R^X]$; this allows us to define $h_{\#}^{R+D}$.

By a similar argument, Λ induces a map $\Lambda^{R+2D} : C_\bullet[X\backslash K_{R+2D}]_n \to C_{\bullet+1}[X\backslash K_R]$ such that $\Lambda r_{\#}^{R+2D} = r_{\#}^R \Lambda^{R+2D}$; it follows that $h_{\#}^{R+D} g_{\#}^{R+2D} - \iota_{\#}^{R,R+2D} = \partial\Lambda^{R+2D} + \Lambda^{R+2D}\partial$. This proves the existence of the commutative diagram shown in Figure 2.1. □

Chapter 3

Coarse detection of group splittings

3.1 Coarse separation

3.1.1 Coarse complementary components

We introduce the notion of coarse complementary components, generalising the notion of an almost invariant set (see Proposition 3.14). It should be remarked that a coarse complementary component C is not necessarily coarsely connected, i.e. $P_r(C)$ may not be connected for any r. Lemmas 3.3 and 3.5 provide some motivation for the term 'coarse complementary components'.

Definition 3.1. Let X be a 1-geodesic metric space, $C \subseteq X$ and $r \geq 0$. We define the *coarse r-boundary* of C to be

$$\partial_r C := \{x \in X \backslash C \mid d(x, C) \leq r\}.$$

If $W \subseteq X$, $r \geq 1$ and $A \geq 0$, we say that $C \subseteq X$ is an (r, A)-*coarse complementary component* of W if

$$\partial_r(C \backslash N_A(W)) \subseteq N_A(W).$$

We say that C is a *coarse complementary component* of W if it is an (r, A)-*coarse complementary component* of W for some $r \geq 1$ and $A \geq 0$.

Remark 3.2. If $\partial_r(C) \subseteq N_A(W)$, then C is an (r, A)-coarse complementary component of W. Conversely, if C is an (r, A)-coarse complementary component of W, then $\partial_r(C) \subseteq N_{A+r}(W)$.

Lemma 3.3. *Let X be a 1-geodesic metric space and $W \subseteq X$. Then $C \subseteq X$ is an (r, A)-coarse complementary component of W if and only if $C \backslash N_A(W)$ is the vertex set of a union of components of $P_r(X) \backslash P_r(N_A(W))$.*

Proof. Assume C is an (r, A)-coarse complementary component of W. For $x \in C$ and $y \in X$, suppose x and y lie in the same component of $P_r(X)\backslash P_r(N_A(W))$. Then there exists an r-chain $x = x_0, x_1, \ldots, x_n = y$ in $X\backslash N_A(W)$. If some $x_i \in C$ and $x_{i+1} \notin C$, then $x_{i+1} \in \partial_r(C\backslash N_A(W)) \subseteq N_A(W)$ which contradicts $x_{i+1} \in X\backslash N_A(W)$. Therefore $y \in C$.

Now suppose that $C\backslash N_A(W)$ is the vertex set of a union of components of $P_r(X)\backslash P_r(N_A(W))$. If $x \in \partial_r(C\backslash N_A(W))$, then there is some $y \in C\backslash N_A(W)$ such that $d(x, y) \leq r$. If $x \notin N_A(W)$, then x and y lie in the same component of $P_r(X)\backslash P_r(N_A(W))$; thus $x \in C\backslash N_A(W)$, which is a contradiction. Therefore $x \in N_A(W)$, showing that $\partial_r(C\backslash N_A(W)) \subseteq N_A(W)$ and hence C is an (r, A)-coarse complementary component of W. □

Corollary 3.4. *The complement, union, intersection or symmetric difference of (r, A)-coarse complementary components of W is itself an (r, A)-coarse complementary component of W.*

Lemma 3.5. *Let X and Y be 1-geodesic metric spaces with $W \subseteq X$. Suppose $C \subseteq X$ is an (r, A)-coarse complementary component of W and $f : X \to Y$ is an (η, ϕ, B)-coarse isometry. Then for each $s \geq 1$, there exists an $A' = A'(A, \eta, \phi, B, s)$ such that $N_B(f(C))$ is an (s, A')-coarse complementary component of $f(W)$.*

Proof. Let $Z := f(W)$ and $Q := N_B(f(C))$. Setting $A' := \phi(A + \tilde{\eta}(s + 2B)) + B$, we claim that Q is an (s, A') coarse complementary component of Z, where $\tilde{\eta}$ is as defined in Remark 2.2. We need to show $\partial_s(Q\backslash N_{A'}(Z)) \subseteq N_{A'}(Z)$.

Let $y \in \partial_s(Q\backslash N_{A'}(Z))$; there exists an $x \in X$ such that $d_Y(f(x), y) \leq B$. If $x \in C$, then $y \in Q$; hence $y \in N_{A'}(Z)$ and we are done. We therefore suppose $x \notin C$. Then there exists a $y' \in Q\backslash N_{A'}(Z)$ such that $d_Y(y, y') \leq s$. We thus pick $x' \in C$ such that $d_Y(f(x'), y') \leq B$. Since $A' \geq \phi(A) + B$ and $y' \notin N_{A'}(Z)$, it follows that $x' \in C\backslash N_A(W)$. As $d_Y(y, y') \leq s$, we see that $d_X(x, x') \leq \tilde{\eta}(s + 2B)$.

We now observe that x' and x can be joined by a 1-geodesic $x' = x_0, \ldots, x_n = x$ such that $n = d_X(x, x') \leq \tilde{\eta}(s+2B)$. We pick the minimal t such that $x_t \notin C\backslash N_A(W)$ and note that

$$d_X(x, x_t) = n - t \leq \tilde{\eta}(s + 2B).$$

By the minimality of t, $x_t \in \partial_r(C\backslash N_A(W)) \subseteq N_A(W)$, so there exists a $w \in W$ such that $d_X(w, x_t) \leq A$. Therefore

$$d_Y(f(w), y) \leq d_Y(f(w), f(x)) + B \leq \phi(A + \tilde{\eta}(s + 2B)) + B = A',$$

showing that $y \in N_{A'}(Z)$. □

We say that C is an *irreducible* (r, A)*-coarse complementary component* whenever $C \backslash N_A(W)$ is the vertex set of a single component of $P_r(X) \backslash P_r(N_A(W))$. However, unlike general coarse complementary components, irreducible coarse complementary components are not preserved under coarse isometries.

Corollary 3.6. *For any finite collection $\{C_i\}_{i \in I}$ of coarse complementary components of W and any $r \geq 1$, there exists an $A \geq 0$ such that every C_i is an (r, A)-coarse complementary component of W.*

Proof. If C is an (r, A)-coarse complementary component of W, then it is an (r, A')-coarse complementary component of W for any $A' \geq A$. Applying Lemma 3.5 with $f = \mathrm{id}_X$, we see that every C_i is an (r, A_i) coarse complementary component for some A_i. Letting $A := \max_{i \in I}(A_i)$, we see that each C_i is an (r, A)-coarse complementary component of W. □

Since every coarse complementary component is a $(1, A)$-coarse complementary component for some $A \geq 0$, one might think that the 'r' parameter in Definition 3.1 is redundant. The following Lemma, which plays a crucial role in the proofs of Propositions 3.8 and 3.19, motivates the need for two parameters in Definition 3.1.

Lemma 3.7. *Let $C \subseteq X$ be an (r, A)-coarse complementary component of W. Then every simplex of $P_r(X)$ lies in either $P_r(N_A(W) \cup C)$ or $P_r(N_A(W) \cup (X \backslash C))$.*

Proof. Let $\Delta = \{x_0, \ldots, x_n\}$ be a simplex of $P_r(X)$. If each $x_i \in N_A(W)$, then Δ is a simplex of both $P_r(N_A(W) \cup C)$ and $P_r(N_A(W) \cup (X \backslash C))$. We therefore assume $x_i \notin N_A(W)$ for some $x_i \in \Delta$. Interchanging C and $X \backslash C$ if necessary, we may assume $x_i \in C \backslash N_A(W)$. For each $x_j \in \Delta$, either $x_j \in C \backslash N_A(W)$ or $x_j \in \partial_r(C \backslash N_A(W)) \subseteq N_A(W)$. Thus Δ is a simplex of $P_r(N_A(W) \cup C)$. □

We now show that if C is a coarse complementary component of W and that both X and W are coarsely uniformly $(n-1)$-acyclic, then $W \cup C$ is also coarsely uniformly $(n-1)$-acyclic. This allows us to define the coarse cohomology of $W \cup C$.

Proposition 3.8. *Let X be a 1-geodesic metric space with $W \subseteq X$, such that both X and W are (λ, μ)-coarsely uniformly $(n-1)$-acyclic over R for some $n > 0$. If C is a $(1, A)$-coarse complementary component of W, then $W \cup C$ is (λ', μ')-coarsely uniformly $(n-1)$-acyclic over R, where λ' and μ' depend only on A, λ and μ.*

To prove this, we use the following lemma:

Lemma 3.9. *Let X be a bounded geometry metric space with $A \subseteq B \subseteq X$. Suppose there exists some $t \geq 0$ such that $B \subseteq N_t(A)$. If $f_\# : C_\bullet(P_i(A); R) \to C_\bullet(P_i(B); R)$ is the inclusion of chain complexes, then there is a chain map*

$$p_\# : C_\bullet(P_i(B); R) \to C_\bullet(P_{i+2t}(A); R)$$

such that $p_\# f_\#$ is the inclusion $C_\bullet(P_i(A); R) \to C_\bullet(P_{i+2t}(A); R)$ and

$$p_\#(C_\bullet(P_i(N_r^B(x)); R)) \subseteq C_\bullet(P_{i+2t}(N_{r+t}^A(x); R))$$

for every $x \in X$ and $r \geq 0$.

Proof. We choose a closest point projection $p : B \to A$, noting that for each $b \in B$, $d(b, p(b)) \leq t$. Then p induces the chain map $p_\# : C_\bullet(P_i(B); R) \to C_\bullet(P_{i+2t}(A); R)$ given by $[x_0, \dots x_n] \mapsto [p(x_0), \dots, p(x_n)]$ on each oriented simplex $[x_0, \dots x_n]$. Since $p|_A = \mathrm{id}_A$, $p_\# f_\#$ is just the inclusion of chain complexes. □

Proof of Proposition 3.8. Without loss of generality, we may assume $\lambda(i) \geq 1$ for every i. For any $k < n$, $i \geq 0$, $r \geq 0$ and $x \in X$, let $\sigma \in C_k(P_i(N_r^{C \cup W}(x)); R)$ be a reduced cycle. There exists a $(k+1)$-chain $\omega \in C_{k+1}(P_{\lambda(i)}(N^X_{\mu(i,r)}(x)); R)$ such that $\partial\omega = \sigma$.

By Corollary 3.6, there is an $A' = A'(\lambda(i), A) \geq 0$ such that C is a $(\lambda(i), A')$-coarse complementary component of W. By Lemma 3.7, each simplex of $P_{\lambda(i)}(X)$ lies in either $P_{\lambda(i)}(N^X_{A'}(W) \cup C)$ or $P_{\lambda(i)}(N^X_{A'}(W) \cup (X \backslash C))$. Therefore, $\omega = \omega_1 + \omega_2$, where $\omega_1 \in C_{k+1}(P_{\lambda(i)}(N^{N^X_{A'}(W) \cup C}_{\mu(i,r)}(x)); R)$ and $\omega_2 \in C_{k+1}(P_{\lambda(i)}(N^{N^X_{A'}(W) \cup (X \backslash C)}_{\mu(i,r)}(x)); R)$. As $\partial\omega_2 = \partial\omega - \partial\omega_1 = \sigma - \partial\omega_1$, we see that $\partial\omega_2 \in C_k(P_{\lambda(i)}(N^{N^X_{A'}(W)}_{\mu(i,r)}(x)); R)$ is a reduced k-cycle. The inclusion $W \to N^X_{A'}(W)$ is an $(\mathrm{id}_{\mathbb{R}_{\geq 0}}, \mathrm{id}_{\mathbb{R}_{\geq 0}}, A')$-coarse isometry. Consequently, Proposition 2.15 tells us that $N^X_{A'}(W)$ is (λ', μ') coarsely uniformly $(n-1)$-acyclic, where λ' and μ' depend only on A, λ and μ. Hence there is a $\gamma \in C_{k+1}(P_{i'}(N^{N^X_{A'}(W)}_{r'}(x)); R)$ such that $\partial\gamma = \partial\omega_2$, where $i' = \lambda'(\lambda(i))$ and $r' = \mu'(\lambda(i), \mu(i, r))$.

Let $f_\#$ be the inclusion $C_\bullet(P_{i'}(C \cup W); R) \to C_\bullet(P_{i'}(C \cup N^X_{A'}(W)); R)$. As $C \cup N^X_{A'}(W) \subseteq N^X_{A'}(C \cup W)$, Lemma 3.9 tells us there is a chain map

$$p_\# : C_\bullet(P_{i'}(C \cup N^X_{A'}(W)); R) \to C_\bullet(P_{i'+2A'}(C \cup W); R)$$

such that $p_\# f_\#$ is the inclusion. We see that

$$p_\#(\omega_1 + \gamma) \in C_{k+1}(P_{i'+2A'}(N^{C \cup W}_{r'+A'}(x)); R)$$

and $\partial p_\#(\omega_1 + \gamma) = p_\# f_\#(\sigma) = \sigma$. Therefore, $C \cup W$ is coarsely uniformly $(n-1)$-acyclic over R, and the uniform acyclicity functions depend only on A, λ and μ. □

3.1.2 Coarse separation

We now describe a generalisation of the notion of coarse separation, originally due to Farb and Schwartz [FS96] and further developed by Kapovich and Kleiner [KK05].

A coarse complementary component C of W is said to be *shallow* if $C \subseteq N_R(W)$ for some $R \geq 0$. A coarse complementary component that is not shallow is called *deep*. A collection $\{C_i\}_{i \in I}$ of coarse complementary components of W is said to be *coarse disjoint* if for every $i \neq i' \in I$, $C_i \cap C_{i'}$ is shallow. Similarly, we say that a component of $P_r(X) \backslash P_r(N_A(W))$ is *deep* (resp. *shallow*) if its vertex set is deep (resp. shallow). A coarse complementary component C is said to be *non-trivial* if both C and its complement $C^* := G \setminus C$ are deep.

We say that W *coarsely n-separates* X if X has at least n deep, coarse disjoint, coarse complementary components of W. We say that W *coarsely separates* X if W coarsely 2-separates X.

Lemma 3.10. *Let X be a 1-geodesic metric space and suppose C is an (r, A)-coarse complementary component of W for some $r \geq 1$ and $A \geq 0$. Then for all $R \geq 0$,*

$$N_R(C) \subseteq C \cup N_{A+R}(W).$$

Proof. Suppose $x \in N_R(C) \backslash C$. There exists an $x_0 \in C$ with $d(x_0, x) \leq R$. Hence there is a 1-geodesic $x_0, \ldots, x_n = x$ with $n = d(x_0, x) \leq R$. We pick the minimal t such that $x_t \notin C \backslash N_A(W)$ and claim $x_t \in N_A(W)$. If $t = 0$, then because $x_0 \in C$, $x_0 \in N_A(W)$. If $t > 0$, then $x_t \in \partial_r(C \backslash N_A(W)) \subseteq N_A(W)$. As $d(x, x_t) = n - t \leq R$, we see that $x \in N_R(N_A(W)) \subseteq N_{A+R}(W)$. □

Proposition 3.11. *Suppose X and Y are 1-geodesic and $f : X \to Y$ is an (η, ϕ, B)-coarse isometry. If $W \subseteq X$ coarsely n-separates X, then $f(W)$ coarsely n-separates Y.*

Proof. Since W coarsely n-separates X, there exist n deep, coarse disjoint, coarse complementary components of W, which we denote $C_1, \ldots, C_n$. By Lemma 3.5, we need only show that $N_B(f(C_1)), \ldots, N_B(f(C_n))$ are deep and coarse disjoint. If $N_B^Y(f(C_i)) \subseteq N_R^Y(f(W))$ for some $R \geq 0$, then $C_i \subseteq N_{\tilde{\eta}(R)}^X(W)$. Consequently, as each C_i is deep, so is each $N_B(f(C_i))$.

For each $1 \leq i < j \leq n$, we choose $R \geq 0$ such that $C_i \cap C_j \subseteq N_R(W)$. By Corollary 3.6, there is some $r \geq 1$ and $A \geq 0$ such that each C_j is an (r, A)-coarse complementary component of W. Suppose $y \in N_B(f(C_i)) \cap N_B(f(C_j))$. We pick

$c_i \in C_i$ and $c_j \in C_j$ such that $d_Y(f(c_i), y), d_Y(f(c_j), y) \leq B$; hence $d_X(c_i, c_j) \leq \tilde{\eta}(2B)$. By Lemma 3.10

$$c_i \in C_i \cap N_{\tilde{\eta}(2B)}(C_j) \subseteq C_i \cap (C_j \cup N_{A+\tilde{\eta}(2B)}(W)) \subseteq N_{A'}(W),$$

where $A' := \max(R, A + \tilde{\eta}(2B))$. Hence $y \in N_{\phi(A')+B}(f(W))$, so $N_B(f(C_i))$ and $N_B(f(C_j))$ are coarse disjoint. □

There is another notion of coarse separation, defined in terms of actual complementary components rather than coarse complementary components: if Γ is the Cayley graph of a group G with respect to some finite generating set, we say that W coarsely separates G if for some $R \geq 0$, $\Gamma \backslash N_R(W)$ has at least two deep components, i.e. two components not contained in $N_S(W)$ for any $S \geq 0$. This is the definition used in [KK05].

The two notions of coarse separation are not necessarily equivalent. To see why, we observe that if C is an (r, A)-coarse complementary component of W, then C is a union $\cup_{i \in I} C_i$ of irreducible (r, A)-coarse complementary components. It is possible for C to be deep and for each C_i to be shallow. The following lemma rules out this behaviour when W is a subgroup of a finitely generated group. More generally, whenever W satisfies the *shallow condition* as defined in [Vav12], then the two notions of coarse separation are equivalent.

Lemma 3.12 ([Vav12, Lemma 7.1])**.** *Let G be a finitely generated group and let $H \leq G$ be a subgroup. Then for all $r \geq 1$ and $A \geq 0$, there exists an $R \geq A$ such that every shallow (r, A)-coarse complementary component of H is contained in $N_R(H)$.*

Proof. Using Lemma 3.3, it is enough to show there is an $R \geq 0$ such that every shallow component of $P_r(G) \backslash P_r(N_A(H))$ is contained in $P_r(N_R(H))$. As G has bounded geometry there are only finitely many components of $P_r(G) \backslash P_r(N_A(H))$ that intersect the ball $N_{A+r}(e)$. Hence there is an $R \geq 0$ such that every shallow component of $P_r(G) \backslash P_r(N_A(H))$ that intersects $N_{A+r}(e)$ is contained in $P_r(N_R(H))$. As $P_r(G)$ is connected, an arbitrary shallow component C of $P_r(G) \backslash P_r(N_A(H))$ intersects the ball $N_{A+r}(h)$ for some $h \in H$. Since $h^{-1}C$ intersects $N_{A+r}(e)$, it is contained in $P_r(N_R(H))$; therefore $C \subseteq hP_r(N_R(H)) = P_r(N_R(H))$. □

3.1.3 Relative ends of groups

We now examine two notions of relative ends of groups, due to Houghton–Scott [Hou74] and Kropholler–Roller [KR89] respectively. Suppose G is a finitely generated group

and $\mathcal{P}(G)$ is the powerset of G. Then $\mathcal{P}(G)$ is a vector space over $\mathbb{Z}_2$ in which the addition operation, denoted $\triangle$, is the symmetric difference of two subsets. If H is a subgroup of G, let $\mathcal{F}_H(G)$ be the subspace of all H-finite subsets of G.

Definition 3.13. We say that $C \subseteq G$ is *H-almost invariant* (or H-a.i.) if for all $g \in G$, $C \triangle Cg$ is H-finite. An H-a.i. subset C is *non-trivial* if neither C nor $G \backslash C$ is H-finite.

In other words, C is H-almost invariant precisely when it projects to an element in the fixed set $(\mathcal{P}(G)/\mathcal{F}_H(G))^G$. The following proposition gives a geometric interpretation of H-a.i. sets in terms of coarse complementary components of W.

Proposition 3.14. *Let G be a group equipped with the word metric with respect a finite generating set S, and suppose $H \leq G$ is a subgroup. A subset $C \subseteq G$ is H-a.i. if and only if it is a coarse complementary component of H.*

Proof. Say $C \subseteq G$ is H-almost invariant. By Lemma 2.9, we see that for every $g \in G$ there exists an $r_g \geq 0$ such that $Cg \triangle C \subseteq N_{r_g}(H)$. We choose A large enough so that for all $s \in S^{\pm 1}$, $Cs \triangle C \subseteq N_A(H)$. We claim that C is a $(1, A)$-coarse complementary component of H. Let $g \in \partial_1(C \backslash N_A(H))$. If $g \in C$, then as $g \notin C \backslash N_A(H)$, $g \in N_A(H)$ and we are done. If $g \notin C$, then there exists an $h \in C$ such that $g = hs$ for some $s \in S^{\pm 1}$, and so $g \in Cs \triangle C \subseteq N_A(H)$.

Conversely, suppose C is a coarse complementary component of H. Corollary 3.6 ensures C is a $(1, A)$-coarse complementary component of H for some A. For any $g \in G$, we write $g^{-1} = s_1 \dots s_n$, where for each i, either s_i or its inverse is contained in S. If $b \in C \triangle Cg$, then exactly one of b or $bs_1 \dots s_n$ lies in C. Therefore, there exists a $1 \leq t \leq n$ such that $b' := bs_1, \dots s_t \in \partial_1(C)$. By Remark 3.2, $b' \in N_{A+1}(H)$ and so $b \in N_{A+n+1}(H)$. Thus $C \triangle Cg \subseteq N_{A+n+1}(H)$, so $C \triangle Cg$ is H-finite. □

Kropholler and Roller [KR89] define the number of relative ends of a pair of groups to be

$$\tilde{e}(G, H) := \dim_{\mathbb{Z}_2}(\mathcal{P}(G)/\mathcal{F}_H(G))^G.$$

This can be characterised geometrically as follows:

Proposition 3.15 ([Vav12, Lemma 7.5])**.** *Let G be a finitely generated group and let $H \leq G$. Then H coarsely n-separates G if and only if $\tilde{e}(G, H) \geq n$.*

Proof. If H coarsely n-separates G, then there exist n deep, coarse disjoint, coarse complementary components of H, which we label $C_1, \dots, C_n$. By Proposition 3.14, these correspond to elements of $(\mathcal{P}(G)/\mathcal{F}_H(G))^G$. Since they are coarse disjoint, they are linearly independent, so $\tilde{e}(G,H) \geq n$.

Now suppose $\tilde{e}(G,H) \geq n$. By definition, there exist n deep, coarse complementary components $C_1, \dots, C_n$ of H that represent linearly independent elements of $(\mathcal{P}(G)/\mathcal{F}_H(G))^G$. We explain how to modify these coarse complementary components so that they are coarse disjoint.

We set $A := C_1 \backslash C_2$, $B := C_1 \cap C_2$ and $C = C_2 \backslash C_1$, observing that $C_1 = A \triangle B$ and $C_2 = B \triangle C$. By the linear independence of $C_1, \dots, C_n$, at least two of A, B and C, say A and B, aren't contained in the span of $C_3, \dots, C_n$. Then $A, B, C_3, \dots, C_n$ are linearly independent and A and B are coarse disjoint. Applying this procedure to all pairs, we obtain n deep, coarse disjoint, coarse complementary components of H, showing that H coarsely n-separates G. □

There is another notion of relative ends of groups due to Houghton [Hou74] and Scott [Sco78] . They define another notion of relative ends,

$$e(G,H) := \dim_{\mathbb{Z}_2}(\mathcal{P}(H\backslash G)/\mathcal{F}(H\backslash G))^G,$$

where $\mathcal{P}(H\backslash G)$ is the powerset of the coset space $H\backslash G$, and $\mathcal{F}(H\backslash G)$ consists of finite subsets of $H\backslash G$. The following is an analogue of Proposition 3.15 and can be proved in the same way.

Proposition 3.16 ([Sco78, Lemma 1.6])**.** *Let G be finitely generated and $H \leq G$. Then $e(G,H) \geq n$ if and only if there exist n deep, coarse disjoint, coarse complementary components of H, denoted $C_1, \dots, C_n$, such that $HC_i = C_i$ for each C_i.*

It follows from Proposition 3.15 and 3.16 that $e(G,H) \leq \tilde{e}(G,H)$ for all $H \leq G$. We say that $H \leq G$ is a *codimension one subgroup* of G if $e(G,H) > 1$. If G splits over H, then $\tilde{e}(G,H) \geq e(G,H) > 1$.

If we know that G is coarsely separated by a subgroup H, the following lemma allows us to construct a codimension one subgroup $H' \leq H$ of G.

Lemma 3.17. *Let G be a finitely generated group and H be a subgroup. Let C be a deep, irreducible $(1,A)$-coarse complementary component of H and suppose $G\backslash C$ is deep. If $H' := \mathrm{Stab}_H(C\backslash N_A(H))$, then $e(G,H') > 1$.*

Proof. We use an argument from the proof of [Vav12, Lemma 7.6]. We first claim that H' and $\partial_1(C\backslash N_A(H))$ are finite Hausdorff distance from one another. Since H' stabilizes $C\backslash N_A(H)$, it stabilizes $\partial_1(C\backslash N_A(H))$. We let $r_1 := d(e, \partial_1(C\backslash N_A(H)))$. Then for each $h \in H'$, we see that

$$d(h, \partial_1(C\backslash N_A(H))) = d(e, h^{-1}\partial_1(C\backslash N_A(H))) = d(e, \partial_1(C\backslash N_A(H))) = r_1,$$

so $H' \subseteq N_{r_1}(\partial_1(C\backslash N_A(H)))$.

As C is an irreducible $(1, A)$-coarse complementary component of H, $C\backslash N_A(H)$ is the vertex set of a component $\tilde{C}$ of $P_1(G)\backslash P_1(N_A(H))$. As G has bounded geometry, there are only finitely many components $\{h\tilde{C}\}_{h\in H}$ of $P_1(G)\backslash P_1(N_A(H))$ that intersect $N_{A+1}(e)$. We label these components $h_1\tilde{C}, \dots, h_m\tilde{C}$ and set $r_2 := A + \max_{1\leq k\leq m} \mathrm{length}_S(h_k)$.

Suppose $x \in \partial_1(C\backslash N_A(H)) \subseteq N_A(H)$; we choose $h \in H$ such that $d(h, x) \leq A$. There exists a $y \in C\backslash N_A(H)$ such that $d(x, y) \leq 1$. As $d(h, y) \leq A + 1$, we see that $h^{-1}\tilde{C}$ intersects $N_{A+1}(e)$, hence $h^{-1}\tilde{C} = h_r\tilde{C}$ for some r. Thus $\tilde{C} = hh_r\tilde{C}$, so $hh_r \in H'$. Hence $x \in N_{r_2}(H')$ and so $\partial_1(C\backslash N_A(H)) \subseteq N_{r_2}(H')$.

Letting $C' := C\backslash N_A(H)$, Remark 3.2 tells us that C' is a $(1, r_2)$-coarse complementary component of H'. Since C and $G\backslash C$ are deep as coarse complementary components of H, C' and $G\backslash C'$ are also deep as coarse complementary components of H. As $H' \leq H$, C' and $G\backslash C'$ are thus deep as coarse complementary components of H'. Since $H'C' = C'$ and $H'(G\backslash C') = (G\backslash C')$, Proposition 3.16 tells us that $e(G, H') > 1$. □

Corollary 3.18. *Let G, H, C and H' be as above, and suppose that $\tilde{e}(G, K) = 1$ for all infinite index subgroups $K \leq H$. Then $[H : H']$ is finite.*

3.1.4 A quantitative coarse Mayer–Vietoris sequence

Throughout this section, we fix a commutative ring R with unity and assume all cohomology is taken with coefficients in R.

In [HRY93], Higson, Roe and Yu describe a coarse Mayer–Vietoris sequence for coarse cohomology. We suppose C_1 and C_2 are disjoint, coarse complementary components of W such that $G = C_1 \cup C_2$ and $W = C_1 \cap C_2$. Lemma 3.10 tells us that for all $B \geq 0$, there is a $B' \geq 0$ such that $N_B(C_1) \cap N_B(C_2) \subseteq N_{B'}(W)$. Thus C_1 and C_2 are *coarsely excisive* in the sense of [HRY93]. Using Roe's coarse cohomology

as described in Section 2.2.4, it follows from Theorem 1 of [HRY93] that there is an exact sequence

$$0 \to H^0_{\text{coarse}}(X) \xrightarrow{q^*} H^0_{\text{coarse}}(C_1) \oplus H^0_{\text{coarse}}(C_2) \xrightarrow{p^*} H^0_{\text{coarse}}(W) \xrightarrow{\delta} \cdots$$
$$\cdots \xrightarrow{q^*} H^{n-1}_{\text{coarse}}(C_1) \oplus H^{n-1}_{\text{coarse}}(C_2) \xrightarrow{p^*} H^{n-1}_{\text{coarse}}(W) \xrightarrow{\delta} H^n_{\text{coarse}}(X).$$

We re-derive this Mayer–Vietoris sequence using metric complexes. In doing so, we preserve quantitative information that cannot be deduced using [HRY93].

Proposition 3.19. *Let X be a 1-geodesic, bounded geometry metric space with $W \subseteq X$. Suppose both X and W are coarsely uniformly $(n-1)$-acyclic over R. Let C_1 be a $(1, A)$-coarse complementary component of W and let $C_2 := X \backslash C_1$. Then there exist uniformly $(n-1)$-acyclic R-metric complexes $(W, D_\bullet)$, $(C_1 \cup W, A_\bullet)$, $(C_2 \cup W, B_\bullet)$ and $(X, C_\bullet)$, and an exact sequence*

$$0 \to H^0_c(C_\bullet) \xrightarrow{q^*} H^0_c(A_\bullet) \oplus H^0_c(B_\bullet) \xrightarrow{p^*} H^0_c(D_\bullet) \xrightarrow{\delta} H^1_c(C_\bullet) \xrightarrow{q^*} \cdots$$
$$\cdots \xrightarrow{q^*} H^{n-1}_c(A_\bullet) \oplus H^{n-1}_c(B_\bullet) \xrightarrow{p^*} H^{n-1}_c(D_\bullet) \xrightarrow{\delta} \hat{H}^n_c(C_\bullet).$$

Let $p^\#_{C_1} : A^\bullet_c \to D^\bullet_c$, $p^\#_{C_2} : B^\bullet_c \to D^\bullet_c$, $q^\#_{C_1} : C^\bullet_c \to A^\bullet_c$ and $q^\#_{C_2} : C^\bullet_c \to B^\bullet_c$ be maps which have finite displacement over the respective inclusions. The above maps in cohomology are induced by the cochain maps $p^\#(\sigma, \tau) = p^\#_{C_1}\sigma + p^\#_{C_2}\tau$ and $q^\#(\lambda) = (q^\#_{C_1}\lambda, -q^\#_{C_2}\lambda)$. Moreover, for $r < n$:

(1) there is an $R = R(X, W, A) \geq 0$ such that whenever $\sigma \in D^r_c$ is a cocycle supported in $K \subseteq W$, then $\delta[\sigma]$ can be represented by a cocycle in C^{r+1}_c supported in $N^X_R(K)$;

(2) there exists an $s_0 = s_0(X, W, A)$ such that for any $s \geq s_0$ the following holds: a cohomology class $[\omega] \in \hat{H}^{r+1}_c(C_\bullet)$ is contained in the image of the boundary map if and only if $[\omega]$ can be represented by two modified cocycles $\alpha, \beta \in C^{r+1}_c$ supported in $C_1 \backslash N_s(W)$ and $C_2 \backslash N_s(W)$ respectively.

The difficulty in deriving the above exact sequence when compared to the singular or simplicial Mayer–Vietoris sequence, lies in the fact that maps which we want to be injective or surjective may not be. For example, although W is a subspace of $C_1 \cup W$, $D_\bullet$ isn't a subcomplex of $A_\bullet$, so the map $D_\bullet \to A_\bullet$, induced by inclusion, may not be injective.

The proof is similar to that found in [HRY93], except that as spaces are coarsely uniformly $(n-1)$-acyclic, one doesn't need to pass to inverse limits. We make use of the following lemma, which relates coarse cohomology to the cohomology of successive Rips complexes.

Lemma 3.20. *Let X be a 1-geodesic, bounded geometry metric space with $W \subseteq X$. Suppose both X and W are coarsely uniformly $(n-1)$-acyclic over R. For every $A \geq 0$, there exist $i \leq j \leq k$ and $A \leq n_i \leq n_j \leq n_k$ such that the following holds: for every $(1, A)$-coarse complementary component C of W, there exists a uniformly $(n-1)$-acyclic R-metric complex $(C \cup W, C_\bullet)$ and proper chain maps*

$$[C_\bullet]_n \xrightarrow{\alpha_C} [C_\bullet(P_i(C_i))]_n \xrightarrow{\iota_1} [C_\bullet(P_j(C_j))]_n \xrightarrow{\iota_2} [C_\bullet(P_k(C_k))]_n \xrightarrow{\beta_C} [C_\bullet]_n \tag{3.1}$$

where:

(1) $C_i := N^X_{n_i}(W) \cup C$, $C_j := N^X_{n_j}(W) \cup C$ and $C_k := N^X_{n_k}(W) \cup C$;

(2) every $(1, A)$-coarse complementary component of W is also an (i, n_i), (j, n_j) and (k, n_k)-coarse complementary component of W;

(3) the maps ι_1 and ι_2 are inclusions;

(4) the map α_C has displacement at most $r_\alpha = r_\alpha(X, W, A)$ over the inclusion $W \to C_i$;

(5) the map β_C has displacement at most $r_\beta = r_\beta(X, W, A)$ over a coarse inverse to the inclusion $W \to C_k$;

(6) on the level of modified cohomology with compact supports in dimensions at most n, the following identities hold:

$$\alpha_C^* \iota_1^* \iota_2^* \beta_C^* = \mathrm{id} \tag{3.2}$$

$$\iota_1^* \iota_2^* \beta_C^* \alpha_C^* \iota_1^* = \iota_1^* \tag{3.3}$$

$$\iota_2^* \beta_C^* \alpha_C^* \iota_1^* \iota_2^* = \iota_2^*. \tag{3.4}$$

Moreover, (3.1) is natural in the following sense: for every pair of $(1, A)$-coarse complementary components $C \subseteq D$ of W, there are proper chain maps

$$\begin{array}{ccccccccc}
[C_\bullet]_n & \xrightarrow{\alpha_C} & [C_\bullet(P_i(C_i))]_n & \xrightarrow{\iota_1} & [C_\bullet(P_j(C_j))]_n & \xrightarrow{\iota_2} & [C_\bullet(P_k(C_k))]_n & \xrightarrow{\beta_C} & [C_\bullet]_n \\
\downarrow{\scriptstyle f} & & \downarrow{\scriptstyle f_i} & & \downarrow{\scriptstyle f_j} & & \downarrow{\scriptstyle f_k} & & \downarrow{\scriptstyle f} \\
[D_\bullet]_n & \xrightarrow{\alpha_D} & [D_\bullet(P_i(D_i))]_n & \xrightarrow{\iota_1} & [D_\bullet(P_j(D_j))]_n & \xrightarrow{\iota_2} & [D_\bullet(P_k(D_k))]_n & \xrightarrow{\beta_D} & [D_\bullet]_n
\end{array}$$

where:

(1) each row is an instance of (3.1) applied to C and D respectively;

(2) *the maps f_i, f_j and f_k are inclusions of chain complexes;*

(3) *the chain map $f : [C_\bullet]_n \to [D_\bullet]_n$ has finite displacement over the inclusion $C \cup W \to D \cup W$;*

(4) *on the level of modified cohomology with compact supports in dimensions at most n, the following identities hold:*

$$\beta_C^* f^* = f_k^* \beta_D^* \tag{3.5}$$

$$\alpha_C^* f_i^* \iota_1^* = f^* \alpha_D^* \iota_1^*. \tag{3.6}$$

Proof. By Proposition 3.8, there exist λ' and μ', depending only on X, W and A, such that for every $(1, A)$-coarse complementary component C, $C \cup W$ is (λ', μ')-coarsely uniformly $(n-1)$-acyclic. Using Proposition 2.42, we construct a $\tilde{\mu}$-uniformly $(n-1)$-acyclic metric complex $(C \cup W, C_\bullet)$ such that $\tilde{\mu}$ and the n-displacement of $C_\bullet$ depend only on X, W and A.

We proceed by repeated applications of Lemmas 2.39 and 2.40. Using the notation of Lemma 2.39, we set $i = i(\lambda')$, and using Lemma 3.5 we pick n_i such that every $(1, A)$-coarse complementary component of W is also an (i, n_i)-coarse complementary component of W. By Lemma 2.39 and our choice of i, for every $(1, A)$-coarse complementary component C of W, there exists a finite displacement chain map $[C_\bullet]_n \to C_\bullet(P_i(W \cup C))$. We postcompose this map with the inclusion to get the required map α_C.

By Proposition 2.15, C_i is (λ'', μ'')-coarsely uniformly $(n-1)$-acyclic for some λ'' and μ'', depending only on λ, μ and n_i. Using the notation of Lemma 2.39, we set $j = j(i, \lambda'')$. We define suitable n_j, k and n_k similarly. We choose a map $C_k \to W \cup C$ which is a coarse inverse to the inclusion, and by Lemma 2.40, we can define a suitable β_C.

By Lemma 2.21, $\beta_C \iota_2 \iota_1 \alpha_C$ has finite displacement over the identity, so Lemma 2.40 tells us it is properly chain homotopic to the identity chain map. Thus $\alpha_C^* \iota_1^* \iota_2^* \beta_C^* = \mathrm{id}$. Similarly, we see that $\alpha_C \beta_C \iota_2 \iota_1$ has finite displacement over the identity. By our choice of j, we see that $\iota_1 \alpha_C \beta_C \iota_2 \iota_1$ is properly chain homotopic to ι_1, hence $\iota_1^* \iota_2^* \beta_C^* \alpha_C^* \iota_1^* = \iota_1^*$. We deduce (3.4) similarly.

To show naturality, we use Lemma 2.40 to define an $f : [C_\bullet]_n \to [D_\bullet]_n$ that has finite displacement over the inclusion $C \subseteq D$. We then deduce (3.5) and (3.6) using Lemmas 2.39 and 2.40 respectively. □

Remark 3.21. In the proof of Lemma 3.20, one can choose n_k arbitrarily large whilst fixing $(X, C_\bullet)$, i, j, k, n_i, n_j and α_C. However, as we vary n_k we lose the bound r_β on the displacement of β_C.

In the proof of Proposition 3.19, we also use the following easy lemma:

Lemma 3.22. *Let X be a bounded geometry metric space and let $(X, C_\bullet)$ be an R-metric complex. Suppose that for some k and n, $\beta : [C_\bullet(P_k(X))]_n \to C_\bullet$ is a chain map of n-displacement at most r_β. Let $L \subseteq X$ and suppose $q_\# : C_\bullet(P_k(L)) \to C_\bullet(P_k(X))$ is the inclusion of chain complexes. If $\operatorname{supp}(\alpha) \cap N^X_{r_\beta}(L) = \emptyset$ for some $\alpha \in C^n_c$, then $q^\# \beta^\# \alpha = 0$.*

Proof. It is enough to show that $q^\#\beta^\#\alpha$ vanishes on every n-simplex of $P_k(L)$. Let Δ be an n-simplex of $P_k(L)$. Then $\operatorname{supp}(q_\#\Delta) \subseteq L$, so $\operatorname{supp}(\beta_\# q_\# \Delta) \subseteq N^X_{r_\beta}(L)$; therefore $q^\#\beta^\#\alpha(\Delta) = \alpha(\beta_\# q_\# \Delta) = 0$. □

Proof of Proposition 3.19. We note that W and X are themselves $(1, A)$-coarse complementary components of W. We use Lemma 3.20 to produce the proper chain maps in Figure 3.1: the 'p maps' are sums of maps induced by the inclusions $W \to W \cup C_1$ and $W \to W \cup C_2$ and the 'q maps' are differences of the maps induced by the inclusions $W \cup C_1 \to X$ and $W \cup C_2 \to X$. We let $W_i = N^X_{n_i}(W)$, $A_i = C_1 \cup W_i$, $B_i = C_2 \cup W_i$, and similarly for W_j, A_j etc.

$$
\begin{array}{ccccc}
[W_\bullet]_n & \xrightarrow{p} & [A_\bullet]_n \oplus [B_\bullet]_n & \xrightarrow{q} & [C_\bullet]_n \\
\downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \alpha} & & \downarrow{\scriptstyle \alpha} \\
[C_\bullet(P_i(W_i))]_n & \xrightarrow{p} & [C_\bullet(P_i(A_i))]_n \oplus [C_\bullet(P_i(B_i))]_n & \xrightarrow{q} & [C_\bullet(P_i(X))]_n \\
\downarrow{\scriptstyle \iota_1} & & \downarrow{\scriptstyle \iota_1} & & \downarrow{\scriptstyle \iota_1} \\
[C_\bullet(P_j(W_j))]_n & \xrightarrow{p} & [C_\bullet(P_j(A_j))]_n \oplus [C_\bullet(P_j(B_j))]_n & \xrightarrow{q} & [C_\bullet(P_j(X))]_n \\
\downarrow{\scriptstyle \iota_2} & & \downarrow{\scriptstyle \iota_2} & & \downarrow{\scriptstyle \iota_2} \\
[C_\bullet(P_k(W_k))]_n & \xrightarrow{p} & [C_\bullet(P_k(A_k))]_n \oplus [C_\bullet(P_k(B_k))]_n & \xrightarrow{q} & [C_\bullet(P_k(X))]_n \\
\downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \beta} & & \downarrow{\scriptstyle \beta} \\
[W_\bullet]_n & \xrightarrow{p} & [A_\bullet]_n \oplus [B_\bullet]_n & \xrightarrow{q} & [C_\bullet]_n
\end{array}
$$

Figure 3.1: The maps used to construct the coarse Mayer–Vietoris sequence

By Lemma 3.7, each simplex of $P_i(X)$ is contained in either $P_i(A_i)$ or $P_i(B_i)$. This means that the second row of Figure 3.1 is a short exact sequence, so the Mayer–Vietoris sequence in simplicial cohomology with compact supports gives us:

$$\cdots \xrightarrow{p^*} H_c^*(P_i(N_{n_i}(W))) \xrightarrow{\tilde{\delta}} H_c^{*+1}(P_i(X)) \xrightarrow{q^*} \cdots . \tag{3.7}$$

We obtain similar Mayer–Vietoris sequences in the third and forth rows, and by naturality of the simplicial Mayer–Vietoris sequence, the boundary map $\tilde{\delta}$ commutes with the maps ι_2^* and ι_1^*.

For $k \leq n-1$, we define the required boundary map $\delta := \alpha^* \iota_1^* \iota_2^* \tilde{\delta} \beta^* = \alpha^* \iota_1^* \tilde{\delta} \iota_2^* \beta^*$. As $H^{n-1}(W_\bullet) = 0$, we see that on the level of ordinary cohomology, i.e. cohomology possibly without compact supports, the boundary map δ is zero. Therefore, δ defines a map $\delta : H_c^{n-1}(W_\bullet) \to \hat{H}_c^n(C_\bullet)$. In lower dimensions, Lemma 2.45 tells us that modified cohomology with compact supports and cohomology with compact supports are equal.

It is easy to verify exactness using the identities from Lemma 3.20. We shall show exactness at $H_c^r(W_\bullet)$ for $r < n$. Let $x \in H_c^r(W_\bullet)$. Using (3.5), we see that β^* and p^* commute. If $x = p^*(x')$, then as (3.7) is exact, we see that

$$\delta(x) = \alpha^* \iota_1^* \iota_2^* \tilde{\delta} \beta^* p^*(x') = \alpha^* \iota_1^* \iota_2^* \tilde{\delta} p^* \beta^*(x') = 0.$$

We now suppose $\delta(x) = 0$. Using (3.4), we see

$$\tilde{\delta} \iota_2^* \beta^*(x) = \iota_2^* \tilde{\delta} \beta^*(x) = \iota_2^* \beta^* \alpha^* \iota_1^* \iota_2^* \tilde{\delta} \beta^*(x) = \iota_2^* \beta^* \delta(x) = 0.$$

By the exactness of (3.7), $p^*(y) = \iota_2^* \beta^*(x)$ for some $y \in H_c^r(P_j(A_j)) \oplus H_c^r(P_j(B_j))$. Let $z = \alpha^* \iota_1^*(y) \in H_c^r(A_\bullet) \oplus H_c^r(B_\bullet)$. Then

$$p^*(z) = p^* \alpha^* \iota_1^*(y) = \alpha^* p^* \iota_1^*(y) = \alpha^* \iota_1^* p^*(y) = \alpha^* \iota_1^* \iota_2^* \beta^*(x) = x,$$

using (3.6) and the fact that p^* and ι_1^* commute. Thus $\ker(\delta) = \mathrm{im}(p^*)$, so we have exactness at $H_c^r(W_\bullet)$.

(1): Let $r < n$. Given an r-cycle $\sigma \in Z_c^r(P_k(W_k))$, we extend σ to $\rho_\sigma \in C_c^r(P_k(A_k))$ by setting ρ_σ to be zero outside of $P_k(W_k)$. We extend $\delta\rho_\sigma \in C_c^{r+1}(P_k(A_k))$ to $\omega_\sigma \in C_c^{r+1}(P_k(X))$ by setting ω_σ to be zero outside $P_k(A_k)$. Since $p^\#(\rho_\sigma, 0) = \sigma$ and $q^\#(\omega_\sigma) = (\delta\rho_\sigma, 0)$, we see that $\tilde{\delta}[\sigma] = [\omega_\sigma]$. It follows that if $\gamma \in D_c^k$ is a cocycle supported in $K \subseteq W$, then $\alpha^\# \iota_1^\# \iota_2^\# \omega_{\beta\#\gamma}$ is a cocycle representing $\delta[\gamma]$ and supported in $N_R^X(K)$, where R depends only on k, r_α and r_β.

(2): We first fix a cocycle $\gamma \in D_c^r$ and some $s \geq 0$. We will show that $\delta[\gamma]$ can be represented by two cocycles supported in $C_1 \backslash N_s(W)$ and $C_2 \backslash N_s(W)$ respectively. Defining $\omega_{\beta\#\gamma}$ as above, we see that $q_{C_2}^\# \omega_{\beta\#\gamma} = 0$, where $q_{C_2}^\# : C_c^\bullet(P_k(X)) \to$

$C_c^\bullet(P_k(N_{n_k}(W) \cup C_2))$ is the restriction map. Thus if $\omega_{\beta\#\gamma}(\Delta) \neq 0$ for some $(r+1)$-simplex $\Delta \in C_{r+1}(P_k(X))$, then Δ intersects $C_1 \backslash N_{n_k}(W)$. Hence by Lemma 3.10, $\omega_{\beta\#\gamma}$ is supported in $N_k(C_1 \backslash N_{n_k}(W)) \subseteq (C_1 \cup N_{k+A}(W)) \backslash N_{n_k - k}(W)$.

As in Remark 3.21, we can choose n_k arbitrarily large while keeping k and α fixed. Since $\mathrm{supp}(\omega_{\beta\#\gamma}) \subseteq (C_1 \cup N_{k+A}(W)) \backslash N_{n_k-k}(W)$ and $\alpha^\# \iota_1^\# \iota_2^\# \omega_{\beta\#\gamma}$ is a cocycle representing $\delta[\gamma]$, we can choose n_k large enough so that $\alpha^\# \iota_1^\# \iota_2^\# \omega_{\beta\#\gamma}$ is supported in $C_1 \backslash N_s(W)$. Reversing the roles of C_1 and C_2, we see that $\delta[\gamma]$ can be represented by a cocycle supported in $C_2 \backslash N_s(W)$.

For the converse, we keep n_k fixed and choose $s_0 = r_\beta + n_k$. By Lemma 3.10, we see $N_{r_\beta}(C_i \cup N_{n_k}(W)) \subseteq C_i \cup N_{s_0}(W)$ for $i = 1, 2$. Let $s \geq s_0$. By Lemma 3.22 and our choice of s, we see that for any $\mu, \nu \in C_c^{n+1}$ supported in $C_1 \backslash N_s(W)$ and $C_2 \backslash N_s(W)$ respectively, then

$$q_{C_1}^\# \beta^\# \nu = q_{C_2}^\# \beta^\# \mu = p_{C_1}^\# q_{C_1}^\# \beta^\# \mu = 0.$$

Suppose there exist modified cocycles $\mu, \nu \in C_c^{n+1}$, supported in $C_1 \backslash N_s(W)$ and $C_2 \backslash N_s(W)$ respectively, each representing $[\omega] \in \hat{H}_c^{n+1}[C_\bullet]$. Then there is a $\tau \in C_c^n$ such that $\delta\tau = \mu - \nu$. We see that $p^\#(q_{C_1}^\# \beta^\# \tau, 0) = p_{C_1}^\# q_{C_1}^\# \beta^\# \tau$ and

$$\delta(q_{C_1}^\# \beta^\# \tau, 0) = (q_{C_1}^\# \beta^\# (\mu - \nu), 0) = (q_{C_1}^\# \beta^\# \mu, -q_{C_2}^\# \beta^\# \mu) = q^\# \beta^\# \mu.$$

Since $\delta p_{C_1}^\# q_{C_1}^\# \beta^\# \tau = p_{C_1}^\# q_{C_1}^\# \beta^\# (\mu - \nu) = 0$, $p_{C_1}^\# q_{C_1}^\# \beta^\# \tau$ is a cocycle with $\tilde{\delta}[p_{C_1}^\# q_{C_1}^\# \beta^\# \tau] = [\beta^\# \mu] = \beta^*[\omega]$. We therefore see that

$$\begin{aligned} \delta(\alpha^* \iota_1^* \iota_2^* [p_{C_1}^\# q_{C_1}^\# \beta^\# \tau]) &= \alpha^* \iota_1^* \tilde{\delta} \iota_2^* \beta^* \alpha^* \iota_1^* \iota_2^* [p_{C_1}^\# q_{C_1}^\# \beta^\# \tau] \\ &= \alpha^* \iota_1^* \iota_2^* \tilde{\delta} [p_{C_1}^\# q_{C_1}^\# \beta^\# \tau] = \alpha^* \iota_1^* \iota_2^* \beta^* [\omega] = [\omega]. \end{aligned}$$

□

3.2 Constructing splittings

3.2.1 Essential components

Throughout this section, we assume that X is a 1-geodesic, bounded geometry metric space that is coarsely uniformly n-acyclic over $\mathbb{Z}_2$, and that $W \subseteq X$ is a coarse $PD_n^{\mathbb{Z}_2}$ space. We assume all metric complexes are $\mathbb{Z}_2$-metric complexes and that cohomology is taken with coefficients in $\mathbb{Z}_2$.

The following definition makes sense in light of Proposition 3.8 which shows that if C is a coarse complementary component of W, then $C \cup W$ is coarsely uniformly n-acyclic; hence we can define the coarse cohomology $H^k_{\mathrm{coarse}}(C \cup W)$ for $k \leq n+1$.

$$\begin{array}{ccc} [D_\bullet]_{n+1} & \xrightarrow{a_\#} & [A_\bullet]_{n+1} \\ {\scriptstyle p_\#}\downarrow & & {\scriptstyle q_\#}\downarrow \\ [D'_\bullet]_{n+1} & \xrightarrow{a'_\#} & [A'_\bullet]_{n+1} \end{array}$$

Figure 3.2: A commutative diagram induced by f

Definition 3.23. A coarse complementary component $C \subseteq X$ of W is said to be *essential* if the map $H^n_{\text{coarse}}(C \cup W) \to H^n_{\text{coarse}}(W)$, induced by the inclusion $W \to C \cup W$, is zero.

The following proposition shows that essential coarse complementary components are preserved under coarse isometries.

Proposition 3.24. *Suppose $C \subseteq X$ is an essential coarse complementary component of W and $f : X \to Y$ is a coarse isometry with $N^Y_B(f(X)) = Y$. Then $N^Y_B(f(C))$ is an essential coarse complementary component of $f(W)$.*

Proof. It follows from Lemma 3.5 that $N^Y_B(f(C))$ is a coarse complementary component of $f(W)$, so we need only show that it is essential. Let $(W, D_\bullet)$, $(f(W), D'_\bullet)$, $(C \cup W, A_\bullet)$ and $(N^Y_B(f(C)) \cup f(W), A'_\bullet)$ be uniformly n-acyclic metric complexes. Using Lemma 2.40, we define proper chain maps shown in Figure 3.2. We require $a_\#$ and $a'_\#$ have finite displacement over the inclusions $W \to C \cup W$ and $f(W) \to N^Y_B(f(C)) \cup f(W)$ respectively, and $p_\#$ and $q_\#$ have finite displacement over $f|_W$ and $f|_{C \cup W}$ respectively.

By Lemma 2.40, we see that $q_\# a_\#$ is properly chain homotopic to $a'_\# p_\#$. Since C is essential, the map $a^* : H^n_c(A_\bullet) \to H^n_c(D_\bullet)$ is zero, therefore $a^* q^* = p^*(a')^* : H^n_c(A'_\bullet) \to H^n_c(D_\bullet)$ is also zero. As $f|_W : W \to f(W)$ is a coarse isometry, p^* is an isomorphism and so $(a')^*$ is zero; hence $N_B(f(C))$ is essential. □

The following proposition is useful for showing that coarse complementary components are or are not essential. We note that in the following Proposition, C need not not be a coarse complementary component.

Proposition 3.25. *Suppose C is a subspace of X such that $C \cup W$ is coarsely uniformly n-acyclic and $H^n_{\text{coarse}}(C \cup W) \to H^n_{\text{coarse}}(W)$, induced by inclusion, is zero. Then:*

(1) for every $R \geq 0$, C is not contained in $N^X_R(W)$;

(2) if D is a coarse complementary component of W and $C \subseteq D$, then D is essential;

(3) if there exist coarse complementary components C_1 and C_2 of W such that $C_1 \cup C_2 = C$, then at least one of C_1 or C_2 is essential.

Proof. (1): If $C \subseteq N_R^X(W)$ for some $R \geq 0$, then the inclusion $W \to C \cup W$ is a coarse isometry, so induces an isomorphism $H^n_{\text{coarse}}(C \cup W) \to H^n_{\text{coarse}}(W) \cong \mathbb{Z}_2$.

(2): Let $f : W \to C \cup W$ and $g : C \cup W \to D \cup W$ be inclusions. Then the composition $H^n_{\text{coarse}}(D \cup W) \xrightarrow{g^*} H^n_{\text{coarse}}(C \cup W) \xrightarrow{f^*} H^n_{\text{coarse}}(W)$ is zero, since f^* is zero. Therefore $(gf)^* = f^*g^* = 0$, so D is essential.

(3): For $i = 1, 2$, let $a_i : W \to W \cup C_i$ and $b_i : W \cup C_i \to W \cup C$ be inclusions. The relevant terms in the coarse Mayer–Vietoris sequence from Proposition 3.19 are

$$\cdots \to H^n_{\text{coarse}}(C \cup W) \xrightarrow{b^* = b_1^* - b_2^*} H^n_{\text{coarse}}(C_1 \cup W) \oplus H^n_{\text{coarse}}(C_2 \cup W) \xrightarrow{a^* = a_1^* + a_2^*} H^n_{\text{coarse}}(W) \to \cdots .$$

If both a_1^* and a_2^* are non-zero, then as $H^n_{\text{coarse}}(W) \cong \mathbb{Z}_2$, there exist x, y such that $a_1^*(x) = a_2^*(y) = 1$. As $a^*(x, -y) = a_1^*(x) - a_2^*(y) = 0$, there exists a $z \in H^n_{\text{coarse}}(C \cup W)$ such that $b^*(z) = (b_1^*(z), -b_2^*(z)) = (x, -y)$. Therefore $(b_1 a_1)^*(z) = a_1^* b_1^*(z) = a_1^*(x) = 1$, contradicting our hypothesis on C. □

Remark 3.26. If $C \cup W$ is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$, then Theorem 2.50 tells us that $H^n_{\text{coarse}}(C \cup W) = 0$, and so the conclusions of Proposition 3.25 hold.

Definition 3.27. Let C be a bounded geometry metric space and $W \subseteq C$ be a coarse $PD_n^{\mathbb{Z}_2}$ space. We say that (C, W) is a *coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-space* if there exists a coarse $PD_{n+1}^{\mathbb{Z}_2}$ space Y with $C \subseteq Y$, such that C and $Y \backslash C$ are both deep coarse complementary components of W in Y.

The coarse Jordan separation theorem [KK05, Corollary 7.8] states that if $f : W \to Y$ is a coarse embedding of a coarse $PD_n^{\mathbb{Z}_2}$ space into a coarse $PD_{n+1}^{\mathbb{Z}_2}$ space, then Y contains two disjoint coarse complementary components of $f(W)$ that are coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-spaces.

Proposition 3.28. *Let $W \subseteq C \subseteq X$ and suppose (C, W) is a coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-space. Then the map $H^n_{\text{coarse}}(C \cup W) \to H^n_{\text{coarse}}(W)$, induced by inclusion, is zero.*

Proof. Since (C, W) is a coarse $PD_{n+1}^{\mathbb{Z}_2}$ half-space, there exists a coarse $PD_{n+1}^{\mathbb{Z}_2}$ space Y with $C \subseteq Y$, such that C and $C' := Y \backslash C$ are deep, coarse disjoint, coarse complementary components of W. Proposition 3.19 then tells us that there is an exact sequence

$$\cdots \to H^n_{\text{coarse}}(C \cup W) \oplus H^n_{\text{coarse}}(C' \cup W) \to H^n_{\text{coarse}}(W) \to H^{n+1}_{\text{coarse}}(Y) \to \cdots .$$

We pick a metric complex $(Y, D_\bullet)$ such that Proposition 3.19 holds; by Proposition 2.53, $(Y, D_\bullet)$ is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ complex. Using the duality map, we observe that there is a number R such that for each $y \in Y$, the non-trivial class $[\sigma]$ of $H^{n+1}_c(D_\bullet) \cong \mathbb{Z}_2$ can be represented by a cocycle supported in $N^Y_R(y)$. Thus we can represent $[\sigma]$ by cocycles supported in $C \backslash N^Y_s(W)$ and $C' \backslash N^Y_s(W)$ for any $s \geq 0$. Then Part (2) of Proposition 3.19 tells us that $H^n_{\text{coarse}}(W) \to H^{n+1}_{\text{coarse}}(Y)$ is non-zero, hence an isomorphism. Therefore the map $H^n_{\text{coarse}}(C \cup W) \to H^n_{\text{coarse}}(W)$, induced by inclusion, is zero. □

Corollary 3.29. *Suppose $D \subseteq X$ is a coarse complementary component of W and there exists a coarse $PD^{\mathbb{Z}_2}_{n+1}$ half-space (C, W) such that $C \subseteq D$. Then D is essential.*

Proof. We apply Proposition 3.28, followed by Part (2) of Proposition 3.25 . □

The following proposition shows that an essential coarse complementary component of W is uniformly close to every point of W.

Proposition 3.30. *For every $A \geq 0$, there is a number $B = B(X, W, A)$ such that whenever $C \subseteq X$ is an essential $(1, A)$-coarse complementary component of W, then $W \subseteq N_B(C \backslash N_A(W))$.*

Proof. By Lemma 3.20, there exists a uniformly acyclic metric complex $(W, D_\bullet)$, and numbers $j \leq k$ and $A \leq n_j \leq n_k$, depending only on X, W and A, such that the following holds: there exists a proper chain map $\beta : [C_\bullet(P_k(N_{n_k}(W)))]_{n+1} \to D_\bullet$ such that if $\iota^*_2 : H^n_c(P_k(N_{n_k}(W))) \to H^n_c(P_j(N_{n_j}(W)))$ is the restriction map, then $\iota^*_2 \beta^*$ is injective. As C is essential, Lemma 3.20 also tells us that $\iota^*_2 p^*_C$ is zero, where $p^\#_C : C^\bullet_c(P_k(N_{n_k}(W) \cup C)) \to C^\bullet_c(P_k(N_{n_k}(W)))$ is the restriction map.

By Proposition 2.53, $(W, D_\bullet)$ is a coarse $PD^{\mathbb{Z}_2}_n$ complex, hence there is a duality map $\overline{P} : D_\bullet \to D^{n-\bullet}_c$. For each $x \in W$, we pick $\sigma_x \in p_0^{-1}(x) \subseteq \Sigma_0$. Each cocycle $\overline{P}(\sigma_x) \in D^n_c$ represents the non-trivial element of $H^n_c(D_\bullet) \cong \mathbb{Z}_2$.

Since β and P have finite displacement, there exists a $D \geq 0$ depending only on X, W and A such that for every $x \in W$, $\beta^\# \overline{P} \sigma_x$ is supported in $N_D(x)$. We extend $\beta^\# \overline{P} \sigma_x$ to a cochain $\omega_x \in C^n_c(P_k(N_{n_k}(W) \cup C))$ by setting ω_x to be zero outside $P_k(N_{n_k}(W))$. We note that $p^\#_C \omega_x = \beta^\# \overline{P} \sigma_x$.

We suppose for contradiction that ω_x is a cocycle. Then $\iota^*_2 \beta^* [\overline{P} \sigma_x] = \iota^*_2 p^*_C [\omega_x] = 0$. This contradicts the fact that $\iota^*_2 \beta^*$ is injective. Thus ω_x is not a cocycle. It follows that there is some $(n+1)$-simplex Δ of $P_k(N_{n_k}(W) \cup C)$, not contained in $P_k(N_{n_k}(W))$, such that $(\delta \omega_x)(\Delta) = \omega_x(\partial \Delta) \neq 0$. Thus there exists some $y_x \in C \backslash N_{n_k}(W) \subseteq C \backslash N_A(W)$ with $d(x, y_x) \leq k + D$. □

A coarse complementary component C is said to be *almost essential* if it satisfies the conclusion of Proposition 3.30, i.e. for every $A \geq 0$ there is a number $B \geq 0$ such that $W \subseteq N_B(C \backslash N_A(W))$. The top component of Figure 1.2 in the introduction gives an almost essential component that is not essential.

Definition 3.31. We say that $W \subseteq X$ is *essentially embedded* if every deep coarse complementary component of W is essential. We say that $W \subseteq G$ is *almost essentially embedded* if for every $A \geq 0$, there exists a $B = B(X, W, A) \geq 0$ such that $W \subseteq N_B(C \backslash N_A(W))$ for every deep $(1, A)$-coarse complementary component C of W.

The almost essentially embedded condition is a reformulation of the *deep condition* of [Vav12]. Using Proposition 3.15, we see a subgroup K coarsely separates G if and only if $\tilde{e}(G, K) > 1$. We can therefore deduce the following from Lemmas 7.6 and 8.12 of [Vav12].

Proposition 3.32 ([Vav12, Lemmas 7.6 and 8.12])**.** *Let G be a finitely generated group and $H \leq G$ be a subgroup. Then H is almost essentially embedded if and only if no infinite index subgroup K of H coarsely separates G.*

It is also shown in [Vav12] that being almost essentially embedded is a quasi-isometry invariant in the following sense:

Proposition 3.33 ([Vav12, Lemma 2.10])**.** *Let G and G' be finitely generated groups with subgroups $H \leq G$ and $H' \leq G$. Suppose there is a quasi-isometry $f : G \to G'$ such that $d_{\text{Haus}}(f(H), H') < \infty$. Then H is almost essentially embedded if and only if H' is almost essentially embedded.*

We are now ready to prove Proposition 3.34, which generalises Proposition 3.32 to give an algebraic characterisation of essentially embedded subgroups.

Proposition 3.34. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$ and let $H \leq G$ be a coarse $PD_n^{\mathbb{Z}_2}$ group. Then H is essentially embedded if and only if no infinite index subgroup of H coarsely separates G.*

Proof. If H is essentially embedded, then Proposition 3.30 says it is almost essentially embedded. Thus Proposition 3.32 tells us that no infinite index subgroup K of H coarsely separates G.

Now suppose that no infinite index subgroup K of H coarsely separates G. Let C be a deep coarse complementary component of H. We will show that C is essential. We explain how to reduce to the case where $HC = C$. As C is deep, Lemma 3.12 tells

us that for some $A \geq 0$, C contains a deep, irreducible $(1, A)$-coarse complementary component $C' \subseteq C$. By Corollary 3.18, there is a finite index subgroup $H' \leq H$ such that $H'(C' \backslash N_A(H)) = C' \backslash N_A(H)$. Proposition 3.25 then tells that if $C' \backslash N_A(H)$ is essential as a coarse complementary component of H', then C is essential as a coarse complementary component of H. Thus we may assume that $HC = C$.

Since G is $(n-1)$-acyclic at infinity, Proposition 2.50 says that $H^k_{\text{coarse}}(G) = 0$ for $k \leq n$. We now apply the coarse Mayer–Vietoris sequence to the coarse complementary components C and $G \backslash C$. This gives the exact sequence

$$0 = H^k_{\text{coarse}}(G) \to H^k_{\text{coarse}}(H \cup C) \oplus H^k_{\text{coarse}}(H \cup (G \backslash C)) \to H^k_{\text{coarse}}(H)$$

for each $k \leq n$. If C is not essential, then the maps $H^k_{\text{coarse}}(H \cup C) \to H^k_{\text{coarse}}(H)$, induced by inclusion, are isomorphisms for each $k \leq n$. Thus by Lemma 2.61, C must be shallow; this contradicts our choice of C, so C is essential. □

Lemma 3.35 (Non-Crossing Lemma)**.** *Let C_1, C_2 and C_3 be deep, disjoint, $(1, A)$-coarse complementary components of W in X, and suppose C_3 is essential. Then there exists an $s_0 = s_0(X, W, A)$ and a uniformly n-acyclic metric complex $(X, C_\bullet)$ such that whenever $[\omega] \in \hat{H}^{n+1}_c(C_\bullet)$ is represented by two cocycles supported in $C_1 \backslash N_{s_0}(W)$ and $C_2 \backslash N_{s_0}(W)$ respectively, then $[\omega] = 0$.*

Proof. Without loss of generality, we can replace C_3 with $X \backslash (C_1 \cup C_2)$, which is essential by Proposition 3.25. Proposition 3.19 says that there exist uniformly n-acyclic metric complexes $(C_1 \cup W, A_\bullet)$, $(C_2 \cup C_3 \cup W, B_\bullet)$ and $(W, D_\bullet)$ and a Mayer–Vietoris sequence

$$\cdots \xrightarrow{q^*} H^n_c(A_\bullet) \oplus H^n_c(B_\bullet) \xrightarrow{p^*} H^n_c(D_\bullet) \xrightarrow{\delta} \hat{H}^{n+1}_c(C_\bullet).$$

We also choose i, j, k, n_i, n_j, n_k and $\beta : C_\bullet[P_k(X)]_n \to C_\bullet$ as in Lemma 3.20.

The proof proceeds in a similar fashion to the proof of (2) in Proposition 3.19, except we have three components rather than two. As in Proposition 3.19, we set $s_0 := r_\beta + n_k$, which depends only on X, W and A. We consider the restriction maps

$$q^\#_{C_i} : C^\bullet_c(P_k(X)) \to C^\bullet_c(P_k(N_{n_k}(W) \cup C_i)) \text{ for } i = 1, 2, 3$$
$$p^\#_{C_i} : C^\bullet_c(P_k(N_{n_k}(W) \cup C_i)) \to C^\bullet_c(P_k(N_{n_k}(W))) \text{ for } i = 1, 2, 3,$$

noting that $p^\#_{C_1} q^\#_{C_1} = p^\#_{C_3} q^\#_{C_3}$ is the restriction map $C^\bullet_c(P_k(X)) \to C^\bullet_c(P_k(N_{n_k}(W)))$.

Suppose there exist cocycles $\mu, \nu \in C^{n+1}_c$, supported in $C_1 \backslash N_{s_0}(W)$ and $C_2 \backslash N_{s_0}(W)$ respectively, each representing $[\omega]$. Then there is a $\tau \in C^n_c$ such that $\delta\tau = \mu - \nu$. Let $\tilde{\delta}$ be the boundary map associated to the Mayer–Vietoris sequence of Rips complexes.

As in the proof of Proposition 3.19, we see that $\tilde{\delta}[p_{C_1}^{\#} q_{C_1}^{\#} \beta^{\#} \tau] = [\beta^{\#} \mu] = \beta^*[\omega]$. By Lemma 3.22 and our choice of s_0, we see that $q_{C_3}^{\#} \beta^{\#} \delta \tau = q_{C_3}^{\#} \beta^{\#}(\mu - \nu) = 0$. Therefore $q_{C_3}^{\#} \beta^{\#} \tau$ is a cocycle and $\tilde{\delta} p_{C_3}^*[q_{C_3}^{\#} \beta^{\#} \tau] = \tilde{\delta}[p_{C_1}^{\#} q_{C_1}^{\#} \beta^{\#} \tau] = \beta^*[\omega]$.

As C_3 is essential, Lemma 3.20 tells us that the map $\iota_2^* p_{C_3}^* : H_c^n(P_k(N_{n_k}(W) \cup C_3)) \to H_c^n(P_j(N_{n_j}(W)))$ is zero. As noted in the proof of Proposition 3.19, the boundary map $\tilde{\delta}$ commutes with the restriction map ι_2^*, so $[\omega] = \alpha^* \iota_1^* \iota_2^* \beta^*[\omega] = \alpha^* \iota_1^* \tilde{\delta} \iota_2^* p_{C_3}^*[q_{C_3}^{\#} \beta^{\#} \tau] = 0$. □

3.2.2 Kleiner's mobility sets

We now outline a method from an unpublished manuscript of Kleiner. For each coarse cohomology class $[\sigma] \in H^k_{\text{coarse}}(G)$, Kleiner produces a subset $\text{Mob}([\sigma]) \subseteq G$ which is the support of all possible cocycles of uniformly bounded diameter representing $[\sigma]$. All the results in this subsection are contained in [Kle02]. We emphasise that throughout this section, all cohomology is taken with coefficients in $\mathbb{Z}_2$.

Suppose G is a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and $(G, C_\bullet)$ is a uniformly n-acyclic $\mathbb{Z}_2$-metric complex admitting a free G-action. We fix $k \leq n+1$ and let $Z_c^k \leq C_c^k$ be the set of modified k-cocycles with compact support. For $D > 0$, we define

$$Z_c^k(D) := \{\alpha \in Z_c^k \mid \text{diam}(\alpha) \leq D\}.$$

If a group G acts cocompactly on $(X, C_\bullet)$, then G acts on $Z_c^k(D)$ via the right action $(\alpha g)(\sigma) := \alpha(g\sigma)$. Note that $\text{supp}(\alpha) = g\text{supp}(\alpha g)$ for all $g \in G$. We define $\text{Stab}([\alpha_0]) \leq G$ to be the subgroup of G which preserves the cohomology class $[\alpha_0]$. For each $\alpha_0 \in Z_c^k$, let $Z([\alpha_0], D) := \{\alpha \in Z_c^k(D) \mid [\alpha] = [\alpha_0]\}$. We now define the *mobility set* to be

$$\text{Mob}([\alpha_0], D) = \bigcup_{\alpha \in Z([\alpha_0], D)} \text{supp}(\alpha).$$

Lemma 3.36. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and suppose $(G, C_\bullet)$ is a uniformly n-acyclic $\mathbb{Z}_2$-metric complex admitting a free G-action. For any $D \geq 0$ and $k \leq n+1$, suppose $\alpha_0 \in Z_c^k(D)$ is a nonzero cocycle. Then there exists an $R \geq 0$ such that*

$$d_{\text{Haus}}(\text{Stab}([\alpha_0])\text{supp}(\alpha_0), \text{Mob}([\alpha_0], D)) \leq R.$$

Proof. We fix some $\alpha_0 \in Z_c^k(D)$ and consider the set

$$Q := \{\alpha g \mid \alpha \in Z([\alpha_0], D), g \in G \text{ and } \text{supp}(\alpha g) \subseteq N_D(e)\}.$$

As there are only finitely many cocycles supported in $N_D(e)$, the set Q is finite. Therefore we can write $Q = \{\alpha_1 g_1, \ldots, \alpha_n g_n\}$, where $\alpha_i \in Z([\alpha_0], D)$ and $g_i \in G$ for each i.

We now choose R_{α_0} sufficiently large such that $\mathrm{supp}(\alpha_i) \subseteq N_{R_{\alpha_0}}(\mathrm{supp}(\alpha_0))$ for each α_i. We claim that $d_{\mathrm{Haus}}(\mathrm{Stab}([\alpha_0])\mathrm{supp}(\alpha_0), \mathrm{Mob}([\alpha_0], D)) \leq R_{\alpha_0}$. Indeed, we see $g\mathrm{supp}(\alpha_0) = \mathrm{supp}(\alpha_0 g^{-1}) \subseteq \mathrm{Mob}([\alpha_0], D)$ for each $g \in \mathrm{Stab}([\alpha_0])$. Conversely, suppose $g \in \mathrm{Mob}([\alpha_0], D)$. Then there exists an $\alpha \in Z([\alpha_0], D)$ such that $g \in \mathrm{supp}(\alpha)$. Since $\mathrm{supp}(\alpha g) = g^{-1}\mathrm{supp}(\alpha)$ and $\alpha g \in Q$, $\alpha g = \alpha_i g_i$ for some i. Hence $g \in \mathrm{supp}(\alpha_i g_i g^{-1}) \subseteq g g_i^{-1} N_{R_{\alpha_0}}(\mathrm{supp}(\alpha_0))$. This proves the claim, since $g g_i^{-1} \in \mathrm{Stab}([\alpha_0])$.

We let $R := \max\{R_\alpha \mid \alpha \in Z_c^{n+1}(D) \text{ and } \mathrm{supp}(\alpha) \subseteq N_D(e)\}$. Suppose $\alpha \in Z_c^{n+1}(D)$ and $g \in \mathrm{supp}(\alpha)$. As $\mathrm{supp}(\alpha g) \subseteq N_D(e)$,

$$d_{\mathrm{Haus}}(\mathrm{Stab}([\alpha g])\mathrm{supp}(\alpha g), \mathrm{Mob}([\alpha g], D)) \leq R.$$

Since

$$\mathrm{Stab}([\alpha g])\mathrm{supp}(\alpha g) = g^{-1}\mathrm{Stab}([\alpha])\mathrm{supp}(\alpha)$$

and $\mathrm{Mob}([\alpha g], D) = g^{-1}\mathrm{Mob}([\alpha], D)$, it follows that

$$d_{\mathrm{Haus}}(\mathrm{Stab}([\alpha])\mathrm{supp}(\alpha), \mathrm{Mob}([\alpha], D)) \leq R.$$

□

The following corollary shows that the mobility set of a coarse cohomology class has finite Hausdorff distance from the subgroup which stabilizes that class. This provides a connection between geometry and algebra that allows us to construct a splitting.

Corollary 3.37. *For any $D > 0$ such that $Mob([\alpha_0], D) \neq \emptyset$, $Stab([\alpha_0])$ has finite Hausdorff distance from $Mob([\alpha_0], D)$.*

Proof. By Lemma 3.36, we see that

$$d_{\mathrm{Haus}}(\mathrm{Mob}([\alpha_0], D), \mathrm{Stab}([\alpha_0])) \leq R + d(e, \mathrm{supp}(\alpha_0)).$$

□

Proposition 3.38. *Suppose $(G, C_\bullet)$ and $(G', D_\bullet)$ are uniformly $(n-1)$-acyclic R-metric complexes. Say $f : G \to G'$ is a coarse isometry and $0 \neq [\alpha_0] \in \hat{H}_c^n(D_\bullet)$. Then if for any D, D' sufficiently large,*

$$d_{\mathrm{Haus}}(f(\mathrm{Mob}(f^*[\alpha_0], D)), \mathrm{Mob}([\alpha_0], D')) < \infty.$$

Proof. Suppose $C_\bullet$ and $D_\bullet$ have n-displacement at most d. Lemma 2.40 tells us that f induces a chain map $f_\# : [C_\bullet]_n \to [D_\bullet]_n$ of n-displacement at most M over f. We assume D' is large enough so that $\mathrm{Mob}([\alpha_0], D')$ is nonempty. Let $\alpha \in Z([\alpha_0], D')$. The proof of Proposition 2.34 then shows that $f(\mathrm{supp}(f^\#\alpha)) \subseteq N_M(\mathrm{supp}(\alpha))$. In particular, as f is a coarse isometry, there is a D sufficiently large such that $\mathrm{diam}(f^\#\alpha) \leq D$ for all $\alpha \in Z([\alpha_0], D')$.

Since as f is a coarse isometry, $f^\#$ induces an isomorphism in modified cohomology; therefore $f^\#\alpha \neq 0$. Thus $\mathrm{supp}(\alpha) \subseteq N_{M+D'}(f(\mathrm{supp}(f^\#\alpha)))$. Therefore

$$\mathrm{Mob}([\alpha_0], D') \subseteq N_{M+D'}(f(\mathrm{Mob}(f^*[\alpha_0], D))).$$

Using a coarse inverse $g : Y \to X$ to f, the same argument shows that

$$f(\mathrm{Mob}(f^*[\alpha_0], D)) \subseteq N_R(\mathrm{Mob}([\alpha_0], D''))$$

for some suitable D'' and R. Lemma 3.36 ensures that, provided both are non-empty, $\mathrm{Mob}([\alpha_0], D'')$ and $\mathrm{Mob}([\alpha_0], D')$ have finite Hausdorff distance, hence $\mathrm{Mob}([\alpha_0], D')$ and $f(\mathrm{Mob}(f^*[\alpha_0], D))$ have finite Hausdorff distance. □

3.2.3 Proof of Theorem 1.1

Lemma 3.39. *Let G be a finitely generated group with $H \leq G$ a subgroup. Suppose there exists a constant $A \geq 0$ and a coarse complementary component C of H such both C and $G \setminus C$ are deep, and for every $g \in G$, either $gH \subseteq C \cup N_A(H)$ or $gH \subseteq (G \setminus C) \cup N_A(H)$. Then there exists a proper H-almost invariant set X such that $XH = X$.*

Proof. Let $X := \{g \in G \mid gH \subseteq C \cup N_A(H)\}$. It is clear that $XH = X$, so we need only show that X is a proper H-almost invariant set. Suppose $g \in C \backslash N_A(H)$. Then $gH \cap (C \backslash N_A(H)) \neq \emptyset$ so $gH \subseteq C \cup N_A(H)$; therefore $g \in X$. Conversely, suppose $g \in X \backslash N_A(H)$. Then because $g \in gH \subseteq C \cup N_A(H)$, it follows $g \in C$. Thus C and X are equal outside $N_A(H)$. In particular, X and $G \setminus X$ are deep coarse complementary components of H (since C and $G \setminus C$ are), so Proposition 3.14 tells us that X is a proper H-almost invariant set. □

The proof of the following Lemma is based on an analogous statement for the $n = 1$ case found in [Kle02].

Lemma 3.40. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and let $W \subseteq G$ be a coarse $PD_n^{\mathbb{Z}_2}$ subspace. Suppose G contains three essential, coarse disjoint, coarse complementary components of W. Then there exists a subgroup $H \leq G$, at finite Hausdorff distance from W, and a proper H-almost invariant set X such that $XH = X$.*

Proof. Say C_1, C_2 and C_3 are three essential, coarse disjoint, coarse complementary components of W. Without loss of generality, we may replace C_1 with $G\backslash(C_2 \cup C_3)$ and C_2 with $C_2\backslash C_3$, so that C_1, C_2 and C_3 are disjoint. Using Corollary 3.6, we may assume C_1, C_2 and C_3 are each $(1, A)$-coarse complementary components for some $A \geq 0$ large enough.

By Proposition 3.19, there exist uniformly $(n-1)$-acyclic metric complexes $(W, D_\bullet)$, $(C_1 \cup W, A_\bullet)$, $(C_2 \cup C_3 \cup W, B_\bullet)$ and $(G, C_\bullet)$, and an exact sequence

$$\cdots \xrightarrow{q^*} H_c^n(A_\bullet) \oplus H_c^n(B_\bullet) \xrightarrow{p^*} H_c^n(D_\bullet) \xrightarrow{\delta} \hat{H}_c^{n+1}(C_\bullet).$$

As C_1 and $C_2 \cup C_3$ are essential, the map $\delta : H_c^n(D_\bullet) \to \hat{H}_c^{n+1}(C_\bullet)$ is injective and has image containing some non-trivial $[\alpha_1] \in \hat{H}_c^{n+1}(C_\bullet)$.

As in the proof of Proposition 3.30, there exists a $D \geq 0$ such that for every $x \in W$, there is a cocycle supported in $N_D^W(x)$ representing the non-trivial element of $H_c^n(D_\bullet) \cong \mathbb{Z}_2$. Thus Proposition 3.19 tells us that there exists a $D' = D'(D, G, W, A) \geq 0$ such that for each $x \in W$, $[\alpha_1] \in \hat{H}_c^{n+1}(C_\bullet)$ can be represented by a cocycle supported in $N_{D'}^G(x)$. Therefore $W \subseteq N_{D'}(\mathrm{Mob}([\alpha_1], D'))$.

Part (2) of Proposition 3.19 and Lemma 3.35 say that for s sufficiently large, $[\alpha_1] \neq 0$ can be represented by a cocycle supported in $C_1\backslash N_s(W)$, but not in either of $C_2\backslash N_s(W)$ or $C_3\backslash N_s(W)$. Therefore, for any $\sigma \in Z([\alpha_1], D')$, $\mathrm{supp}(\sigma)$ cannot intersect either $C_2\backslash N_{s+D'}(W)$ or $C_3\backslash N_{s+D'}(W)$. Thus $\mathrm{Mob}([\alpha_1], D') \subseteq C_1 \cup N_{s+D'}(W)$.

The constants s and D' depend only on G, W, and A. Thus we can interchange the roles of C_1 and C_2 to obtain a non-trivial cohomology class $[\alpha_2] \in H_{\mathrm{coarse}}^{n+1}(X)$ such that $W \subseteq N_{D'}(\mathrm{Mob}([\alpha_2], D'))$ and $\mathrm{Mob}([\alpha_2], D') \subseteq C_2 \cup N_{s+D'}(W)$.

Let $H_i = \mathrm{stab}([\alpha_i])$ and $F_i = \mathrm{Mob}([\alpha_i], D')$ for $i = 1, 2$. Since each F_i is non-empty, Corollary 3.37 says that each H_i has finite Hausdorff distance from F_i. Thus there is an $R \geq D'$ large enough so that $H_i \subseteq N_R(F_i)$ and $F_i \subseteq N_R(H_i)$ for $i = 1, 2$. By Lemma 3.10,

$$H_i \subseteq N_R(F_i) \subseteq N_R(C_i \cup N_{s+D'}(W)) \subseteq C_i \cup N_{s+D'+A+R}(W)$$

for $i = 1, 2$. Hence

$$H_1 \cap H_2 \subseteq N_R(F_1) \cap N_R(F_2) \subseteq N_{s+D'+A+R}(W)$$

and

$$W \subseteq N_{D'}(F_1) \cap N_{D'}(F_2) \subseteq N_{R+D'}(H_1) \cap N_{R+D'}(H_2).$$

By Lemma 2.10, $H := H_1 \cap H_2$ has finite Hausdorff distance from $N_{R+D'}(H_1) \cap N_{R+D'}(H_2)$. Therefore, W has finite Hausdorff distance from $H := H_1 \cap H_2$.

Let $g \in G$. We recall that for each $x \in W$, there exists a cocycle supported in $N_{D'}(x)$ representing $[\alpha_1] \neq 0$. Hence for each $x \in gW$, there is a cocycle supported in $N_{D'}(x)$ representing $[\alpha_1 \cdot g^{-1}] \neq 0$; therefore $gW \subseteq N_{D'}(\mathrm{Mob}([\alpha_1 g^{-1}], D'))$.

By Lemma 3.35 and our choice of s, $[\alpha_1 \cdot g^{-1}]$ cannot be represented by cocycles supported in any two of $C_1 \backslash N_s(W)$ and $C_2 \backslash N_s(W)$ and $C_3 \backslash N_s(W)$. Therefore, either $\mathrm{Mob}([\alpha_1 g^{-1}], D') \subseteq C_1 \cup N_{s+D'}(W)$ or $\mathrm{Mob}([\alpha_1 g^{-1}], D') \subseteq C_2 \cup C_3 \cup N_{s+D'}(W)$. Hence by Lemma 3.10, either $gW \subseteq C_1 \cup N_{s+D'+A}(W)$ or $gW \subseteq C_2 \cup C_3 \cup N_{s+D'+A}(W)$.

Since H and W are at finite Hausdorff distance, C_1 is a coarse complementary component of H. Moreover, both C_1 and $G \backslash C_1 = C_2 \cup C_3$ are deep. By a further application of Lemma 3.10, there exists a sufficiently large A' such that for every $g \in G$, either $gH \subseteq C_1 \cup N_{A'}(H)$, or $gH \subseteq (G \backslash C_1) \cup N_{A'}(H)$. Thus C_1 satisfies the hypotheses of Lemma 3.39, so there exists a proper H-almost invariant set X such that $XH = X$. □

The following theorem is what is needed to construct a splitting. It is a special case of a more general statement known as the Kropholler conjecture. See [Dun17] for more details.

Theorem 3.41 ([DR93],[DS00])**.** *Let G be a finitely generated group and let $H \leq G$ be a subgroup. Suppose $\tilde{e}(G, K) = 1$ for every infinite index subgroup K of H. If G contains a proper, H-almost invariant subset X such that $XH = X$, then G admits a splitting over a subgroup $C \leq G$ that is commensurable to H.*

We note that this theorem is stated in [DR93] under the assumption that H is a polycyclic-by-finite group. However, the argument of [DR93] holds for general H under the additional assumption that $\tilde{e}(G, K) = 1$ for every infinite index subgroup K of H (as stated in [DS00, Theorem 3.4]). This theorem also appears in [SS00] and [Nib02].

We are now in a position to prove our main theorem, a special case of which was stated in the introduction as Theorem 1.1.

Theorem 3.42. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and let $W \subseteq G$ be a coarse $PD_n^{\mathbb{Z}_2}$ subspace. Suppose G contains three essential, coarse disjoint, coarse complementary*

components of W and that W is almost essentially embedded. Then there exists a subgroup $H \leq G$, at finite Hausdorff distance from W, such that G splits over H.

Proof. We deduce from Lemma 3.40 that W is at finite Hausdorff distance from some subgroup H, and there exists a proper H-almost invariant set X such that $XH = X$. By Proposition 3.32, we deduce that $\tilde{e}(G, K) = 1$ for every infinite index subgroup K of H. Theorem 3.41 tells us that G admits a splitting over a subgroup C which is commensurable with a subgroup of H. It follows from Proposition 2.12 that C has finite Hausdorff distance from W. □

3.2.4 Applications

Theorem 1.2, stated in the introduction, can easily be deduced from Theorem 3.42 and Proposition 3.34. We note that the subgroups H and H' in Theorem 1.2 are themselves quasi-isometric. Using [BM00], we can deduce Corollary 1.3 from Theorem 1.2.

For $k > n$, coarse $PD_k^{\mathbb{Z}_2}$ groups are $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$. Thus Corollary 1.3 has some overlap with Theorem 1.7 from [MSW11] and Theorem 3.4 from [Pap07], which concern graphs of groups whose vertices and edges are coarse Poincaré duality groups. However, Corollary 1.3 applies to vertex groups that are $(n-1)$-acyclic at infinity which need not be coarse Poincaré duality groups.

Under the additional assumption that H is virtually polycyclic, we can improve Theorem 1.2 by dropping the hypothesis that no infinite index subgroup coarsely separates.

Corollary 3.43. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$. Suppose $H \leq G$ is a VPC-n subgroup that coarsely 3-separates G. Then for any quasi-isometry $f : G \to G'$, there is a subgroup $H' \leq G'$, at finite Hausdorff distance from $f(H)$, such that G' splits over H'.*

Proof. Example 2.59 says that a torsion-free polycyclic group of Hirsch length r is a $PD_r^{\mathbb{Z}}$ group, hence it must also be a coarse $PD_r^{\mathbb{Z}_2}$ group. Every virtually polycyclic group of Hirsch length r contains a finite index, torsion-free polycyclic subgroup, hence is a coarse $PD_r^{\mathbb{Z}_2}$ group.

By Theorem 1.2, it is sufficient to show that $\tilde{e}(G, K) = 1$ for every infinite index $K \leq H$. Let $K \leq H$ be a subgroup of minimal Hirsch length such that $\tilde{e}(G, K) > 1$. We suppose for contradiction that $r := h(K) < h(H) = n$. As K is itself a coarse $PD_r^{\mathbb{Z}_2}$ group, Proposition 3.34 says that K is essentially embedded. Therefore, there

exists essential coarse complementary components C_1 and $C_2 = G \setminus C_1$ of H. Hence there exists a coarse Mayer–Vietoris sequence

$$\cdots \to H^r_{\text{coarse}}(C_1 \cup K) \oplus H^r_{\text{coarse}}(C_2 \cup K) \to H^r_{\text{coarse}}(K) \to H^{r+1}_{\text{coarse}}(G) \to \cdots .$$

As C_1 and C_2 are essential, the boundary map $H^r_{\text{coarse}}(K) \to H^{r+1}_{\text{coarse}}(G)$ is injective. Since G is $(n-1)$-acyclic at infinity, $H^{r+1}_{\text{coarse}}(G) = 0$, which is a contradiction. Hence there is no infinite index subgroup $K \leq H$ such that $\tilde{e}(G, K) > 1$. □

Theorem 3.44. *Let G be a group of type $FP^{\mathbb{Z}_2}_{n+1}$ that is the fundamental group of a finite graph of groups $\mathcal{G}$. Suppose $\mathcal{G}$ contains an edge e with associated edge monomorphisms $i_0 : G_e \to G_v$ and $i_1 : G_e \to G_w$ such that the following holds: G_e is a coarse $PD^{\mathbb{Z}_2}_n$ group and $|\text{Comm}_{G_v}(i_0(G_e)) : i_0(G_e)|$ and $|\text{Comm}_{G_w}(i_1(G_e)) : i_1(G_e)|$ are both greater than one and not both equal to two. If $f : G \to G'$ is a quasi-isometry, then $f(G_e)$ has finite Hausdorff distance from some subgroup $H' \leq G'$, and G' splits over a subgroup commensurable to a subgroup of H'.*

We note that unlike Theorem 1.2 and Corollary 3.43, which assume geometric hypotheses, the hypotheses in Theorem 3.44 are more algebraic in nature. In the following proof, we use the concept of a tree of spaces as in [Sco78]. This will be discussed in more detail in Section 4.1.

Proof. Let $A = G_v$, $B = G_w$ and $H = G_e$. We may choose $a \in \text{Comm}_A(H) \backslash H$ and $b \in \text{Comm}_B(H) \backslash H$. Thus aH and bH have finite Hausdorff distance from H. Let $\tilde{X}$ be a tree of spaces associated to the graph of groups $\mathcal{G}$ and let $p : \tilde{X} \to T$ be an equivariant projection map onto the Bass-Serre tree T. As G acts properly and cocompactly on the 1-skeleton of $\tilde{X}$, the Milnor-Švarc lemma tells us there is a quasi-isometry $g : \tilde{X}^{(1)} \to G$ such that $g(p^{-1}(e))$ has finite Hausdorff distance from $H \leq G$.

We choose a geodesic L in T with edges $\ldots, babH, baH, bH, H, aH, abH, abaH, \ldots$. The edge spaces associated to any two consecutive edges of L have uniformly finite Hausdorff distance from one another. We define $W := \cup_{e \in L} p^{-1}(e)$, as e ranges over all the edges in the line L. Then W is a *coarse fibration* as described by Kapovich and Kleiner in Section 11.5 of [KK05].

This is similar to Example 11.11 in [KK05], except we don't require H to have finite index in A and B. We take the line $\hat{L}$ dual to the geodesic L in the graph of groups; each vertex of $\hat{L}$ is an edge of L, and two vertices in $\hat{L}$ are joined by an edge if the corresponding edges in L share a vertex. To each vertex of $\hat{L}$ we associate the

corresponding edge space X_e, which is a coarse $PD^{\mathbb{Z}_2}_n$ space. We thus see that W is a coarse fibration whose base is $\hat{L}$ and whose fibres are edge spaces. It follows from Theorem 11.13 of [KK05] that W is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ space.

An edge e of L separates T into two components T^+ and T^-. Thus $C_1 := p^{-1}(T^+)$ and $C_2 := p^{-1}(T^-)$ are two deep, coarse disjoint, coarse complementary components of X_e in $\tilde{X}$. As C_1 and C_2 contain the coarse $PD^{\mathbb{Z}_2}_{n+1}$ half-spaces $(p^{-1}(T^+) \cap W, X_e)$ and $(p^{-1}(T^-) \cap W, X_e)$ respectively, Corollary 3.29 tells us that C_1 and C_2 are essential.

As either $|\mathrm{Comm}_A(H) : H|$ or $|\mathrm{Comm}_B(H) : H|$ is at least three, we can choose another geodesic L' in T that contains the edge e and intersects L in a finite subtree of T. A slight modification of the above argument shows that there are three essential, coarse disjoint, coarse complementary components of X_e. Since $g : \tilde{X}^{(1)} \to G$ is a quasi-isometry and $g(X_e)$ has finite Hausdorff distance from H, Proposition 3.24 tells us that there exist three essential, coarse disjoint, coarse complementary components of $H \leq G$. We therefore apply Theorem 1.1. □

We remark that if $|G_v : i_0(G_e)|$ and $|G_w : i_1(G_e)|$ are finite, then $\mathrm{Comm}_{G_v}(i_0(G_e)) = G_v$ and $\mathrm{Comm}_{G_w}(i_1(G_e)) = G_w$. Therefore, if G is the fundamental group of a finite graph of groups $\mathcal{G}$ in which all edge and vertex groups are coarse $PD^{\mathbb{Z}_2}_n$ groups (for some fixed n), our result overlaps with Theorem 2 of [MSW03] and Theorem 3.1 of [Pap07]. However, Theorem 3.44 applies in situation in which edge groups don't necessarily have finite index in adjacent vertex groups.

Under an even stronger algebraic condition, namely that the subgroup over which the group splits commensurises the group, we deduce Theorem 3.45.

Theorem 3.45. *Let G be a group of type $FP^{\mathbb{Z}_2}_{n+1}$ and suppose G splits over a coarse $PD^{\mathbb{Z}_2}_n$ subgroup H and that $\mathrm{Comm}_G(H) = G$. Let $f : G \to G'$ be a quasi-isometry. Then one of the following occurs:*

(1) G' is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ group;

(2) $f(H)$ has finite Hausdorff distance from some $H' \leq G'$, and G' splits over H'.

Proof. We consider the tree of spaces $\tilde{X}$ associated to the splitting of G over H, and let $p : \tilde{X} \to T$ be the equivariant projection to the Bass-Serre tree. As $\mathrm{Comm}_G(H) = H$, any two edge spaces have finite Hausdorff distance from one another. Therefore, either T is a line or H coarsely 3-separates G. In the former case, the argument in the proof of Theorem 3.44 tells us that G is coarse fibration whose base is a line and whose fibres are coarse $PD^{\mathbb{Z}_2}_n$ groups. Thus Theorem 11.13 of [KK05] tells us that G is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ group, so G' is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ group.

In the case where H coarsely 3-separates G, we see that for every coarse complementary component C of H, $p(C)$ contains a ray in the Bass-Serre tree in which consecutive edge spaces have uniformly bounded Hausdorff distance from one another. The argument in the proof of Theorem 3.44 tells us that C contains a coarse $PD^{\mathbb{Z}_2}_{n+1}$ half-space, so must be essential. Thus H is essentially embedded. Since H is essentially embedded and coarsely 3-separates G, Theorem 3.42 tells us G' splits over a subgroup at finite Hausdorff distance from $f(H)$. $\square$

Chapter 4

The geometry of JSJ decompositions

Bowditch [Bow98] shows that the structure of cut pairs in the boundary of a hyperbolic group can be encoded in a tree known as the JSJ tree. Papasoglu and Swenson use a similar construction for CAT(0) groups [PS09]. See also [PS06] for a general construction associating an $\mathbb{R}$-tree to the cut-point structure of continuum. One limitation of these methods is that the group under consideration is required to have a boundary. Moreover, even if the group does have a boundary, splittings over distorted subgroups cannot be detected.

The main construction of this section is motivated by work of Dunwoody–Swenson [DS00], Bowditch [Bow02], and Scott–Swarup [SS03], who use the structure of almost invariant sets to construct a tree. These methods are less geometric than those using the boundary. In particular, we need to know that the group has a codimension one subgroup, so the resulting tree is not necessarily a quasi-isometry invariant of the group. Instead of working with almost invariant sets, we work with coarse complementary components.

In Section 4.2, we associate a combinatorial structure known as a *pretree* to a collection of coarse complementary components of a finitely generated group. In certain situations, this pretree can be embedded into a simplicial tree. We describe this pretree construction in more generality than is needed, with the hope that the framework can be used in the quasi-isometric rigidity of group splittings in other contexts.

The following theorem, which we prove in Section 4.3, is the main result of this section. It partially generalises the Scott-Swarup regular neighbourhood [SS03] to the coarse setting. We do not state the precise definition of enclosing here, but refer the reader to Definition 4.37.

Theorem 4.1. *Let G be a group of type $FP^{\mathbb{Z}_2}_{n+1}$ and suppose $\{W_1, \dots, W_m\}$ is a non-empty collection of essentially embedded coarse $PD^{\mathbb{Z}_2}_n$ subspaces that coarsely separate G. Then there is a minimal G-tree $T = T(W_1, \dots, W_m)$, bipartite on $VT = V_0 \sqcup V_1$, such that:*

(1) Every non-trivial coarse complementary component of gW_i, for any $g \in G$, is enclosed by a V_0-vertex of T. Conversely, every V_0-vertex of T encloses such a coarse complementary component.

(2) Suppose that G_v is the vertex stabilizer of a V_0-vertex v . Then either:

(a) a conjugate of G_v is at finite Hausdorff distance from some W_i;

(b) $G_v = \mathrm{Comm}_G(H)$, where $H \leq G$ is a subgroup such that some conjugate of H is at finite Hausdorff distance from some W_i;

(c) $G_v = \mathrm{Stab}([\alpha])$ for some non-zero coarse cohomology class $[\alpha] \in H^{n+1}_{\mathrm{coarse}}(G)$.

Moreover, if T consists of a single vertex, then it is a V_0-vertex of type (b) or (c).

The main shortcoming of this theorem is that one doesn't know what the edge stabilizers adjacent to vertices of type (c) are. Vertices of type (c) that are not also of type (a) or (b), generalise quadratically hanging vertices of JSJ theory. The cohomology class $[\alpha]$ can be thought of as a coarse fundamental class supported on a subspace of G, namely the mobility sets used in Section 3.2.2. Motivated by work of Bowditch [Bow04] and Kleiner [Kle02], one might wish to think of these mobility sets as coarsely resembling the fundamental group of an $(n+1)$-manifold, whose boundary components correspond to incident edge groups. One would significantly strengthen this analogy, and Theorem 4.1, if one could show that edge stabilizers incident to such vertices are coarse $PD^{\mathbb{Z}_2}_n$ groups.

To help us interpret Theorem 4.1, we introduce a class of groups generalising coarse Poincaré duality groups:

Definition 4.2. A group G is a *coarse n-manifold group* if it is of type $FP^{\mathbb{Z}_2}_n$ and $H^n(G, \mathbb{Z}_2 G)$ has a non-zero, finite dimensional, G-invariant subspace.

It is shown in [Kle02] that a group G is a coarse 2-manifold group if and only if it is virtually a surface group, i.e. the fundamental group of a closed surface that is not a sphere or projective plane. This is also proved in [Bow04]. A result of Farrell [Far75] states that for any $n \geq 1$, if G is a coarse n-manifold group of type FP that is uniformly $(n-2)$-acyclic at infinity over $\mathbb{Z}_2$, then it is a $PD^{\mathbb{Z}_2}_n$ group.

The following lemma shows that we may think of coarse n-manifold groups as a generalisation of coarse $PD_n^{\mathbb{Z}_2}$ groups.

Lemma 4.3. *Let G be a group of type $FP_n^{\mathbb{Z}_2}$ and let $(G, C_\bullet)$ be a uniformly $(n-1)$-acyclic R-metric complex admitting a G-action. Then G is a coarse n-manifold group if and only if there exists a nonzero $[\alpha_0] \in \hat{H}_c^n(C_\bullet)$ and a $D \geq 0$ such that $G = N_D(\mathrm{Mob}([\alpha_0], D))$. In particular, if G is a coarse $PD_n^{\mathbb{Z}_2}$ group, then it is a coarse n-manifold group.*

Proof. Suppose that G is a coarse n-manifold group and that $F \leq H^n(G, \mathbb{Z}_2 G)$ is a non-zero, finite dimensional, G-invariant subspace. Since F is finite dimensional, there is a nonzero $[\alpha_0] \in F$ such that $\mathrm{Stab}([\alpha_0])$ has finite index in G. By Proposition 2.49, we identify $[\alpha_0]$ with an element of $\hat{H}_c^n(C_\bullet)$. It follows from Lemma 3.36 that for some D sufficiently large, $G = N_D(\mathrm{Mob}([\alpha_0], D))$.

Now suppose $G = N_D(\mathrm{Mob}([\alpha_0], D))$ for some non-zero $[\alpha_0] \in \hat{H}_c^n(C_\bullet)$. Then Lemma 3.36 tells us that $\mathrm{Stab}([\alpha_0])$ has finite index in G. Hence $[\alpha_0]G$ can be identified with a finite dimensional subspace of $H^n(G, \mathbb{Z}_2 G)$ by Proposition 2.49. □

Proposition 3.38 and Lemma 4.3 imply that if G is a coarse n-manifold group and G' is quasi-isometric to G, then G' is also a coarse n-manifold group. We deduce the following corollary of Theorem 4.1, which we think of as a coarse analogue of the Algebraic Torus Theorem [DS00].

Corollary 4.4. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ that is coarsely 2-separated by an essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ subspace W. Then either:*

(1) G is a coarse $(n+1)$-manifold group;

(2) G splits over a subgroup at finite Hausdorff distance from W;

(3) G acts non-trivially (i.e. not fixing a vertex) and minimally on a tree T, containing a vertex $v \in T$ enclosing W, where $G_v = \mathrm{Stab}([\alpha])$ for some non-trivial cohomology class $[\alpha] \in H^{n+1}(G, \mathbb{Z}_2 G)$.

Proof. If W coarsely 3-separates G, then we apply Theorem 3.42. If not, then we apply Theorem 4.1 and obtain a tree $T = T(W)$. It will be clear from the construction of T in Section 4.3 that if W doesn't coarsely 3-separate G, then every vertex of T is of type (c). As W is enclosed by a V_0-vertex, T has at least one V_0-vertex. If G fixes this vertex, then G is a coarse $(n+1)$-manifold group. Otherwise, we have (3). □

4.1 Groups acting on trees

Before proving these theorems, we first cover some preliminary topics concerning Bass-Serre theory and JSJ decompositions.

We fix notation and terminology for graphs. A graph Γ consists of the set $V\Gamma$ of *vertices*, the set $E\Gamma$ of *oriented edges* and maps $\iota, \tau : E\Gamma \to V\Gamma$ which map an oriented edge to its initial and terminal vertices respectively. For every $e \in E\Gamma$, we let $\bar{e}$ denote the edge with opposite orientation to e. We call the pair $\{e, \bar{e}\}$ an *unoriented edge*. A *graph isomorphism* $\chi : \Gamma \to \Gamma'$ is map which bijectively sends vertices to vertices, edges to edges, and commutes with the functions ι and τ. A graph is *simplicial* if it has no one edge loops and at most one unoriented edge between any pair of vertices.

A G-tree is a tree T which admits a group action of G by isomorphisms. A G-tree is *minimal* if it has no proper G-invariant subtree; in particular, unless T has a single vertex, a minimal G-action cannot fix a vertex of T. Given a G-tree, we let G_v and G_e denote the stabilizers of a vertex $v \in VT$ and an edge $e \in ET$ respectively. If G acts minimally on T and without edge inversions, then Bass-Serre theory gives a description of G as the fundamental group of a graph of groups; see [Ser77] and [SW79] for details.

Vertex and edge spaces

Suppose G is finitely generated and that Γ is the Cayley graph of G with respect to some finite generating set. Given a minimal G-tree, we fix a basepoint $v_0 \in VT$ and define the G-equivariant map $p : \Gamma \to T$ via $g \mapsto g \cdot v_0$ and extending linearly on edges of Γ. For each vertex $v \in VT$, we define the *vertex space* X_v to be $p^{-1}(v)$. Similarly, for each edge $e \in ET$, the *edge space* X_e is defined to be $p^{-1}(m_e)$, where m_e is the midpoint of the edge e. We will, for the moment at least, just think of vertex and edge spaces as subspaces of the Cayley graph Γ, equipped with the subspace metric.

Remark 4.5. Care must be taken when working with vertex and edge spaces, as even in some natural situations, such as when working with the Scott–Swarup regular neighbourhood of a finitely presented group over two-ended subgroups (see Section 4.2.3), one may have vertex and edge stabilizers that are not finitely generated. If we build a tree of spaces as in [SW79] and [MSW11], there is no natural way to metrise vertex and edge spaces. To circumvent this issue, we think of vertex and edge spaces as subsets of the ambient Cayley graph, which does have a natural well-defined metric up to quasi-isometry.

Suppose $\mathcal{G}$ is a finite graph of groups and $G := \pi_1(\mathcal{G})$, with Γ a Cayley graph of G. Let T be the Bass-Serre tree of $\mathcal{G}$. Every vertex or edge of T is identified with a coset gG_v or gG_e of a vertex or edge group of $\mathcal{G}$. The following proposition relates vertex and edge spaces to their corresponding cosets in G.

Proposition 4.6. *Suppose $\mathcal{G}$, G, and T are as above. There is a constant $A \geq 0$ such that the following holds. Suppose that G_v and G_e are vertex and edge groups of $\mathcal{G}$, and that $\tilde{v} = gG_v \in VT$ and $\tilde{e} = gG_e \in ET$ are the corresponding vertex and edges in the Bass-Serre tree. Then $d_{\mathrm{Haus}}(gG_v, X_{\tilde{v}}) \leq A$ and $d_{\mathrm{Haus}}(gG_e, X_{\tilde{e}}) \leq A$.*

Proof. Let G_v be a vertex group of $\mathcal{G}$, and let $\tilde{v}$ denote the vertex of T corresponding to the coset G_v; thus G_v is the vertex stabilizer of $\tilde{v}$. As the action of G on Γ is cobounded and $G_v X_{\tilde{v}} = X_{\tilde{v}}$, we can choose an $A_v \geq 0$ such that the quotient space $G_v \backslash X_{\tilde{v}}$ has diameter A_v. Let $x \in X_{\tilde{v}}$. We therefore see that $d_{\mathrm{Haus}}(G_v x, X_{\tilde{v}}) \leq A_v$. Thus G_v, thought of as a subspace of the Cayley graph equipped with the subspace metric, has finite Hausdorff distance from $X_{\tilde{v}}$. The same argument holds for each edge group G_e of $\mathcal{G}$. One chooses A large enough so that every edge and vertex group of $\mathcal{G}$ (there are only finitely many), has Hausdorff distance at most A from the corresponding vertex or edge space. The result follows by the equivariance of p. □

One consequence of this proposition is that vertex and edge spaces are uniquely defined up to Hausdorff distance, i.e. they don't depend on the choice of basepoint v_0.

We now consider the more restrictive setting where $\mathcal{G}$ is a graph of finitely generated groups. In this situation, we can build a tree of spaces that is quasi-isometric to G. We use the following notation for trees of spaces as in [CM17].

Definition 4.7. A *tree of spaces* $X := X(T, \{X_v\}_{v \in VT}, \{X_e\}_{e \in ET}, \{\alpha_e\}_{e \in ET})$ consists of:

(1) a simplicial tree T known as the *base tree*;

(2) a metric space X_v for each vertex $v \in VT$ known as a *vertex space*;

(3) a subspace $X_e \subseteq X_{\iota e}$ for each oriented edge e of T known as an *edge space*;

(4) maps $\alpha_e : X_e \to X_{\bar{e}}$ for each edge $e \in ET$, such that $\alpha_{\bar{e}} \circ \alpha_e = \mathrm{id}_{X_e}$ and $\alpha_e \circ \alpha_{\bar{e}} = \mathrm{id}_{X_{\bar{e}}}$.

We consider X as a metric space as follows: we take the disjoint union of all the X_v and then, for all unoriented edges $\{e, \bar{e}\}$ and every $x \in X_e$, we attach a unit interval between $x \in X_e$ and $\alpha_e(x) \in X_{\bar{e}}$.

As noted in [CM17], it is easy to verify that X is actually a metric space, i.e. that the quotient pseudometric on X is actually a metric. In [CM17], a slightly weaker definition is used, where each $\alpha_{\tilde{e}} \circ \alpha_e$ is only assumed to be close to the identity. That will not be needed here, since all trees of spaces will arise 'algebraically' using Proposition 4.8.

The following proposition, whose argument is standard, relates the algebra of the Bass-Serre tree to the geometry of the corresponding tree of spaces obtained by 'blowing up' the Bass-Serre tree.

Proposition 4.8 (Section 2.5 of [CM17])**.** *Suppose $\mathcal{G}$, G, Γ and T are as above. Then there exists a tree of spaces X with base tree T and an equivariant quasi-isometry $q : \Gamma \to X$.*

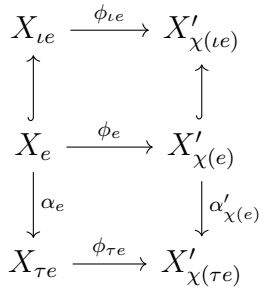
Proof. We choose a generating set of G that is a union of generating sets of each of the vertex groups of $\mathcal{G}$, as well as generators of the fundamental group of the underlying graph of $\mathcal{G}$. We define the vertex space of each $\tilde{v} = gG_v \in VT$ to be the coset gG_v, thought of as a subspace of G equipped with the word metric. Suppose $\tilde{e} = gG_e$ is an edge of T. There are monomorphisms $\beta_0 : G_e \to G_{\iota(e)}$ and $\beta_1 : G_e \to G_{\tau(e)}$ associated with the edge e of $\mathcal{G}$. We define $X_{\tilde{e}} := g \cdot \mathrm{im}(\beta_0)$ and $\alpha_{\tilde{e}}$ by $gx \mapsto g(\beta_1 \circ \beta_0^{-1})(x)$. It is easy to verify that X is a proper geodesic metric space (see Lemma 2.13 of [CM17]) and that G acts properly discontinuously and cocompactly on X. Hence for any $x \in X$, the orbit map $g \mapsto g \cdot x$ is a quasi-isometry. □

If we choose basepoints correctly, the map $p : \Gamma \to T$ defined above factors as $\Gamma \xrightarrow{q} X \xrightarrow{r} T$, where r is the equivariant collapsing map that sends vertex spaces of X to vertices of T, and edge spaces of X to edges of T. Note that we now have two different definitions of vertex spaces, both of which are denoted X_v. We live with this equivocacy, observing that the equivariant quasi-isometry $\Gamma \xrightarrow{q} X$ identifies these two different X_v up to finite Hausdorff distance.

Our reason for introducing trees of spaces, is that it allows one to build quasi-isometries between groups by patching together quasi-isometries between vertex spaces. This can be done if quasi-isometries on adjacent vertex spaces agree up to uniformly bounded error on their common edge space. This method will be used in Chapter 5 to classify certain RAAGs up to quasi-isometry.

Proposition 4.9 ([CM17, Corollary 2.16])**.** *Let $K \geq 1$ and $A \geq 0$. Suppose $X := X(T, \{X_v\}, \{X_e\}, \{\alpha_e\})$ and $X' := X'(T', \{X'_v\}, \{X'_e\}, \{\alpha'_e\})$ are trees of spaces, and that there is a tree isomorphism $\chi : T \to T'$. Suppose for every $v \in VT$ and $e \in ET$*

there is a (K, A)-quasi-isometry $\phi_v : X_v \to X'_{\chi(v)}$ and $\phi_e : X_e \to X'_{\chi(e)}$. Suppose also that for every $e \in ET$, the following diagram commutes up to uniformly bounded error A. Then there is a quasi-isometry $\phi : X \to X'$ such that $\phi|_{X_v} = \phi_v$ for every $v \in VT$.



JSJ decompositions

We now define JSJ decompositions, following the approach of Guirardel and Levitt [GL17], though in less generality. See also [RS97],[DS99] and [FP06] etc. We fix a finitely generated group G and a class of subgroups $\mathcal{E}$ such that $\mathcal{E}$ is closed under conjugation and passing to finite index subgroups. We assume that no element of $\mathcal{E}$ is contained in another with infinite index, and that G does not split over any infinite index subgroup of an element of $\mathcal{E}$.

For example, if G is a one-ended group and $\mathcal{E}$ is the set of all two-ended subgroups, then the above hypotheses are indeed satisfied. We say that a G-tree T is an $\mathcal{E}$-tree if edge stabilizers of T are in $\mathcal{E}$. For the remainder of this section, unless otherwise stated, all trees will be minimal G-trees whose edge stabilizers are in $\mathcal{E}$.

Given two trees S and T, we say that T *dominates* S if there exists an equivariant map $T \to S$. This is equivalent to saying that every vertex stabilizer of T fixes a vertex of S. Bestvina and Feighn show that one can obtain S from T as a sequence of finitely many folding and collapse moves, provided that edge stabilizers are finitely generated [BF91]. Informally, one should think of this as saying that T is "bigger" than S. We define an equivalence class on trees, where $T \sim S$ if both T dominates S and S dominates T, and call equivalence classes *deformation spaces*. A deformation space $\mathcal{D}$ *dominates* a deformation space $\mathcal{D}'$ if there exists a $T \in \mathcal{D}$ and $S \in \mathcal{D}'$ such that T dominates S.

In order to understand all ways in which a group splits, one might hope to take a "largest" deformation space that dominates every other deformation space. Such a deformation space exists when $\mathcal{E}$ consists of all finite subgroups of a finitely presented

group G, see [Dun85, GL17]. However, such deformation spaces do not necessarily exist: in general no such deformation space exists when G is one-ended and $\mathcal{E}$ is the class of two-ended subgroups, as we will shortly see when G is the fundamental group of a closed hyperbolic surface.

An action of a group on a tree T is said to be *elliptic* if it fixes a vertex of T. A group is said to be *universally elliptic* if it acts elliptically on every $\mathcal{E}$-tree; we define $\mathcal{E}_{\text{ell}}$ to be the set of all universally elliptic elements of $\mathcal{E}$.

Definition 4.10. The *JSJ deformation space* $\mathcal{D}_{\text{JSJ}}$ is the deformation space of $\mathcal{E}_{\text{ell}}$-trees that dominates every other deformation space of $\mathcal{E}_{\text{ell}}$-trees. Elements of $\mathcal{D}_{\text{JSJ}}$ are called *JSJ trees over* $\mathcal{E}$.

Alternatively, one can define a JSJ tree directly: T is a JSJ tree over $\mathcal{E}$ if and only if it is an $\mathcal{E}_{\text{ell}}$-tree that dominates every other $\mathcal{E}_{\text{ell}}$-tree. Note that the deformation space $\mathcal{D}_{\text{JSJ}}$, if it exists, is unique, whilst the choice of JSJ tree is generally not unique.

Theorem 4.11 ([Dun85],[GL17])**.** *Let G be a group of type $FP_2^{\mathbb{Z}_2}$ and $\mathcal{E}$ a family of subgroups closed under conjugation and taking subgroups. Then the JSJ deformation space $\mathcal{D}_{\text{JSJ}}$ exists and contains a tree with finitely generated vertex and edge stabilizers.*

A vertex of a JSJ tree is said to be *rigid* if its stabilizer is universally elliptic, and *flexible* otherwise. For example, let G be the fundamental group of a closed hyperbolic surface S and let $\mathcal{E}$ be the class of two-ended subgroups of G. For any $g \in \pi_1(S) \setminus \{1\}$, let C_g be an immersed closed geodesic representing g. We may pick a simple closed geodesic C in S that intersects C_g transversely. Note that C gives a splitting of G over a cyclic subgroup $E \in \mathcal{E}$, but g acts hyperbolically (without fixing a point) on the Bass-Serre tree associated to this splitting. Thus no non-trivial subgroup of G is universally elliptic, so the (unique) JSJ tree of G is a single flexible vertex. We will shortly see that for JSJ trees over certain classes of groups, this example of a flexible vertex is essentially the only possible one.

We now describe a different sort of JSJ tree (which may not be a JSJ tree as defined above) called the *JSJ tree of cylinders*. This was introduced in [GL11]. Let T be any $\mathcal{E}$-tree. The corresponding *tree of cylinders* T_{cyl} is defined as follows. We define an equivalence class on edges in T, where two edges are equivalent if and only if their edge stabilizers are commensurable. The union of all edges in an equivalence class is shown to be a subtree of T, and is called a *cylinder*.

The tree of cylinders is then defined to be a bipartite tree with vertex set $V_0 \sqcup V_1$, where V_0 is the set of cylinders of T and V_1 consists of vertices of T that lie in at

least two cylinders. There is an edge between $v \in V_0$ and $w \in V_1$ precisely when w is contained in v. We remark that even though edge stabilizers of T are in $\mathcal{E}$, edge stabilizers of T_{cyl} may not be.

In light of Propositions 2.12 and 4.6, there is an alternative characterisation of cylinders: two edges are equivalent if and only if the associated edge spaces of X are coarsely equivalent. This is because every subgroup of the form gHg^{-1} is coarsely equivalent to the coset gH. Cylinders thus correspond to coarse equivalence classes of edge spaces of T. We will mainly use this formulation of cylinders, which is more natural in the coarse geometric context.

Indeed, suppose we are given the set of coarse equivalence classes of edge spaces of a tree of cylinders T_{cyl} and forget the tree structure. We can then recover this tree structure by examining when an edge space is "between" two other edge spaces. This idea will be developed fully later in this chapter. However, if we are given coarse equivalence classes of edge spaces of the original tree T, we have no way of saying that a coarse equivalence class $[X_{e'}]$ lies between $[X_e]$ and $[X_{e''}]$ when e, e' and e'' all lie in the same cylinder of T.

Guirardel and Levitt [GL11] show that T_{cyl} depends only on the deformation space $\mathcal{D}$ containing T, and is independent of the choice of $T \in \mathcal{D}$. Given a JSJ tree T, the *JSJ tree of cylinders* is thus defined to be the tree of cylinders T_{cyl}. The JSJ tree of cylinders is canonical; in particular, it is $\text{Aut}(G)$ invariant.

We now restrict to the case where $\mathcal{E}$ consists of virtually polycyclic groups of Hirsch length n. We call such groups VPC-n. We remark the following theorem holds for more general classes of edge groups. However for simplicity, we state it only for JSJ trees over VPC-n subgroups. We recall the definition of quadratically hanging vertices (Definition 1.4).

Theorem 4.12 ([DS99] and [FP06])**.** *The stabilizer of a flexible vertex in a JSJ tree over VPC-n subgroups is quadratically hanging.*

Let T be a JSJ tree of G over VPC-n subgroups. As edge stabilizers incident to quadratically hanging vertices correspond to conjugacy classes of parabolics, we deduce that no two such edges have commensurable stabilizers. Thus a flexible vertex of T corresponds to a quadratically hanging vertex in the JSJ tree of cylinders. In particular, every edge in T_{cyl} adjacent to such a vertex has VPC-n stabilizer.

In some cases, e.g. hyperbolic groups, the JSJ tree of cylinders is itself a JSJ tree. In particular, the JSJ tree described by Bowditch [Bow98] is a JSJ tree of cylinders. As we see in Chapter 5, this will never be the case for one-ended RAAGs — edge

stabilizers of the JSJ tree of cylinders over infinite cyclic subgroups will always be one-ended, and so the JSJ tree of cylinders is never a JSJ tree.

There is a related construction by Scott-Swarup [SS03] called the *regular neighbourhood*. We outline this construction in Section 4.2.3. It is shown in [GL10] that the JSJ tree of cylinders is equivariantly isomorphic to a subdivision of the Scott-Swarup regular neighbourhood. This subdivision is important: it gives a bipartite structure that makes it easier to state various enclosing properties (see [SS03]). However, sometimes the JSJ tree of cylinders will be more useful, since then every edge is adjacent to a cylinder.

4.2 A pretree structure on coarse complementary components

Let G be a finitely generated group equipped with the word metric. Let $\mathcal{W}$ denote a collection of subspaces that are almost-essentially embedded (see Definition 3.31). For instance, $\mathcal{W}$ may contain finitely generated subgroups of G that coarsely separate G, yet have no infinite index subgroups that coarsely separate G. However, the definition of $\mathcal{W}$ is purely geometric.

In Proposition 3.14, we observed that $C \subseteq G$ is a coarse complementary of a subgroup $H \leq G$ if and only if it is an H-almost invariant subset. We use this observation to generalise the machinery of [Bow02] and [SS03] to the more general setting of subsets of G that are coarse complementary components of some $W \in \mathcal{W}$.

Convention 4.13. Throughout this section, we assume that all coarse complementary components in G are coarse complementary components of some $W \in \mathcal{W}$.

For each coarse complementary component C in G, we let C^* denote its complement $G - C$, which is also a coarse complementary component. We recall that a coarse complementary component C of $W \in \mathcal{W}$ is said to be *non-trivial* if both C and C^* are deep.

The following lemma shows that if C is a non-trivial coarse complementary component of W and Z, both of which are contained in $\mathcal{W}$, then W and Z are at finite Hausdorff distance. It thus makes sense to talk about a coarse complementary component C without referencing a particular choice of $W \in \mathcal{W}$.

Lemma 4.14. *Suppose that C is a non-trivial coarse complementary component of $W \in \mathcal{W}$. Then W is at finite Hausdorff distance from $\partial_1 C$.*

Proof. Suppose C is a non-trivial $(1, A)$-coarse complementary component of W; thus C and C^* are deep $(1, A)$-coarse complementary components of W. As W is almost essentially embedded, there exists a B such that $W \subseteq N_B(D \backslash N_A(W))$ for every deep $(1, A)$-coarse complementary component D of W. By Remark 3.2, it is clear that $\partial_1 C \subseteq N_{A+1}(W)$. For the converse, suppose that $x \in W \cap C$. Then there exists a $y \in C^* \backslash N_A(W)$ with $d(y, x) \leq B$. A 1-geodesic from x to y must contain some point in $\partial_1 C$, so we deduce that $x \in N_B(\partial_1 C)$. We use a similar argument for each $x \in W \cap C^*$, and thus deduce $W \subseteq N_B(\partial_1 C)$. □

Definition 4.15. Let C and D be non-trivial coarse complementary components of W and Z respectively, with both W and Z contained in $\mathcal{W}$. We say that $C \cap D$ is *small* if for some $r > 0$, $C \cap D$ is contained in $N_r(Z)$.

Lemma 4.14 ensures that this definition depends only on C and D, not the choice of W and Z. We now show that smallness is symmetric in C and D.

Proposition 4.16. *For C and D as above, $C \cap D$ is small if and only if for some $r > 0$, $C \cap D \subseteq N_r(W)$. In particular, $C \cap D$ is small if and only if $D \cap C$ is small.*

Proof. Suppose that there exists an $r' > 0$ such that $C \cap D \subseteq N_{r'}(Z)$. By the assumption that Z is almost essentially embedded, there exists a $B \geq 0$ such that $Z \subseteq N_B(D \backslash N_{r'}(Z))$. Thus $C \cap D \subseteq N_{B+r'}(D \backslash N_{r'}(Z))$. Noting that $D \backslash N_{r'}(Z) \subseteq D \cap C^*$, we deduce that $C \cap D \subseteq N_{B+r'}(\partial_1 C) \subseteq N_r(W)$ for some $r > 0$. The other direction follows by symmetry. □

Lemma 4.17. *Suppose C and D are non-trivial coarse complementary components that are coarsely equivalent to C' and D' respectively. Then $C \cap D$ is small if and only if $C' \cap D'$ is small.*

Proof. It is sufficient to show that $C \cap D$ is small if and only if for every $r, r' > 0$, $N_r(C) \cap N_{r'}(D)$ is small. By Lemma 3.10, we see that $C \cap D$ is small if and only if $C \cap N_{r'}(D)$ is small. Thus by reversing the roles of C and D, Proposition 4.16 and a second application of Lemma 3.10, we deduce that $C \cap N_{r'}(D)$ is small if and only if $N_r(C) \cap N_{r'}(D)$ is small. □

The following generalises Definition 2.2 of [Sco98], which defines crossing of almost invariant sets. Proposition 4.16 ensures the following is symmetric:

Definition 4.18 ([SS00, Definition 2.2])**.** Let C and D be coarse complementary components. Then we say that C *crosses* D if none of the four sets $C \cap D$, $C^* \cap D$, $C \cap D^*$ or $C^* \cap D^*$ is small.

There is another notion of crossing:

Definition 4.19 ([SS00, Definition 3.1])**.** Let C and D be coarse complementary components of W and Z respectively. Then we say that C *strongly crosses* D if for every $r \geq 0$, neither $W \cap D$ nor $W \cap D^*$ are contained in $N_r(Z)$. If C crosses D but does not strongly cross D we say that C *weakly crosses* D.

We will shortly prove that as the nomenclature suggests, strong crossing is indeed a stronger property than crossing.

Proposition 4.20. *If C and D are coarse complementary components that strongly cross, then they cross.*

Suppose that D is a non-trivial coarse complementary component of $Z \in \mathcal{W}$. Since Z is almost essentially embedded, it follows that D is coarsely equivalent to $D \setminus N_r(Z)$ for every $r \geq 0$. We use this to deduce the following:

Lemma 4.21. *Suppose that C and D are non-trivial coarse complementary components of elements of $\mathcal{W}$ and that $C \cap D^*$ is small. Then there exists a coarse complementary component D' that is coarsely equivalent to D such that $C \subseteq D'$.*

Proof of Proposition 4.20. Suppose C and D are strongly crossing non-trivial coarse complementary components of $W, Z \in \mathcal{W}$ respectively. We assume for contradiction that C and D do not cross. Thus one of the four sets in Definition 4.18, say $C \cap D^*$ without loss of generality, must be small. Then by Lemma 4.21, $C \subseteq N_r(D)$ for some $r \geq 0$. By Lemmas 3.10 and 4.14, $W \subseteq D \cup N_{r'}(Z)$ for some $r' \geq 0$, and so $W \cap D^* \subseteq N_{r'}(Z)$, contradicting the fact that W and Z strongly cross. □

We note that in the generality considered considered in [SS00], strong crossing is *not* symmetric. However, the almost essentially embedded hypothesis ensures that it is symmetric in the present setting.

Proposition 4.22. *Let C and D be non-trivial coarse complementary components. If C strongly crosses D, then D strongly crosses C.*

Proof. Suppose C and D are non-trivial coarse complementary components of $W, Z \in \mathcal{W}$ respectively. If C strongly crosses D, then C crosses D, and hence Proposition 4.16 tells us that D crosses C. We suppose that D does not strongly cross C. Then there exists an $r > 0$ such that at least one of $Z \cap C$ or $Z \cap C^*$ is contained in $N_r(W)$.

We may suppose without loss of generality that $Z \cap C$ is contained in $N_r(W)$, thus by Lemmas 3.10 and 4.14, we deduce $C \cap \partial_1 D \subseteq N_{r_2}(W)$ for some $r_2 > 0$. Then

$$\partial_1(C \cap D) \subseteq \partial_1 C \cup (C \cap \partial_1 D) \subseteq N_R(W)$$

for some $R > 0$. Thus $C \cap D$ is a coarse complementary component of W. Since C crosses D, $C \cap D$ is not small, hence is deep. As W is almost essentially embedded, Lemma 3.10 ensures that

$$W \subseteq N_B(C \cap D) \subseteq N_B(D) \subseteq D \cup N_{B'}(Z)$$

for some $B, B' > 0$. Thus $W \cap D^* \subseteq N_{B'}(Z)$, contradicting our assumption that C strongly crosses D. □

If G splits over H, we may associate an H-almost invariant subset C to this splitting as follows: let T be the Bass-Serre associated to the splitting, let $e \in ET$ be an edge with stabilizer H, and let T^+ be a component of $T \backslash e$. Fix a vertex $v \in VT$. Then $C := \{g \in G | gv \in T^+\}$ is a non-trivial coarse complementary component of H, i.e. an H-almost invariant subset. We say that C is the *coarse complementary component associated to the given splitting* of G over H.

Lemma 4.23 ([SS03, Lemma 2.27])**.** *Suppose G is a finitely generated group that splits over the subgroups H_1 and H_2. Let T_1 and T_2 be the associated Bass-Serre trees, and let C_1 and C_2 be the coarse complementary components associated to these splittings. We say that T_1* is hyperbolic with respect to *T_2 if H_1 doesn't fix a vertex of T_2. Then C_1 strongly crosses C_2 if and only if T_1 is hyperbolic with respect to T_2.*

The following proposition characterises weak crossing:

Proposition 4.24. *Let C and D be non-trivial coarse complementary components of W and Z respectively. If C crosses D but does not strongly cross D, then W and Z are at finite Hausdorff distance.*

Proof. Since C does not strongly cross D, Proposition 4.22 tells us that D cannot strongly cross C. Replacing C and D by their complements if necessary, we thus choose $r > 0$ such that $W \cap D \subseteq N_r(Z)$ and $Z \cap C \subseteq N_r(W)$. Then Lemma 4.14 ensures that

$$\partial_1(C \cap D) \subseteq (\partial_1 C \cap D) \cup (C \cap \partial_1 D) \subseteq N_R(W) \cap N_R(Z)$$

for some $R > 0$. By our assumption that C crosses D, we deduce that $C \cap D$ is a non-trivial coarse complementary component of both W and Z. Thus by Lemma 4.14, W and Z are at finite Hausdorff distance. □

We now construct our partial order. Let $\mathcal{B}$ denote the coarse equivalence classes of non-trivial coarse complementary components in G. Given $[C], [D] \in \mathcal{B}$ we write $[C] \leq [D]$ if and only if $C \cap D^*$ is small. By Lemma 4.17, this is independent of the choice of $C \in [C]$ and $D \in [D]$. Lemma 4.21 tells us that $[C] \leq [D]$ if and only if for every $C' \in [C]$, there exists a $D' \in [D]$ such that $C' \subseteq D'$.

Proposition 4.25. $(\mathcal{B}, \leq)$ *is a partial order.*

Proof. It is clear that $\leq$ is reflexive. Suppose that $[C] \leq [D]$ and $[D] \leq [C]$. By Lemma 4.21, there exists a $D' \in [D]$ such that $C \subseteq D'$, and a $C' \in [C]$ such that $D' \subseteq C'$. Since $C \subseteq D' \subseteq C'$ and $[C'] = [C]$, it follows that $[C] = [D]$. Thus $\leq$ is symmetric. Now suppose $[C] \leq [D]$ and $[D] \leq [E]$. By Lemma 4.21, we may choose $D' \in [D]$ and $E' \in [E]$ such that $C \subseteq D' \subseteq E'$; thus $C \cap (E')^* = \emptyset$ so that $[C] \leq [E'] = [E]$. Thus $\leq$ is transitive. □

Unlike in [SS00] and [SS03], we no longer have to worry about almost invariant sets being in *good position*. This is because if the following fact established in the preceding proof: if C and D are non-trivial coarse complementary components such that two of the four sets $C \cap D$, $C^* \cap D$, $C \cap D^*$ and $C^* \cap D^*$ are small, then either $[C] = [D]$ or $[C^*] = [D]$. We write $[C] < [D]$ if $[C] \leq [D]$ and $[C] \neq [D]$.

Lemma 4.26 ([SS03, Lemma 3.6])**.** *Suppose* $[C], [D], [E] \in \mathcal{B}$*,* $[C] < [D]$*, and that* E *crosses* C *but does not cross* D*. Then either* $[E] < [D]$ *or* $[E^*] < [D]$*.*

Proof. As $[C] < [D]$, we use Lemma 4.21 to deduce that there exists a $D' \in [D]$ such that $C \subseteq D'$. As E crosses C and $E \cap C \subseteq E \cap D'$, we deduce $E \cap D'$, and hence $E \cap D$, cannot be small. Similarly, $E^* \cap D$ cannot be small. However, as E and D do not cross, at least one of $E \cap D^*$ or $E^* \cap D^*$ is small, and so either $[E] \leq [D]$ or $[E^*] \leq [D]$. As $E \cap D$ and $E^* \cap D$ are not small, $[D]$ cannot be equal to either $[E]$ or $[E^*]$. □

We now fix a collection X of non-trivial coarse complementary components of elements of $\mathcal{W}$. For each $A \in X$, we define $\overline{[A]} := \{[A], [A^*]\}$. We let $\overline{X} := \{\overline{[A]} \mid A \in X\}$. For ease of notation, we will refer to $\overline{[A]}$ simply as $[A]$ when unambiguous.

Definition 4.27. Let $R, S, T \in \overline{X}$. We say that S *is between* R *and* T if there exist $[A] \in R$, $[B] \in S$ and $[C] \in T$ such that $[A] < [B] < [C]$. If S is between R and T, we write RST.

We now define an equivalence relation $\sim$ on $\overline{X}$ generated by the relation $\overline{[A]} \sim \overline{[B]}$ if A crosses B. We call the equivalence classes *cross-connected components*. Let $P(X)$ denote the set of all cross-connected components of $\overline{X}$.

Definition 4.28. A cross-connected component L is called a *strong cross-connected component* if it contains two pairs $\overline{[C]}$ and $\overline{[D]}$ such that C strongly crosses D. We say that L is *isolated* if L contains a single element $\overline{[C]}$. Otherwise, we say that L is a *weak cross-connected component*.

We now classify strong and weak coarse complementary components.

Proposition 4.29. *Suppose L is a cross connected component. Then L is strong if and only if for every $\overline{[C]} \in L$, there is some $\overline{[D]} \in L$ such that C strongly crosses D. If L is weak, then for any $\overline{[C]}, \overline{[D]} \in L$, $\partial_1 C$ and $\partial_1 D$ are at finite Hausdorff distance.*

Proof. Suppose L is strong and $\overline{[C]} \in L$. As L is strong, there exists $\overline{[D]}, \overline{[E]} \in L$ such that D strongly crosses E. By the definition of cross-connected components, there exists a sequence $\overline{[C]} = \overline{[C_0]}, \dots, \overline{[C_n]} = \overline{[D]}, \overline{[C_{n+1}]} = \overline{[E]}$ such that for each i, C_i crosses each C_{i+1}. We pick the minimal i such that C_i strongly crosses C_{i+1}, noting that such an i exists as D strongly crosses E. By Proposition 4.24, $\partial_1 C$ has finite Hausdorff distance from $\partial_1 C_i$. Thus as C_{i+1} strongly crosses C_i, it strongly crosses C. It follows that for every $\overline{[C]} \in L$, there is some $\overline{[D]} \in L$ such that C strongly crosses D. The rest of the statement follows from Proposition 4.24. □

Definition 4.30. Let $L, M, N \in P(X)$ be distinct cross-connected components. We say that M *is between* L *and* N if for some $R \in L$, $S \in M$ and $T \in N$, S is between R and T. If M is between L and N, we write LMN.

The following proposition shows that the above betweenness notion is independent of the choice of R and T.

Proposition 4.31. *Suppose that L, M, N are distinct elements of $P(X)$. Let $R, R' \in L$, $S \in M$ and $T, T' \in N$. If RST, then $R'ST'$.*

Proof. It suffices to show RST implies $R'ST$ when R crosses R'. If RST, there exist $\overline{[A]} \in R$, $\overline{[B]} \in S$ and $\overline{[C]} \in T$ such that $[A] < [B] < [C]$. Suppose $[A'] \in R'$. Then by Lemma 4.26, either $[A'] < [B] < [C]$ or $[(A')^*] < [B] < [C]$; in either case, we deduce that $R'ST$. □

In [War81], Ward described axioms for a pretree — a ternary relation that generalises the betweenness relation of the vertex set of a tree. Pretrees have played important roles in the work of Bowditch [Bow98, Bow02], Scott–Swarup [SS03] and Papasoglu–Swenson [PS06, PS09].

Definition 4.32. Let $\mathcal{P}$ be a set endowed with a ternary relation, which we denote xyz for $x, y, z \in \mathcal{P}$. We say that $\mathcal{P}$ is a *pretree* if it satisfies the following axioms:

(T0) If the relation xyz holds, then $x \neq z$.

(T1) The relation xyz holds if and only if zyx does.

(T2) The relation xyz and xzy cannot both hold.

(T3) If the relation xyz holds and $w \neq y$, then either wyz or xyw.

Proposition 4.33 ([SS03, Theorem 3.8])**.** *The set of cross-connected components $P(X)$, equipped with the betweenness relation described above, is a pretree.*

Proof. Axioms (T0) and (T1) are trivial. We now show (T2). Suppose for contradiction that LMN and LNM hold for some distinct $L, M, N \in P(X)$. Using Proposition 4.31, there exist $R \in L$, $S \in M$ and $T \in N$ such that RST and RTS. We may therefore choose $[A] \in R$, $[B] \in S$ and $[C] \in T$ such that $[A] < [B] < [C]$. By checking all possibilities and using Proposition 4.25, we conclude that RTS cannot hold; thus (T2) holds.

Now suppose LMN for distinct $L, M, N \in P(X)$ and that $O \neq M$. We choose $R \in L$, $S \in M$ and $T \in N$ such that RST, and choose $[A] \in R$, $[B] \in S$ and $[C] \in T$ such that $[A] < [B] < [C]$. Let $[D] \in Q \in O$. Since D does not cross B, one of the four sets $D \cap B$, $D^* \cap B$ $D \cap B^*$ and $D^* \cap B^*$ is small. If $D \cap B$ is small, then $[A] < [B] < [D^*]$, so that RSQ and therefore LMO. We argue similarly when one of $D^* \cap B$, $D \cap B^*$ or $D^* \cap B^*$ is small. This proves (T3). □

More information on pretrees can be found in [AN98], [Bow98] and [Bow99]. In particular, every finite subset of a pretree can be embedded as a finite subset of a simplicial tree (with the same betweenness relation).

Definition 4.34. Let $\mathcal{P}$ be a pretree. A *star* in $\mathcal{P}$ is a subset $\omega \subseteq \mathcal{P}$ such that:

(1) $|\omega| \geq 2$;

(2) for any $x, z \in \omega$ and $y \in \mathcal{P}$, the relation xyz doesn't hold;

(3) ω is maximal with respect to (1) and (2).

Let $\Omega(\mathcal{P})$ denote the collection of stars of $\mathcal{P}$. A pretree is *discrete* if for any $x, y \in \mathcal{P}$, there are only finitely many $z \in \mathcal{P}$ such that the relation xzy holds.

If a pretree $\mathcal{P}$ is discrete, then $\mathcal{P}$ can naturally be embedded in a bipartite simplicial tree T with vertex set $\mathcal{P} \sqcup \Omega(\mathcal{P})$. We say a vertex $v \in VT$ is a V_0-*vertex* if it corresponds to an element of $\mathcal{P}$, and a V_1-*vertex* if it is a star. See [Bow98].

The sort of situation one might wish to investigate is as follows. Let G be a one-ended finitely generated group that does not split over a two-ended subgroup. Let X be a collection of non-trivial coarse complementary components C such that $\partial_1 C$ is a coarsely embedded plane (either Euclidean or hyperbolic). Can the pretree $P(X)$ be embedded in a simplicial tree, or perhaps a real tree? Can this be used to determine whether G splits over a planar group? We provide a partial resolution to these questions in Section 4.3.

We now relate this pretree construction to various JSJ theories that appear in the literature.

4.2.1 Groups with infinitely many ends

Let G be a finitely generated group with infinitely many ends, equipped with the word metric for some choice of finite generating set. We define

$$k_0 := \min\{|\partial_1 C| \mid C \subseteq G \text{ and } C^* \text{ are infinite}\}.$$

As G has infinitely many ends, $k_0 < \infty$. We say that $C \subseteq G$ is a k-*cut* if C and C^* are infinite and $|\partial_1 C| = k$. We observe that every k-cut is a non-trivial coarse complementary component of any finite subset of G. Furthermore, $\partial_1 C$ trivially satisfies the almost-essentially embedded condition. A *minimal end cut* is a k_0-cut. (This differs from the definition of cuts in [EP14].)

Let X be the set of minimal end cuts. It is evident that G acts naturally on X. It is shown in [EP14] that $P(X)$ is a discrete pretree and that G acts without fixed points on $P(X)$. Using this, one can thus deduce Stallings' theorem. Crucial to this is the observation that coarse equivalence classes of minimal end cuts in a cross-connected component have a circular structure, so that the simplicial tree into which $P(X)$ embeds is called a *cactus tree*.

This tree is not invariant under a change in generating set of G. To see this, we consider the Cayley graph of the free group F_2 with respect to a free basis. It is easy to see that the cactus tree is not invariant under $\mathrm{Out}(F_2)$, thus is not an invariant under

a change of generating set. In particular, the cactus tree is certainly not a geometric invariant of the group. See also Lemma 6.9 of [SS03], which says more generally that one cannot define canonical almost invariant sets of a group with infinitely many ends.

4.2.2 Bowditch's JSJ tree

Let G be a one-ended hyperbolic group and let Γ be its Cayley graph. Let $\partial\Gamma$ be the Gromov boundary of Γ. Using the cut point structure of $\partial\Gamma$, Bowditch defines the JSJ tree of G. We explain how to recover this tree using the preceding construction. This can be done using the double-dagger condition of Bestvina-Mess [BM91], and is made explicit using work of Barrett [Bar18].

We consider two-ended subspaces of a hyperbolic group that coarsely separate. By [Gro87, Corollary 8.1.D], every 2-ended subgroup of a hyperbolic group is quasi-convex, thus is at finite Hausdorff distance from a geodesic. Therefore, in order to examine splittings over 2-ended subgroups, it is sufficient to investigate geodesics in Γ that coarsely separate the group.

Proposition 4.35 ([Bar18, Proposition 1.22])**.** *Let $\gamma \subseteq \Gamma$ be a geodesic, and let $\Lambda\gamma \subseteq \partial\Gamma$ be the limit set of Γ. Let $\mathcal{C}_{r,K}(\gamma)$ denote the set of components of $\Gamma \backslash N_r(\gamma)$ that are not contained in $N_K(\gamma)$. For suitable $K \gg r \gg 0$, there is a 1-1 correspondence between $\mathcal{C}_{r,K}(\gamma)$ and components of $\partial\Gamma \setminus \Lambda\gamma$. In particular, a geodesic in Γ coarsely separates if and only if its limit set is a cut pair of $\partial\Gamma$.*

It is worth emphasising that Proposition 4.35 uses the double-dagger condition, so depends crucially on the fact that ∂G has no global cut points [Swa96], which is a highly non-trivial theorem.

Let X be the set of non-trivial coarse complementary components of geodesics that coarsely separate Γ. We consider the pretree $P(X)$. Note that X, and hence the pretree $P(X)$, is invariant under self quasi-isometries of G. Using Proposition 4.35, it is not hard to verify that $P(X)$ is isomorphic to the pretree described in [Bow98]. In particular, $P(X)$ is discrete so embeds in the *Bowditch JSJ tree* $T(X)$. Indeed, we recall that in [Bow98], the set of local cut points of $\partial\Gamma$ is partitioned into three sets: infinite $\sim$-classes, $\sim$-pairs and $\approx$-pairs. We do not explain this notation here, but refer the reader to [Bow98]. Strong cross-connected components in $P(X)$ correspond to infinite $\sim$-classes, whilst weak and isolated cross-connected components in $P(X)$ correspond to either $\sim$-pairs and $\approx$-pairs. The betweenness relation of $P(X)$ is the same as the betweenness relation used in [Bow98].

4.2.3 The regular neighbourhood of Scott-Swarup

We outline the main construction of Scott-Swarup [SS03]. Let G be a finitely generated group. By Proposition 3.14, if $C \subseteq G$ is the coarse complementary component of a finitely generated subgroup $H \leq G$, it is an H-almost invariant subset. Restricting to almost invariant subsets rather then general coarse complementary components, the preceding pretree construction is the same as the regular neighbourhood construction of [SS03].

Recall a subgroup is said to be VPC-n if it has a finite index subgroup that is polycyclic of Hirsch length n. Under the assumption that G does not split over any VPC-k group for $k < n$, Proposition 3.32 tells us that every VPC-n subgroup of G is almost essentially embedded.

Theorem 4.36 ([SS03, Theorem 12.3],[GL10, Theorem 4.1 and 4.2])**.** *Fix $n \geq 1$. Let G be a one-ended, finitely presented group that is not VPC-$(n+1)$ and does not split over any VPC-k group for $k < n$. Let $\mathcal{H}$ be the collection of VPC-n subgroups of G. Let X be the collection of all non-trivial H-almost invariant subsets of G for all $H \in \mathcal{H}$. Then the pretree $P(X)$ is discrete and so embeds in a simplicial tree $T(X)$. Moreover, V_0-vertices of $T(X)$ — those that correspond to a cross-connected components of $P(X)$ — can be described as follows.*

Suppose that G_v is the vertex stabilizer of a V_0-vertex corresponding to a cross connected component containing $\overline{[C]}$, where $C \in X$ is an H-almost invariant subset for some $H \in \mathcal{H}$. Then either:

(1) No $[D] \in X$ crosses $[C]$, in which cases G_v is commensurable to H. We say that v is an isolated vertex.

(2) Some $[D] \in X$ weakly crosses $[C]$, but no $[E] \in X$ strongly crosses $[C]$. Then $G_v = \mathrm{Comm}_G(H)$ and we say that v is of commensurator type. We say that H has small *commensurator if H has finite index in $\mathrm{Comm}_G(H)$; we then say that v is of small commensurator type.*

(3) Some $[D] \in X$ strongly crosses $[C]$. Then incident edge stabilizers of G_v are VPC-n, and G_v is quadratically hanging *as defined in Definition 1.4.*

We call $T(X)$ the regular neighbourhood of $\mathcal{H}$.

The tree $T(X)$ is $Aut(G)$ invariant, as the collection of almost invariant sets and their crossing patterns are $Aut(G)$ invariant. The tree $T(X)$ satisfies the properties described in Theorem 12.5 of [SS03]. Although edge stabilizers of $T(X)$ are not

necessarily VPC-n, they will be if all vertices of commensurator type are small. One should bear in mind that vertices of large commensurator type and their adjacent edges are not necessarily finitely generated; see Example 11.1 of [SS03].

4.2.4 Coarse separation by quasi-lines

Let G be a one-ended finitely presented group. A *quasi-line* is defined to be a coarsely embedded copy of $\mathbb{R}$. Consider the set X of all non-trivial coarse complementary components of almost essentially embedded quasi-lines in G. We fix a JSJ decomposition $\mathcal{G}$ of G with two-ended edge groups. Using [Pap05], we can describe cross-connected components of $\overline{X}$ as follows:

(1) for each $C \in X$, C is not strongly crossed by any other coarse complementary component in X if and only if $\partial_1 C$ is finite Hausdorff distance from the conjugate of an edge group of $\mathcal{G}$;

(2) if $C, D \in X$ and C strongly crosses D, then both $\partial_1 C$ and $\partial_1 D$ are coarsely contained in the conjugate of a quadratically hanging vertex group of $\mathcal{G}$;

(3) for every quadratically hanging vertex group G_v of $\mathcal{G}$, there exist $C, D \in X$ such that C strongly crosses D and both $\partial_1 C$ and $\partial_1 D$ are coarsely contained in G_v.

Using this description, we observe that strong cross-connected components correspond to conjugates of quadratically hanging vertex groups of $\mathcal{G}$. If $[C]$ is contained in a weak cross-connected component, then $\partial_1 C$ is finite Hausdorff distance from the conjugate of an edge group of $\mathcal{G}$.

The above description of crossing ensures that $P(X)$ is isomorphic to the underlying pretree in the Scott-Swarup regular neighbourhood. Indeed, even though there are quasi-lines that coarsely separate and don't lie within finite Hausdorff distance of a two-ended subgroup, all such quasi-lines are coarsely contained in the conjugate of a quadratically hanging vertex group.

4.3 The regular neighbourhood

To state our results, we first generalise the notion of *enclosing* from [SS03].

Definition 4.37. Let T be a G-tree (not necessarily minimal). We fix a *basepoint* $v \in VT$. Let $e \in ET$ be an oriented edge of T and let T_e^+ be the component of $T \backslash \mathring{e}$ containing the terminal vertex of e, where $\mathring{e}$ is the open edge corresponding to e. Let

$C_e := \{g \in G \mid gv \in T_e^+\}$. We say that $C \in X$ is *enclosed* by a vertex $v \in VT$ if for every $e \in ET$ with terminal vertex v, either $C_e^* \cap C$ or $C_e^* \cap C^*$ is small.

It is shown in Lemma 2.10 of [SS03] that the coarse equivalence class $[C_e]$ is independent of the choice of v.

Remark 4.38. We remark that G_e is not assumed to be almost essentially embedded, so one has to make sure that the smallness required in the above definition is well-defined. We observe that since $G_e C_e = C_e$, we have that C_e is almost essential whenever it is deep. Indeed, if C_e is deep, then for any $A \geq 0$ there exists an $x \in C_e \backslash N_A(G_e)$, such that $G_e x \subseteq C_e \backslash N_A(G_e)$, and therefore $G_e \subseteq N_B(C_e \backslash N_A(G_e))$. Thus the statement "$C_e^* \cap C$ is small" is well-defined and symmetric. Note however, that strong crossing is not necessarily symmetric in such a situation.

We now begin our construction of the tree in Theorem 4.1. We fix some group G of type $FP_{n+1}^{\mathbb{Z}_2}$ and a uniformly n-acyclic $\mathbb{Z}_2$-metric complex $(G, C_\bullet)$. Let $\mathcal{W} := \cup_{i=1}^m \cup_{g \in G} \{gW_i\}$, where $W_1, \ldots, W_m$ is a collection of essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ subspaces that coarsely separate G. Using techniques of Chapter 3, we are thus able to choose constants A, D, s that satisfy the following properties.

(1) There exists an $A \geq 0$ such that each $W \in \mathcal{W}$ has two essential, coarse disjoint $(1, A)$-coarse complementary components.

(2) For every $W \in \mathcal{W}$ that coarsely 3-separates G, W has three essential, coarse disjoint $(1, A)$-coarse complementary components.

(3) Then there is a $D \geq 0$ such that the following holds. Suppose C is a non-trivial $(1, A)$-coarse complementary component of $W \in \mathcal{W}$. Let $[\alpha_C] \in \hat{H}_c^{n+1}(C_\bullet)$ be the nonzero cohomology class corresponding to the image of δ in the Mayer-Vietoris sequence

$$\cdots \to H^n_{\text{coarse}}(C \cup W) \oplus H^n_{\text{coarse}}(C^* \cup W) \to H^n_{\text{coarse}}(W) \xrightarrow{\delta} H^{n+1}_{\text{coarse}}(G).$$

Then for every $x \in W$, there is an $[\alpha] = [\alpha_C]$ such that $\text{supp}(\alpha) \subseteq N_D(x)$. [1]

(4) **The crossing criterion:** There is an $s \geq A$ such that the following holds. Suppose that C is a non-trivial $(1, A)$-coarse complementary component of some $W \in \mathcal{W}$. Then for any $s' \geq s$, $[\omega] \in \hat{H}_c^{n+1}(C_\bullet)$ can be represented by two

[1] We note that as W is a coarse $PD_n^{\mathbb{Z}_2}$ space, $H^n_{\text{coarse}}(W) \cong \mathbb{Z}_2$. Moreover, since $C \cup W$ and $C^* \cup W$ are both essential, the map $H^n_{\text{coarse}}(W) \xrightarrow{\delta} H^{n+1}_{\text{coarse}}(G)$ is injective, so its image does contain a unique non-zero cohomology class $[\alpha_C]$.

cocycles supported in $C\backslash N_{s'}(W)$ and $C^*\backslash N_{s'}(W)$ respectively if and only if $[\omega] = [\alpha_C]$.

Moreover, if C_1, C_2 and C_3 are coarse disjoint non-trivial $(1, A)$-coarse complementary component of some $W \in \mathcal{W}$, and $[\omega]$ can be represented by two cocycles supported in $C_1\backslash N_s(W)$ and $C_2\backslash N_s(W)$ respectively, then $[\omega] = 0$. (See Part (2) of Proposition 3.19 and Lemma 3.35.)

Remark 4.39. Throughout this section, we use the term *crossing criterion*, to describe the application of Part (2) of Proposition 3.19 and Lemma 3.35 as stated above in (4).

Definition 4.40. We say a collection $\mathcal{W}$ of essentially embedded coarse $PD_n^{\mathbb{Z}_2}$ subspaces is a *uniform $PD_n^{\mathbb{Z}_2}$-set* if it is G-invariant, and conditions (1)–(4) hold as above for some constants A, D and s.

As noted above, if $\mathcal{W} := \cup_{i=1}^m \cup_{g\in G} \{gW_i\}$, where $W_1, \ldots, W_m$ are as in Theorem 4.1, then $\mathcal{W}$ is a uniform $PD_n^{\mathbb{Z}_2}$-set. Our main motivation for working with more general uniform $PD_n^{\mathbb{Z}_2}$-sets that do not a priori arise in this way, is the following:

Proposition 4.41. *If $\mathcal{W}$ is a uniform $PD_n^{\mathbb{Z}_2}$-set in G and $f : G \to G'$ is a quasi-isometry, then $\mathcal{W}' := \{g'f(W) \mid W \in \mathcal{W}, g' \in G'\}$ is a uniform $PD_n^{\mathbb{Z}_2}$-set in G'.*

We now state the main theorem of this section. Theorem 4.1 can be easily deduced as a special case.

Theorem 4.42. *Let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and suppose $\mathcal{W}$ is a non-empty uniform $PD_n^{\mathbb{Z}_2}$-set. Then there is a minimal G-tree $T = T(\mathcal{W})$, bipartite on $VT = V_0 \sqcup V_1$, such that:*

(1) Every non-trivial coarse complementary component of every $W \in \mathcal{W}$ is enclosed by a V_0-vertex of T. Conversely, every V_0-vertex of T encloses such a coarse complementary component.

(2) Suppose that G_v is the vertex stabilizer of a V_0-vertex v . Then either:

(a) G_v is at finite Hausdorff distance from some $W \in \mathcal{W}$;

(b) $G_v = \mathrm{Comm}_G(H)$, where $H \leq G$ is a subgroup at finite Hausdorff distance from some $W \in \mathcal{W}$;

(c) $G_v = \mathrm{Stab}([\alpha])$ for some non-zero coarse cohomology class $[\alpha] \in H^{n+1}_{\mathrm{coarse}}(G)$.

Moreover, if T consists of a single vertex, then it is a V_0-vertex of type (b) or (c).

4.3.1 Construction of the regular neighbourhood

To prove Theorem 4.42, we fix constants A, D and s as in Definition 4.40. Let X denote the set of non-trivial $(1, A)$-coarse complementary components of all the $W \in \mathcal{W}$. For each $C \in X$, we let W_C denote an element of $\mathcal{W}$ such that C is a $(1, A)$-coarse complementary component of W_C. Although W_C might not be unique, Lemma 4.14 ensures the coarse equivalence class $[W_C]$ is well-defined. We say that $W \in \mathcal{W}$ *bisects* G if it coarsely 2-separates G, but does not coarsely 3-separate G.

For each $C \in X$, we let $[\alpha_C] \in \hat{H}_c^{n+1}(C_\bullet)$ be the coarse cohomology class defined as in (3) above. Note that $[\alpha_C] = [\alpha_{C^*}]$, as we are working with $\mathbb{Z}_2$ coefficients. Proposition 3.19 tells us that $[\alpha_C] = [\alpha_{C'}]$ if $[C']$ is equal to either $[C]$ or $[C^*]$. We use these cocycles to investigate the pretree $P(X)$. We recall the following useful fact.

Proposition 4.43 ([KK05, Part 4 of Theorem 7.7])**.** *If W and W' are coarse $PD_n^{\mathbb{Z}_2}$ spaces and $W \subseteq N_R(W')$ for some $R \geq 0$, then there exists $R' \geq 0$ such that $W' \subseteq N_{R'}(W)$.*

We first investigate stabilizers of cross-connected components of $\overline{X}$ containing C, where W_C coarsely 3-separates G.

Proposition 4.44. *Suppose $W \in \mathcal{W}$ coarsely 3-separates G and that $C \in X$ is a non-trivial $(1, A)$-coarse complementary component of W. By Lemma 3.40, we know that there is a coarse $PD_n^{\mathbb{Z}_2}$ subgroup $H \leq G$ at finite Hausdorff distance from W. Let $L \in P(X)$ be the cross connected component containing C. Then $H_0 \leq \mathrm{Stab}_G(L) \leq \mathrm{Comm}_G(H)$, where H_0 is a finite index subgroup of H.*

Proof. We choose B large enough so that C is a $(1, B)$-coarse complementary component of H. We may write C as a disjoint union $\sqcup_{i \in I} C_i$ of irreducible $(1, B)$-coarse complementary components of H. By Lemma 3.12, there is a $B' \geq B$ such that every shallow C_i is contained in $N_{B'}(H)$. For each deep C_i, let $H_i := \mathrm{stab}_G(C_i \backslash N_B(H))$, which is a finite index subgroup of H by Corollary 3.18. By Proposition 3.30 and the fact that G is locally finite, there are only finitely many C_i that are deep. Hence the intersection $H_0 := \bigcap_{C_i \text{ is deep}} H_i$ is a finite index subgroup of H satisfying $[hC] = [C]$ for all $h \in H_0$. Thus $\mathrm{Stab}_G(L) \geq H_0$.

The crossing criterion ensures C cannot strongly cross any other element of X. Thus if C and gC are contained in the same cross-connected component, then Proposition 4.29 tells us that W and gW are at finite Hausdorff distance, and hence $g \in \mathrm{Comm}_G(H)$ by Proposition 2.12. □

We now investigate the case where $W \in \mathcal{W}$ bisects G. When W bisects G, we let $C_W \in X_0$ denote a non-trivial coarse complementary component. The choice of C_W is unique up to coarse equivalence and interchanging C_W with C_W^*. Let L_W denote the unique cross connected component containing $\{[C_W], [C_W^*]\}$, and let $[\alpha_W] := [\alpha_{C_W}] = [\alpha_{C_W^*}]$. The following proposition is the linchpin of this section, demonstrating why stabilizers of cohomology classes appear in the statement of Theorem 4.42.

Proposition 4.45. *If $W \in \mathcal{W}$ bisects G, then* $\mathrm{Stab}(L_W) \leq \mathrm{stab}([\alpha_W])$.

Proof. Let $g \in \mathrm{Stab}(L_W)$. First suppose $[gC_W] = [C_W]$ or $[C_W^*]$. Then it is evident from the way in which α_W is defined that $[\alpha_W] = [\alpha_{gW}] = [\alpha_W g^{-1}]$, thus $g \in \mathrm{Stab}([\alpha_W])$. Now suppose that $C_{W'}$ crosses C_W, where $W' \in \mathcal{W}$ also bisects G. If the crossing were weak, Proposition 4.24 would imply that W and gW would be finite Hausdorff distance apart. But as W bisects G, we would then have $[C_{W'}] = [C_W]$ or $[C_W^*]$, which cannot occur as we have assumed that C_W and $C_{W'}$ cross. Thus $C_{W'}$ strongly crosses C_W. Recall for every $x \in W'$, there exists a cocycle α such that $[\alpha] = [\alpha_{W'}]$ and $\mathrm{supp}(\alpha) \subseteq N_D(x)$. This means $[\alpha_{W'}]$ can be represented by cocycles supported in $C_W \setminus N_s(W)$ and $C_W^* \setminus N_s(W)$ for any $s \geq 0$. Thus the crossing criterion ensures that $[\alpha_W] = [\alpha_{W'}]$. □

We explain the significance of Proposition 4.45. One of the major steps required in a fruitful JSJ theory is an understanding the ways in which hyperbolic–hyperbolic splittings can interact, i.e. a classification of what are called *flexible* vertices in [GL17]. In the case of splittings over VPC-n groups, and more generally over slender groups, these flexible vertices are quadratically hanging.

In the examples considered in Section 4.2, all cross-connected components have a 'circular' or 'necklace' arrangement, and it is this that allows one to deduce that the enclosing vertex group is quadratically hanging. This classification ensures that cross-connected components cannot be too large. For instance, if a finitely presented one-ended group isn't virtually a surface group, a cross-connected component of quasi-lines cannot be stabilized by the ambient group [Pap05]. Proposition 4.45 plays the same role when considering coarse $PD_n^{\mathbb{Z}_2}$ spaces: if G isn't a coarse $(n+1)$-manifold group as in Definition 4.2, then the group cannot fix a cross-connected component in the pretree $P(X)$.

Proposition 4.45 gives new information even in the case where one is considering almost invariant sets rather than coarse complementary components. For instance, when considering splittings of a finitely generated group over hyperbolic surface groups,

it is not known how hyperbolic–hyperbolic splittings can interact; we remark that they do not necessarily have a 'circular' arrangement, see [KM12]. Very little seems to be known about flexible vertices of JSJ decompositions over non-slender groups. Work of Markovic [Mar13] and Beeker–Lazarovich [BL18] may be useful for classifying flexible vertex groups in JSJ trees over hyperbolic surface subgroups.

We now want to obtain an actual tree rather than a pretree. As $P(X)$ itself is unlikely to always be discrete, we collapse certain vertices of $P(X)$ to form a pretree that is discrete. In a nutshell, the idea is that if W_C and $W_{C'}$ both bisect G and $[\alpha_{W_C}] = [\alpha_{W_{C'}}]$, then there is no cohomological obstruction to $[C]$ and $[C']$ lying in the same cross-connected component. We can thus form a new pretree by artificially declaring that $[C]$ and $[C']$ lie in same "cross-connected component", regardless of whether or not they actually do. This will be made precise in Definition 4.47. The following lemma demonstrates that this construction preserves the betweenness relation.

Lemma 4.46. *Suppose we have $C, C', C'' \in X$ such that $[C] < [C'] < [C'']$ and $[\alpha_C] = [\alpha_{C''}]$. Then $W_{C'}$ bisects G and $[\alpha_{C'}] = [\alpha_C]$.*

Proof. As $[C] < [C'] < [C'']$, the crossing criterion tells us that $[\alpha_C]$ can be represented by cocycles supported in $C \backslash N_r(W_C)$ and $(C'')^* \backslash N_r(W_{C''})$ for any r sufficiently large. Thus there are coarse disjoint, coarse complementary components C_1 and C_2 of $W_{C'}$, such that $[C] \leq [C_1]$, $[(C'')^*] \leq [C_2]$, and $[\alpha_C]$ can be represented by cocycles supported in $C_1 \backslash N_r(W_{C'})$ and $C_2 \backslash N_r(W_{C'})$ for any r sufficiently large. Furthermore, we can choose C_1 and C_2 such that $[C_1] = [C_2^*]$ if and only if $W_{C'}$ bisects G. Then as $[\alpha_C] \neq 0$, the crossing criterion ensures that $W_{C'}$ bisects G and $[\alpha_{C'}] = [\alpha_C]$. □

We now define the new pretree that we mentioned above. We recall that $\overline{X} := \{\overline{[C]}\} = \{\{[C], [C^*]\} \mid C \in X\}$. It will be helpful for technical reasons to identify $(1, A)$-coarse complementary that aren't necessary in the same cross-connected component of $\overline{X}$, but will eventually be in the same cross-connected component as the parameter A increases.

Definition 4.47. We define an equivalence relation $\sim$ on the set $\overline{X}$ generated by the following relations:

(1) if $\overline{[B]}, \overline{[C]} \in \overline{X}$ such that W_B and W_C both bisect G and $[\alpha_B] = [\alpha_C]$, then $\overline{[B]} \sim \overline{[C]}$;

(2) if $\overline{[B]}, \overline{[C]} \in \overline{X}$ such that $[W_B] = [W_C]$ coarsely 3-separates G, and there exist coarse complementary components $B = B_0, \dots, B_n = C$ of W_C such that each B_i crosses B_{i+1}, then $\overline{[B]} \sim \overline{[C]}$.

Let $\hat{P}(X)$ be the set of equivalence classes. Given distinct $\mathbf{L}, \mathbf{M}, \mathbf{N} \in \hat{P}(X)$, we say that $\mathbf{M}$ is between $\mathbf{L}$ and $\mathbf{N}$ if there exist $A \in \mathbf{L}, B \in \mathbf{M}$ and $C \in \mathbf{N}$ such that $[A] < [B] < [C]$. Let $\pi : X \to \hat{P}(X)$ be the map taking an element of X to the equivalence class of $\hat{P}(X)$ containing $\overline{[X]}$.

Building on Proposition 4.45 and Lemma 4.46, we now show that the betweenness relation on $\hat{P}(X)$ and $P(X)$ is the same amongst vertices that aren't collapsed.

Proposition 4.48. *(1) Suppose $B, C \in X$ and B crosses C. Then $\pi(B) = \pi(C)$.*

(2) Suppose $B, C, E \in X$ with $[B] < [C] < [E]$ and $\pi(B) = \pi(E)$. Then $\pi(C) = \pi(B)$.

Proof. (1): If B crosses C strongly, then the crossing criterion ensures that both W_B and W_C bisect G. As in the proof of Proposition 4.45, we deduce $[\alpha_B] = [\alpha_C]$, showing $\pi(B) = \pi(C)$. If B crosses C weakly, then $[W_B] = [W_C]$. If W_C bisects G, then $[B]$ is equal to $[C]$ or $[C^*]$. Otherwise, we have a sequence $B = B_0, B_1 = C$ where B_0 crosses B_1, and thus $\pi(B) = \pi(C)$.

(2): In the case where W_B bisects G, then $\pi(B) = \pi(E)$ implies that W_E bisects G and that $[\alpha_B] = [\alpha_E]$. Thus Lemma 4.46 ensures that W_C bisects G and $[\alpha_B] = [\alpha_C]$, so we deduce $\pi(B) = \pi(C)$.

If W_B coarsely 3-separates G, then $\pi(B) = \pi(E)$ implies that $[W_B] = [W_E]$ coarsely 3-separates G and there exist coarse complementary components $B = B_0, \dots, B_n = E$ of W_B such that each B_i crosses B_{i+1}. If C does not cross any B_i, then we may apply Lemma 4.26 n times to deduce that either $[B_n] < [C]$ or $[B_n^*] < [C]$; this contradicts $[C] < [E] = [B_n]$. Thus C weakly crosses some B_i, and hence $[W_C] = [W_B]$ and $\pi(B) = \pi(C)$. □

The upshot of Proposition 4.48 is that every element of $\hat{P}(X)$ is a disjoint union of cross-connected components in $P(X)$. Moreover, for any $L, N \in P(X)$ such that L and N are both contained in some $\mathbf{L} \in \hat{P}(X)$, then M is contained in $\mathbf{L}$ whenever M is between L and N. Informally, $\hat{P}(X)$ can be be obtain from $P(X)$ by collapsing convex "sub-pretrees" to points while preserving the betweenness relation.

The following lemma provides us with a quantitative bound that is essential to proving discreteness of $\hat{P}(X)$.

Lemma 4.49. *Suppose $B, C \in X$ are $(1, A)$-coarse complementary components of $W_B, W_C \in \mathcal{W}$ respectively. If $[B] < [C]$ and $\pi(B) \neq \pi(C)$, then $W_B \subseteq C \cup N_{D+s}(W_C)$.*

Proof. If $W_B \not\subseteq C \cup N_{D+s}(W_C)$, then there exists an $x \in W_B \cap (C^* \setminus N_{D+s}(W_C))$. Thus $[\alpha_B]$ can be represented by a cocycle supported in $N_D(x) \subseteq C^* \setminus N_s(W_C)$. By Lemma 4.21 and the assumption $[B] < [C]$, there is a $Q \geq 0$ such that $B \setminus N_Q(W_B) \subseteq C \setminus N_s(W_C)$. The crossing criterion thus ensures $[\alpha_B]$ can be represented by a cocycle supported in $C \setminus N_s(W_C)$, so by another application of the crossing criterion, $[\alpha_C] = [\alpha_B]$. We now show this leads to a contradiction.

This is clearly the case if W_B and W_C both bisect G, as we have assumed $\pi(B) \neq \pi(C)$. If W_B coarsely 3-separates G, then there exists a $B' \in X$ such that B and B' are coarse disjoint, $(1, A)$-coarse complementary components of W_B and $[B'] < [C]$. We repeat the argument with B' to deduce that $[\alpha_B] = [\alpha_{B'}] = [\alpha_C]$, and then use the crossing criteria at W_B to show that $[\alpha_B] = 0$, which is the desired contradiction. Lastly, we suppose W_C coarsely 3-separates G and W_B bisects G; in particular, W_B and W_C cannot be at finite Hausdorff distance. Then $x \in D$, where $D \in X$ is a $(1, A)$ coarse complementary component of W_C. Thus by the crossing criteria applied at W_C, we deduce that $[\alpha_B] = 0$, which is again a contradiction. □

Lemma 4.50. *Suppose $B, C \in X$ and $[B] < [C]$. Then there is a finite set $C_0, C_1, \ldots, C_k \in X$ with $[B] = [C_0] < [C_1] < \cdots < [C_k] = [C]$ such that for any $E \in X$ with $[B] < [E] < [C]$, at least one of the following holds:*

(1) E crosses some C_i;

(2) W_E bisects G and there is some i such that $[\alpha_{C_i}] = [\alpha_E]$ and W_{C_i} bisects G;

(3) there is some i such that $[E] = [C_i]$.

Proof. Let $x \in W_B$, $y \in W_C$, and let $x = x_0, x_1, \ldots, x_t = y$ be a 1-geodesic from x to y in G. Let $K := \{x_0, \ldots, x_t\}$. Suppose that there are $B, C, E \in X$ such that $[B] < [E] < [C]$. We claim that $W_E \cap N_{s+D}(K) \neq \emptyset$.

Indeed, we assume for contradiction that $W_E \subseteq G \setminus N_{s+D}(K)$, and so $K \subseteq G \setminus N_{s+D}(W_E)$. This implies that $P_1(K)$, which is connected, is contained in a single component of $P_1(G) \setminus P_1(N_{s+D}(W_E))$. As E is a $(1, A)$-coarse complementary component and $s + D \geq A$, Lemma 3.3 tells us that K is contained in either $E \setminus N_{s+D}(W_E)$ or $E^* \setminus N_{s+D}(W_E)$. By Lemma 4.49, we know that $W_B \subseteq E \cup N_{s+D}(W_E)$, so as $x \in K \cap W_B$, we deduce that $K \subseteq E \setminus N_{s+D}(W_E)$. Since $[C^*] < [E^*]$, Lemma 4.49 tells us that $W_C \subseteq E^* \cup N_{s+D}(W_E)$, contradicting the fact that $y \in K \cap W_C$.

Let $x \in W_E \cap N_{s+D}(K)$. There exists a cocycle α_x such that $[\alpha_x] = [\alpha_E]$ and $\mathrm{supp}(\alpha_x) \subseteq N_D(x) \subseteq N_{s+2D}(K)$. Note that $N_{s+2D}(K)$ is finite, so as in the proof of Lemma 3.36, there are only finitely many, say at most Q, cocycles with support in $N_{s+2D}(K)$.

We define $C_0, C_1, C_2, \ldots, C_k$ inductively as follows. Suppose we have $[B] = [C_0] < [C_1] < \cdots < [C_k] = [C]$ and that $[B] < [E] < [C]$. If none of (1)–(3) hold, then $[C_i] < [E] < [C_{i+1}]$ for some i, so we add E to the sequence and repeat. By Lemma 4.46, whenever we have more than three consecutive C_i that have the same cohomology class $[\alpha_{C_i}]$, we have a redundant C_i that can be deleted. After deleting these redundant vertices, we deduce that $k \leq 3Q$, and so the process of adding elements will eventually terminate. □

Corollary 4.51. *The pretree $\hat{P}(X)$ is discrete, thus embeds in a simplicial tree $T(\mathcal{W})$.*

Proof. Let $\mathbf{L}$ and $\mathbf{N}$ be distinct elements of $\hat{P}(X)$. We may choose $B, C \in X$ so that $[B] < [C]$, $\pi(B) = \mathbf{L}$ and $\pi(C) = \mathbf{N}$. Choose $B = C_0, \ldots, C_n = C \in X$ as in Lemma 4.50. Suppose $\mathbf{M}$ is between $\mathbf{L}$ and $\mathbf{N}$. Then there exist $B' \in \pi^{-1}(\mathbf{L})$, $E \in \pi^{-1}(\mathbf{M})$ and $C' \in \pi^{-1}(\mathbf{L})$ such that $[B'] < [E] < [C']$. By Proposition 4.31, $[B] < [E] < [C]$, so by Lemma 4.50, $\pi(E) = \pi(C_i)$ for some i. □

We now analyse the tree $T(\mathcal{W})$.

Lemma 4.52 (cf. Lemma 5.1 of [SS03])**.** *Let $C \in X$, and let v be the vertex of $T(\mathcal{W})$ corresponding to the element $\pi(C) \in \hat{P}(X)$. Then C is enclosed by v.*

The proof is essentially the same as that of Lemma 5.1 of [SS03], except we have to modify the argument to take account of the fact that C is only a coarse complementary component and not an almost invariant set.

Proof. We recall the definition of the set C_e in Definition 4.37. Let e be an edge with terminal vertex v, and let v be the basepoint used to define C_e. If $g \in C_e^*$, then $gv \neq v$, so that $\pi(C) \neq \pi(gC)$. Proposition 4.48 ensures that C and gC don't cross. Replacing C with its complement if necessary, we assume that either $[gC] < [C]$ or $[gC^*] < [C]$. Suppose that $h \in C_e^*$. Then v is not between gv and hv, so $\pi(C)$ is not between $\pi(gC)$ and $\pi(hC)$. Hence we have either $[hC] < [C]$ or $[hC^*] < [C]$. As this holds for every $h \in C_e^*$, Lemma 4.49 tells us that $C_e^* W_C \subseteq C \cup N_{s+D}(W_C)$. In particular, $C_e^* \subseteq C \cup N_R(W_C)$ for some sufficiently large R, and thus $C_e^* \cap C^*$ is small. □

Lemma 4.53 (Lemma 4.8 of [SS03])**.** *If a vertex v of $T(\mathcal{W})$ encloses a non-trivial coarse complementary component C, then v is contained in any minimal subtree of T.*

Note that a minimal subtree need not be unique when G fixes at least two vertices.

Proof. Let T_0 denote a minimal subtree of $T(\mathcal{W})$. We suppose for contradiction that $v \notin T_0$. We may pick an edge e of $T(\mathcal{W})$ with terminal vertex v, such that e lies on a geodesic from T_0 to v. Note that T_0 lies in the component of $T(\mathcal{W})\backslash\mathring{e}$ that contains the initial vertex of e. Picking a basepoint $w \in T_0$, we see that the resulting set C_e is empty. Thus $C_e^* = G$, so that since C is non-trivial, neither $C_e^* \cap C$ nor $C_e^* \cap C^*$ can be small. This contradicts the fact that C is enclosed by v. □

Corollary 4.54 (Lemma 5.2 of [SS03])**.** *The action of G on $T(\mathcal{W})$ is minimal.*

Proof. Let T_0 be a minimal subtree of $T(\mathcal{W})$. Lemmata 4.52 and 4.53 imply that every V_0-vertex is contained in T_0. Since $T(\mathcal{W})$ was constructed from $\hat{P}(X)$ by adding stars of size at least two, every V_1-vertex of $T(\mathcal{W})$ is adjacent to at least two V_0-vertices, so is contained in T_0. Hence $T_0 = T(\mathcal{W})$. □

Proof of Theorem 4.42. Most of this was shown above. All that remains is to analyse vertex stabilizers and verify the stated enclosing property. We first determine the vertex stabilizer of a V_0-vertex v corresponding to $\pi(C) \in \hat{P}(X)$. There are three cases to consider. If W_C bisects G, it is evident from the definition of $\hat{P}(X)$ that

$$G_v = \{g \in G \mid \pi(C) = \pi(gC)\} = \{g \in G \mid [\alpha_C] = [\alpha_{gC}]\} = \mathrm{Stab}([\alpha_C]).$$

Now suppose that W_C coarsely 3-separates G. As in Proposition 4.44, we know that there is a subgroup H_0 at finite distance from W_C such that $[hC] = [C]$ for every $h \in H_0$. We say that C is *atomic* if C cannot be written as $C = C_1 \cup C_2$, where C_1 and C_2 are deep, coarse disjoint, coarse complementary of W_C.

If C is atomic and D is a coarse complementary component of W_C, then writing $C = (C \cap D) \sqcup (C \cap D^*)$, we observe that C cannot cross D. Thus if $\pi(D) = \pi(C)$ for some $D \in X$, $[D]$ is equal to either $[C]$ or $[C^*]$. If $[gC] = [C]$, then $gW_C \cap N_{2A}(W_C) \neq \emptyset$. Indeed, if $gW_C \cap N_{2A}(W_C) = \emptyset$, then W_C is contained in a single $(1, A)$-coarse complementary component of gW_C and vice versa, hence $[gC] \neq [C]$. As W_C has finite Hausdorff distance from H_0, it it easy to see that the set of all $g \in G$ such that $gW_C \cap N_{2A}(W_C) \neq \emptyset$ is contained in finitely many double cosets of the form $H_0 g H_0$. Moreover, if $[gC] = [C]$, then $g \in \mathrm{Comm}_G(H_0)$, and hence the stabilizer of $\pi(C)$ must thus be commensurable to H_0. The same argument applies if C^* is atomic, so we now suppose neither C^* nor C is atomic.

If $\pi(C) = \pi(gC)$, then Proposition 4.29 ensures that W_C and gW_C are at finite Hausdorff distance and therefore $g \in \mathrm{Comm}_G(H_0)$. Conversely, suppose $g \in \mathrm{Comm}_G(H_0)$,

so that gC is also a coarse complementary component of W_C and neither gC nor gC^* is atomic. We may write $C = C_1 \cup C_2$, $C^* = C_3 \cup C_4$, $gC = D_1 \cup D_2$ and $gC^* = D_3 \cup D_4$, where each union consists of coarse disjoint, deep, coarse complementary components of W_C. Then $C_1 \cup D_1 \cup C_3 \cup D_3$ crosses both C and gC, thus $\pi(C) = \pi(gC)$. Hence the vertex corresponding to $\pi(C)$ has stabilizer $\mathrm{Comm}_G(H)$.

Finally, we verify the enclosing property. We have already seen in Lemma 4.52 that every non-trivial $(1, A)$-coarse complementary component of every gW_i is enclosed by a V_0-vertex of $T = T(\mathcal{W})$. Let $A' \geq A$. We now construct a new tree T' using the preceding construction applied to the set X' of all $(1, A')$-coarse complementary components of $\mathcal{W}$. By our choice of A, this tree T' is easily seen to be a subdivision of T. Indeed, the only additional V_0-vertices are those that correspond to atomic coarse complementary components of some $W \in \mathcal{W}$ that coarsely 4-separates G. By applying Lemma 4.52 to this new tree T', we see that every $(1, A')$-coarse complementary component $C \in X'$ is enclosed by some V_0-vertex of T' and is thus enclosed by some V_0-vertex of the original tree T. □

Remark 4.55. We comment briefly on the dependence of $T = T(\mathcal{W})$ on the parameter A. The final paragraph of the preceding proof demonstrates that as we increase the parameter A, we subdivide certain edges of T. We thus deduce that up to subdivision, this tree T is defined independently of A.

We conclude this section with a discussion about whether the tree $T(W_1, \dots W_m)$ in Theorem 4.1 stabilizes as we increase m. To show this, we require some sort of accessibility result. A group G is said to be *accessible* if the process of iteratively splitting along finite subgroups eventually terminates. Dunwoody showed that groups of type $FP_2^{\mathbb{Z}_2}$ are accessible [Dun85], but not every finitely generated group is accessible [Dun93]. Bestvina and Feighn showed that groups of type $FP_2^{\mathbb{Z}_2}$ are accessible over larger classes of edge groups [BF91]. We recall the following cohomological characterisation of accessibility, also due to Dunwoody.

Theorem 4.56 ([Dun79])**.** *Let G be a group and R a commutative ring. Then $H^1(G, RG)$ is finitely generated as an RG module if and only if G is accessible.*

We also mention the formulation of accessibility due to Thomassen and Woess [TW93]. They showed that a finitely generated group is accessible if and only if there exists a k such any two ends of its Cayley graph can be separated by a k-cut (as defined in Section 4.2.1). This is equivalent to requiring that $H^1_{\mathrm{coarse}}(G, \mathbb{Z}_2)$ be generated by cocycles whose support has diameter at most $D = D(k)$.

Inspired by these cohomological interpretations of accessibility, one might hope that one way to show that the tree constructed in Theorem 4.1 stabilizes as we increase m by determining if $H^{n+1}_{\text{coarse}}(G, \mathbb{Z}_2)$ is finitely generated as a $\mathbb{Z}_2 G$ module. Alternatively, one might hope to use the formulation of Dunwoody accessibility as stated in [GL17] to prove that this tree stabilizes. We will not pursue such questions in this thesis, but leave them for future investigation.

4.3.2 Quasi-isometric rigidity of the regular neighbourhood

In this section we show that the regular neighbourhood constructed in the previous section is a quasi-isometry invariant, provided the uniform $PD^{\mathbb{Z}_2}_n$-set is invariant under quasi-isometry. We recall the definition of vertex and edge spaces X_v and X_e in Section 4.1.

Theorem 4.57. *Let G and G' be groups of type $FP^{\mathbb{Z}_2}_{n+1}$. Let $f : G \to G'$ be a quasi-isometry. Suppose $\mathcal{W}$ is a uniform $PD^{\mathbb{Z}_2}_n$-set in G, $\mathcal{W}'$ is a uniform $PD^{\mathbb{Z}_2}_n$-set in G' and that there is constant $B \geq 0$ such that:*

- *for any $W \in \mathcal{W}$, there is a $W' \in \mathcal{W}'$ such that $d_{Haus}(f(W), W') \leq B$;*
- *for any $W' \in \mathcal{W}$, there is a $W \in \mathcal{W}$ such that $d_{Haus}(f(W), W') \leq B$.*

Let $T = T(\mathcal{W})$ and $T' = T(\mathcal{W}')$ be as in Theorem 4.42. Then f induces a tree isomorphism $f_* : T \to T'$ such that

(1) f_* preserves V_0-vertices and vertex type (a),(b) and (c);

(2) there is a constant $R \geq 0$ such that for every vertex or edge $r \in VT \sqcup ET$, $d_{\text{Haus}}(f(X_r), X'_{f_*(r)}) \leq R$.

To prove this, we show that the vertex and edge spaces are completely determined by the sets $\mathcal{W}$ and $\mathcal{W}'$. We fix a uniform $PD^{\mathbb{Z}_2}_n$-set $\mathcal{W}$ of G, and let A, D and s be the associated constants. Let X be the collection of all $(1, A)$-coarse complementary components of elements of $\mathcal{W}$, and let $T = T(\mathcal{W})$ denote the resulting tree. There is a map $\pi : X \to T$ that takes an element of X to the corresponding V_0-vertex of T.

We recall that we have fixed a uniformly n-acyclic $\mathbb{Z}_2$-metric complex $(G, C_\bullet)$. As in Section 3.2.2, let $Z^{n+1}_c(D)$ denote the set of modified $(n+1)$-cocycles whose support has diameter at most D. For each $(1, A)$-coarse complementary component C of some $W \in \mathcal{W}$, we have associated a cocycle $\alpha_C \in Z^{n+1}_c(D)$. Since the action of G on both T and $Z^{n+1}_c(D)$ is cocompact, there exist finitely many $C_1, \ldots, C_k$ such that

(1) each C_i is a $(1, A)$-coarse complementary component of some $W_i \in \mathcal{W}$;

(2) for every $(1, A)$-coarse complementary component C of some $W \in \mathcal{W}$, there exists an i and $g \in G$ such that $\pi(C) = \pi(gC_i)$ and $[\alpha_{gC_i}] = [\alpha_C]$.

For each i, let $v_i := \pi(C_i)$ be the corresponding V_0-vertex of T, and let $H_i := \mathrm{Stab}_G(v_i)$. Note that every V_0-vertex of T is in the G-orbit of some v_i.

Lemma 4.58. *For each V_0-vertex v of T, we define $Y_v := \cup\{W_C \mid \pi(C) = v\}$. There is a $B \geq 0$, such that for every V_0-vertex $v \in VT$, $d_{\mathrm{Haus}}(X_v, Y_v) \leq B$.*

Proof. We may write $v = gv_i$ for some $g \in G$ and v_i. Note that $Y_{gv_i} = gY_{v_i}$, so it is sufficient to only consider vertices of the form $v = v_i$ by Proposition 4.6. We need only show that each H_i has finite Hausdorff distance from Y_{v_i}. We first consider the case where W_i bisects G, so that $H_i = \mathrm{Stab}([\alpha_{C_i}])$. The definition of the pretree ensures that $[\alpha_C] = [\alpha_{C_i}]$ for every C with $\pi(C) = v_i$. Note that for every $x \in W_C$, there exists a cocycle α supported in $N_D(x)$ with $[\alpha] = [\alpha_{C_i}]$, thus $Y_{v_i} \subseteq N_D(\mathrm{Mob}([\alpha_{C_i}], D))$. Since $H_i Y_{v_i} = Y_{v_i}$, we deduce using Corollary 3.37 that Y_{v_i} has finite Hausdorff distance from $H_i = \mathrm{Stab}([\alpha_{C_i}], D)$.

We now consider the case where W_i coarsely 3-separates G. Then W_i has finite Hausdorff distance from a subgroup $K \leq G$; by Proposition 4.44, we may assume $K \leq H_i$. Suppose $C \in \pi^{-1}(v_i)$. Then W_C and K have finite Hausdorff distance. By an argument of Vavrichek (see Proposition 4.3 of [Vav13] and its generalisation in Section 7 of [Vav12]), we deduce that there exists a $g \in \mathrm{Comm}_G(K) \cap H_i$ and uniform constant A_i such that W_C has Hausdorff distance at most A_i from gK. Thus Y_{v_i} has finite Hausdorff distance from $H_i W_i$. As W_i has finite Hausdorff distance from $K \leq H_i$, it follows that H_i has finite Hausdorff distance from Y_{v_i}. □

Recall that every V_1-vertex of T is a star ω, i.e. a collection of adjacent V_0-vertices. Thus every edge of the tree is of the form (w, ω), where w is a V_0-vertex and ω is a star with $v \in \omega$. We fix a set of representatives $e_1, \ldots, e_s$ for the finitely many G-orbits of edges in T, where $e_i = (w_i, \omega_i)$.

Lemma 4.59. *For any r sufficiently large, there exists an $R \geq 0$ such that the following holds. Let $e = (w, \omega)$ be any edge of T, with w and ω as above. Let $\mathcal{W}_e := \{(W_C, W_{C'}) \mid \pi(C) = w, \pi(C') \in \omega \setminus \{w\}\}$. We set*

$$Y_e := \bigcup_{(W,W') \in \mathcal{W}_e} W \cap N_r(W').$$

Then $d_{\mathrm{Haus}}(X_e, Y_e) \leq R$.

This can be proved via the following lemma.

Lemma 4.60. *Let $H, K, L \leq G$ be subgroups of G with $K \leq L$. Then for any $A, B \geq 0$, $\bigcup_{g \in L} N_A(H) \cap N_B(gK)$ is contained in finitely many right cosets of $H \cap L$.*

Proof. Let $X := \bigcup_{g \in L} N_A(H) \cap N_B(gK)$. For each $x \in X$, there exist $h_x \in H$, $g_x \in L$ and $k_x \in K$ such that $d(x, h_x) \leq A$ and $d(x, g_x k_x) \leq B$. We thus pick a finite subset $X_0 \subseteq X$ such that

$$\{(h_x^{-1}x, k_x^{-1}g_x^{-1}x) \mid x \in X\} = \{(h_x^{-1}x, k_x^{-1}g_x^{-1}x) \mid x \in X_0\}.$$

For each $x \in X$, there exists an $x_0 \in X_0$ with $(h_x^{-1}x, k_x^{-1}g_x^{-1}x) = (h_{x_0}^{-1}x_0, k_{x_0}^{-1}g_{x_0}^{-1}x_0)$. Then $x = h_x h_{x_0}^{-1} x_0$. Moreover, as $h_x^{-1}g_x k_x = h_{x_0}^{-1} g_{x_0} k_{x_0}$ and $K \leq L$, we deduce that $h_x h_{x_0}^{-1} \in H \cap L$, and so X is contained in at most $|X_0|$ right cosets of $H \cap L$. □

Proof of Lemma 4.59. As in the proof of Lemma 4.58, it is sufficient to only consider edges of the form $e = e_i = (w_i, \omega_i)$. We pick r large enough so that for every e_i, there exists some $(W, W') \in \mathcal{W}_{e_i}$ such that $W \cap N_r(W') \neq \emptyset$. By the choice of r, $Y_{e_i} \neq \emptyset$, so $G_{e_i} x \subset Y_{e_i}$ for any $x \in Y_{e_i}$. Therefore, we need only show that Y_{e_i} is contained in some metric neighbourhood of G_{e_i}.

Let $w_i = g' v_i$ and let $\{g_1 v_{j_1}, \dots g_k v_{j_k}\}$ be a set of representatives of G_{ω_i} orbits of vertices in the star ω_i. If $\pi(C) = w_i$, then $W_C \subseteq N_B(g' H_i)$ by Proposition 4.8 and Lemma 4.58. Similarly, if $\pi(C') \in \omega_i$, then $\pi(C') = g g_t v_{j_t}$ for some $g \in G_\omega$ and $1 \leq t \leq k$; so $W_{C'} \subseteq N_B(g g_t H_{j_t})$ by Proposition 4.8 and Lemma 4.58 as above. We thus deduce that

$$Y_{e_i} \subseteq \bigcup_{t=1}^{k} \bigcup_{g \in G_{\omega_i}} N_B(g' H_i) \cap N_{B+r}(g g_t H_{j_t}).$$

For each t, we apply Lemma 4.60, setting $H = g' H_i (g')^{-1} = G_{w_i}$, $L = G_\omega$ and $K = g_t H_{j_t} g_t^{-1} \leq L$. We thus see that Y_{e_i} is contained in finitely many right $G_{w_i} \cap G_{\omega_i} = G_{e_i}$ orbits as required. □

The following lemma has an almost identical proof, which we omit.

Lemma 4.61. *For any r sufficiently large, there exists an $R \geq 0$ such that the following holds. Let ω be a V_1-vertex of T. Let $\mathcal{W}_\omega = \{(W_C, W_{C'}) \mid \pi(C) \in \omega, \pi(C') \in \omega, \pi(C) \neq \pi(C')\}$. We set*

$$Y_\omega = \bigcup_{(W, W') \in \mathcal{W}_\omega} N_r(W) \cap N_r(W').$$

Then $d_{\mathrm{Haus}}(X_\omega, Y_\omega) \leq R$.

Proof of Theorem 4.57. There is a constant $A' \geq 0$ such that for every $(1, A)$-coarse complementary C of some $W \in \mathcal{W}$, $f(C)$ is a $(1, A')$-coarse complementary component of some $W' \in \mathcal{W}'$. We thus define $f_*(\pi(C)) := \pi(f(C))$, which induces a well-defined map between V_0-vertices. Since crossing is preserved by quasi-isometries, f_* preserves vertex type. It is clear from Lemmas 4.58, 4.59 and 4.61 that there is a constant R such that $f(X_r) \subseteq N_R(X'_{f_*(r)})$ for all r. We fix a coarse inverse $h : G' \to G$ and apply the same argument to h. We thus deduce that f_* is a tree isomorphism and $f(X_r)$ has uniformly finite Hausdorff distance from $X'_{f_*(r)}$. □

4.3.3 A toolkit for computing the regular neighbourhood

One drawback of the machinery developed in the previous section is that one doesn't know the edge stabilizers of the tree $T(\mathcal{W})$ in Theorem 4.42. Consequently, it is not entirely clear how this tree relates to other trees such as the JSJ tree or the JSJ tree of cylinders. To remedy this, we explain how one can read off mobility sets in a graph of groups, provided the graph of groups satisfies certain cohomological assumptions. We expect the techniques developed here to be crucial to applications of Theorem 4.42. In particular, we expect to be able to prove theorems of the sort found in [MSW03], [MSW11] and [Pap07].

Definition 4.62. Let G be a group and $\mathcal{S}$ a collection of finitely many coarse $PD_n^{\mathbb{Z}_2}$ subgroups. Let $G\mathcal{S}$ denote the set of left cosets of elements of $\mathcal{S}$. Then we say that $(G, \mathcal{S})$ is a *coarse* $PD_{n+1}^{\mathbb{Z}_2}$ *pair* if the following hold:

- G is of type $FP_\infty^{\mathbb{Z}_2}$, i.e. of type $FP_n^{\mathbb{Z}_2}$ for every n;
- $H^k_{\text{coarse}}(G) = 0$ for $k \neq n$;
- there is a short exact sequence

$$1 \to H^n_{\text{coarse}}(G) \xrightarrow{\oplus p_{gS}} \bigoplus_{gS \in G\mathcal{S}} H^n_{\text{coarse}}(gS) \xrightarrow{\epsilon} \mathbb{Z}_2 \to 0,$$

  where each p_{gS} is induced by the inclusion $gS \to G$, and ϵ is the augmentation map that sends the generator of each $H^n_{\text{coarse}}(gS) \cong \mathbb{Z}_2$ to 1.

A Poincaré duality pair in the sense of [BE78] is a coarse $PD_{n+1}^{\mathbb{Z}_2}$ pair. If v is a quadratically hanging vertex group in a JSJ decomposition of a group over VPC-n subgroups, then $(G_v, \mathcal{S}_v)$ is a coarse $PD_{n+1}^{\mathbb{Z}_2}$ pair, where $\mathcal{S}_v$ is the collection of edge groups incident to v.

There is another, perhaps more intuitive, notion of coarse $PD^{\mathbb{Z}_2}_{n+1}$ pair, requiring the double of the pair $(G, \mathcal{S})$ to be a $PD^{\mathbb{Z}_2}_{n+1}$ group. The above definition includes this notion of coarse $PD^{\mathbb{Z}_2}_{n+1}$ pairs, but is presumably more general. Our definition is easier to work with in practice; in particular, there is a combination theorem of coarse $PD^{\mathbb{Z}_2}_{n+1}$ pairs, analogous to that of $PD^{\mathbb{Z}_2}_{n+1}$ pairs (see Theorem 8.1 and 8.3 of [BE78]). The main advantage of our definition is that in order to show something is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ pair, one doesn't need to worry about constructing augmentation maps on metric complexes as in Remark 2.52 — everything can be done entirely using cohomology.

Standing Assumptions. We assume $\mathcal{G}$ is a graph of groups with the following properties:

(1) it is reduced the sense of [BF91];

(2) every edge group is coarse $PD^{\mathbb{Z}_2}_n$;

(3) every vertex group G_v is one of three types:

- $(G_v, \mathcal{S}_v)$ is a coarse $PD^{\mathbb{Z}_2}_{n+1}$ pair, where $\mathcal{S}_v$ is the collection of edge groups incident to v. In this case we say v is *hanging*.
- G_v is a coarse $PD^{\mathbb{Z}_2}_n$ group.
- G_v is of type $FP^{\mathbb{Z}_2}_{n+1}$ and is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$.

Remark 4.63. Much of what follows can be deduced algebraically using relative cohomology and the Mayer-Vietoris sequences as described in Sections 3 and 8 of [BE78]. However, in the spirit of this thesis, we give a geometric argument using coarse cohomology.

We fix a graph of groups $\mathcal{G}$ satisfying the standing assumptions and let $G = \pi_1(\mathcal{G})$ be its fundamental group and T its associated Bass-Serre tree.

Definition 4.64. A subtree T_0 of T is said to be *topologically homogeneous* if every vertex of T_0 whose valence (in T_0) is greater than one is hanging.

Example 4.65. The Bass-Serre tree of a non-geometric graph manifold is topologically homogeneous. More generally, suppose we have a JSJ tree T of a finitely presented group G over VPC-n subgroups. Any subtree whose interior vertices are all quadratically hanging, and arranged in a 'manifold-like way' so there is no coarse 3-separation, is topologically homogeneous.

Suppose that $G = A *_C B$, where all groups are of type $FP_{n+1}^{\mathbb{Z}_2}$. There is an associated long exact sequence in group cohomology

$$\cdots \to H^n(G, \mathbb{Z}_2 G) \to H^n(A, \mathbb{Z}_2 G) \oplus H^n(B, \mathbb{Z}_2 G) \to H^n(C, \mathbb{Z}_2 G) \to \cdots .$$

The non-connecting maps (i.e. the maps that don't shift dimension) are restriction maps. A similar long exact sequence holds for HNN-extensions and more generally, for graphs of groups. We refer the reader to Section 3 of [BE78] and Section VII.9 of [Bro82] for more details.

In Section 2.2.5, we saw that group cohomology $H^n(G, \mathbb{Z}_2 G)$ may be thought of as the compactly supported cohomology of a suitable CW complex or metric complex. How ought one to think of $H^n(A, \mathbb{Z}_2 G)$? We observe that $\mathbb{Z}_2 G$ is a direct sum of copies of $\mathbb{Z}_2 A$, where each summand corresponds to a left coset of A. Thus the cohomology $H^n(A, \mathbb{Z}_2 G)$ can be written as a direct sum of copies of $H^n(A, \mathbb{Z}_2 A)$, each copy corresponding to a left A coset.

Using the methods described in [Sco78], we can build a graph of spaces with finitely many cells in each dimension, corresponding to the splitting $G = A *_C B$. We think of $H^n(A, \mathbb{Z}_2 G)$ as the direct sum of coarse cohomology of each of the vertex spaces associated to cosets of A in G. We interpret $H^n(B, \mathbb{Z}_2 G)$ and $H^n(C, \mathbb{Z}_2 G)$ similarly. The above Mayer-Vietoris sequence can thus be interpreted geometrically as a Mayer-Vietoris in terms of vertex and edge spaces. We state this Mayer-Vietoris sequence more generally for a graph of $FP_{n+1}^{\mathbb{Z}_2}$ groups:

$$\cdots \to H^k_{\text{coarse}}(G) \xrightarrow{\oplus q_v} \bigoplus_{v \in VT} H^k_{\text{coarse}}(X_v) \xrightarrow{\oplus p_e^v} \bigoplus_{e \in ET} H^k_{\text{coarse}}(X_e) \xrightarrow{\delta} H^{k+1}_{\text{coarse}}(G) \xrightarrow{\delta} \cdots . \tag{4.1}$$

This sequence holds in dimensions at most $n+1$. The maps q_v and p_e^v are induced by inclusions, as the corresponding maps on the level of group cohomology are restriction maps. We also remark that as we are working with $\mathbb{Z}_2$ coefficients, there are no sign issues. This Mayer-Vietoris can also be deduced directly using a spectral sequence argument, see Section VII.4 of [Bro82].

We return to the example at hand: a graph of groups $\mathcal{G}$, with Bass-Serre tree T, satisfying our standing assumptions. For each edge $e \in ET$, let $[\beta_e] \in H^n_{\text{coarse}}(X_e)$ be the non-trivial cohomology class. Let $[\alpha_e] = \delta([\beta_e])$, where δ is the connecting homomorphism in the Mayer-Vietoris sequence (4.1). In order to apply our regular neighbourhood construction, we first need to show that each $[\alpha_e] \neq 0$.

Lemma 4.66. *G is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$, hence every edge space is essentially embedded.*

Proof. This can be deduced using the Mayer-Vietoris sequences in [BE78]. Alternatively, we can apply the above Mayer-Vietoris sequence (4.1) to the graph of groups $\mathcal{G}$ as follows. The hypotheses on vertex groups and edge groups ensure that $H^k(G, \mathbb{Z}_2 G) = 0$ for $k \leq n-1$, thus G is $(n-2)$-acyclic at infinity over $\mathbb{Z}_2$.

The part of the Mayer-Vietoris sequence we now need to consider is

$$0 \to H^n_{\text{coarse}}(G) \xrightarrow{\oplus q_v} \bigoplus_{v \in VT} H^n_{\text{coarse}}(X_v) \xrightarrow{\oplus p_e^v} \bigoplus_{e \in ET} H^n_{\text{coarse}}(X_e) \to \cdots .$$

We claim for each $v \in VT$, the map $H^n_{\text{coarse}}(X_v) \xrightarrow{\oplus_e p_e^v} \bigoplus_{e \in \text{star}(v)} H^n_{\text{coarse}}(X_e)$ is injective. If v is hanging, this is clear from the definition of a coarse $PD^{\mathbb{Z}_2}_{n+1}$ pair. If G_v is $(n-1)$-acyclic at infinity, then $H^n_{\text{coarse}}(X_v) = 0$ so that this map is also injective. If G_v is coarse $PD^{\mathbb{Z}_2}_n$, then Proposition 4.43 ensures that each map $p_e^v : H^n_{\text{coarse}}(X_v) \to H^n_{\text{coarse}}(X_e)$ is an isomorphism. Therefore $H^n_{\text{coarse}}(G) = 0$, and so G is $(n-1)$-acyclic at infinity over $\mathbb{Z}_2$. □

We can now describe precisely when the cohomology classes $[\alpha_e]$ and $[\alpha_{e'}]$ are equal.

Proposition 4.67. *For any $e, e' \in ET$, $[\alpha_e] = [\alpha_{e'}]$ if and only if e and e' lie in a topologically homogeneous subtree.*

Let $\mathcal{W}$ be the collection of edge spaces of T. The point of Proposition 4.67 is that it gives us everything we need to obtain the tree $T(\mathcal{W})$ from the Bass-Serre tree T.

To do this, we first partially collapse maximal topologically homogeneous subtrees by mapping all vertices of valence more than one to a point. In doing so, the resulting central vertex group and its incident edge groups will still form a coarse $PD^{\mathbb{Z}_2}_n$ pair, so the resulting graph of groups still satisfies the standing assumptions. When taking the tree of cylinders of this tree, we obtain the regular neighbourhood $T(\mathcal{W})$. The vertices of $T(\mathcal{W})$ that have vertex stabilizer of the form $\text{Stab}([\alpha])$ correspond to partially collapsed maximal topologically homogeneous subtrees. We note that edges adjacent to V_0-vertices of this form have coarse $PD^{\mathbb{Z}_2}_n$ stabilizers.

Before proving Proposition 4.67, we need the following technical lemma. It says that in the above Mayer-Vietoris sequence, an element of $\bigoplus_{e \in ET} H^k_{\text{coarse}}(X_e)$ that looks like it is contained in the image of some $H^n_{\text{coarse}}(X_w)$ is necessarily contained in the image of $H^n_{\text{coarse}}(X_w)$.

Lemma 4.68. *Let $w \in VT$. Suppose that $[\beta] \in \bigoplus_{e \in \text{star}(w)} H^n_{\text{coarse}}(X_e)$, and $[\beta] = (\oplus p_e^v)([\sigma])$ for some $[\sigma] \in \bigoplus_{v \in VT} H^n_{\text{coarse}}(X_v)$. Then there exists a $[\sigma'] \in H^n_{\text{coarse}}(X_w)$ such that $[\beta] = (\oplus p_e^v)([\sigma'])$.*

Proof. There exist finitely many distinct vertices $w = v_1, \dots, v_m$ such that $[\sigma] = [\sigma_1] + \cdots + [\sigma_m]$ and $[\sigma_i] \in H^n_{\text{coarse}}(X_{v_i})$. Replacing $[\sigma]$ with $[\sigma] - [\sigma_i]$ when necessary, we may assume $(\oplus p_e^{v_i})[\sigma_i] \neq 0$ for all $i > 1$. Thus G_{v_i} cannot be $(n-1)$-acyclic at infinity for $i > 1$. We pick v_i with $d_T(w, v_i)$ maximal. It is straightforward to check that when v_i is either hanging or G_{v_i} is coarse $PD^{\mathbb{Z}_2}_n$ group, $(\oplus p_e^{v_i})[\sigma_i]$ must be supported on at least two edges e_1 and e_2, i.e. $(\oplus p_e^{v_i})[\sigma_i] \in \bigoplus_{e \in \text{star}(w)} H^n_{\text{coarse}}(X_e)$ has nonzero $H^n_{\text{coarse}}(X_{e_1})$ and $H^n_{\text{coarse}}(X_{e_2})$ coordinates. One of these, say e_1, is not contained in the ball of radius $N = d(w, v_i)$ around w. By our choice of v_i, we see that e_1 is adjacent to v_i and no other v_j. Thus for $j \neq i$, $(\oplus p_e^v)[\sigma_j]$ is not supported on e_1, so that $(\oplus p_e^v)[\sigma] = [\beta]$ must be supported on e_1. By hypotheses, $[\beta]$ is only supported on edges in $\text{star}(w)$, so $N = 0$ and therefore $m = 1$. □

Proof of Proposition 4.67. It is sufficient to consider the case where e and e' are distinct edges incident to a single vertex v; the general case can then be obtained inductively. If v is hanging, then it follows from the definition of coarse $PD^{\mathbb{Z}_2}_n$ pairs that $[\beta_e] - [\beta_{e'}]$ is contained in the image of the map

$$H^n_{\text{coarse}}(X_v) \xrightarrow{\oplus_e p_e^v} \bigoplus_{e \in \text{star}(v)} H^n_{\text{coarse}}(X_e) \tag{4.2}$$

in the Mayer-Vietoris sequence. Thus $[\alpha_e] = [\alpha_{e'}]$.

If G_v is $(n-1)$-acyclic at infinity, then $[\beta_e] - [\beta_{e'}]$ is not contained in the image of the map shown in (4.2). By Lemma 4.68, we deduce that $[\alpha_e] \neq [\alpha_{e'}]$. If G_v is coarse $PD^{\mathbb{Z}_2}_n$ space, we let $e_1, \dots, e_m$ be the edges incident to v. The image of the map given in (4.2) is therefore the span of $\{[\beta_{e_1}] + \cdots + [\beta_{e_m}]\}$. Note $n > 2$, as the underlying graph of groups is assumed to be reduced. Thus $[\beta_e] - [\beta_{e'}]$ is not contained in the image of the above map. Applying Lemma 4.68 as above, we deduce that $[\alpha_e] \neq [\alpha_{e'}]$. □

4.4 More refined data from the JSJ decomposition

In this section we prove Theorem 4.74, which gives a necessary criterion for certain groups to be quasi-isometric in terms of their JSJ decompositions. This can then be combined with the work of Cashen and Martin [CM17], where the combinatorics of the JSJ tree of cylinders provides us with even more data that is preserved under quasi-isometry. These methods will be crucial in Chapter 5, in which we determine the quasi-isometry type of certain RAAGs.

4.4.1 Relative quasi-isometries

Quasi-isometries are often hard to work with directly because of their inherent coarseness. To circumvent this difficulty, a common technique in geometric group theory is to take a quasi-isometry $f : G \to G'$ and show that it induces a "nicer" map $f_* : Q(G) \to Q(G')$ between objects $Q(G)$ and $Q(G')$ that are associated to G and G'. For instance, quasi-isometries between finitely generated groups induce bi-Lipschitz equivalences between any two of their asymptotic cones, quasi-isometries between hyperbolic groups induce homeomorphisms between their boundaries etc.

Suppose we know that a quasi-isometry $f : G \to G'$ preserves JSJ trees in the following sense: if T and T' are JSJ trees of G and G', then f sends vertex (resp. edge) spaces of T to within uniform Hausdorff distance of vertex (resp. edge) spaces of T'. We want to upgrade this data to something more combinatorial. It is not necessarily the case that f induces a tree isomorphism; for instance, $\mathrm{Out}(G)$ does not necessarily act on a JSJ tree of G.

We thus take the tree of cylinders and show that f induces a tree isomorphism between JSJ trees of cylinders. However, in doing this, we often lose lots of information, especially in the situation where cylinders are large, i.e. the stabilizer of an edge has infinite index in the stabilizer of its cylinder. To keep track of this extra information, we attach extra data to vertices $v \in T_{\mathrm{cyl}}$ that remembers certain edge spaces $\{X_e\}_{e \in ET}$ of the original JSJ tree.

In [CM17], relative quasi-isometries are defined. To do this, a *peripheral structure* — a collection of coarse equivalence classes of subspaces — is assigned to a space. In [CM17], peripheral structures are defined as follows: given a tree of cylinders T, a vertex space X_v is given the peripheral structure consisting of the collection of all edge spaces X_e for edges e that are incident to v. However, as remarked above, one often loses information in this peripheral structure when cylinder stabilizers are large. The situation we want to avoid is the following:

Example 4.69. Suppose a graph of groups decomposition has a vertex group $\mathbb{Z}^2 = \langle a, b \mid [a, b] \rangle$, incident to exactly two edges whose edge groups are $\langle a \rangle$ and $\langle b \rangle$ respectively. In the Bass-Serre tree T, a lift v of this vertex is contained in exactly two cylinders, one containing edges corresponding to cosets of the form $b^i\langle a \rangle$, and the other containing cosets of the form $a^i\langle b \rangle$. In the tree of cylinders T_{cyl}, the V_1-vertex v is incident to exactly two edges. Both of the corresponding edge spaces are coarsely equivalent to the vertex space X_v; this is because $\cup_{i \in \mathbb{Z}}(b^i\langle a \rangle) = \cup_{i \in \mathbb{Z}}(a^i\langle b \rangle) = \mathbb{Z}^2$. If we simply give X_v a peripheral structure consisting of coarse equivalence classes of adjacent edge

spaces, then both elements in the peripheral structure are equal, even though they lie in different cylinders.

The right thing to do is to think of an element of the peripheral structure as a family of coarsely parallel subspaces. A relative quasi-isometry ought to preserve coarse equivalence classes of subspaces themselves, not just the coarse equivalence class of their union. Doing this allows us to distinguish between the two incident edge spaces.

This example motivates the following definition. We caution the reader that as far as we are aware, this definition is not particularly intuitive or useful outside the specific context in which we use it: namely vertex spaces of a tree of cylinders.

Definition 4.70. A *peripheral structure* $\mathcal{P}_X$ on a metric space X is a set containing elements of the form $\{A_i\}_{i\in I}$, where each $A_i \subseteq X$ and $[A_i] = [A_j]$ for all $i, j \in I$. A (K, A)-*relative quasi-isometry* $(f, f_*) : (X, \mathcal{P}_X) \to (Y, \mathcal{P}_Y)$ consists of a (K, A)-quasi-isometry $f : X \to Y$ and a bijection $f_* : \mathcal{P}_X \to \mathcal{P}_Y$ such that if $f_*(\{A_i\}_{i\in I}) = \{B_j\}_{j\in J}$, then:

(1) for every $i \in I$, there is some $j \in J$ such that $d_{\text{Haus}}(f(A_i), B_j) \leq A$;

(2) for every $j \in J$, there is some $i \in I$ such that $d_{\text{Haus}}(f(A_i), B_j) \leq A$.

We say that (f, f_*) is a *relative quasi-isometry* if it is a (K, A)-relative quasi-isometry for some $K \geq 1$ and $A \geq 0$. To simplify notation, we often suppress the map between peripheral structures, denoting (f, f_*) by f. We recover the standard notion of relative quasi-isometry (as used in [CM17] for instance) if for every $\{A_i\}_{i\in I} \in \mathcal{P}_X$, $|I| = 1$.

Example 4.71. Let $G = \mathbb{Z}^2 = \langle a, b \mid [a, b]\rangle$. Let $\mathcal{P}_1 := \{\{b^i\langle a\rangle \mid i \in \mathbb{Z}\}\}$, $\mathcal{P}_2 := \{\{a^i\langle b\rangle \mid i \in \mathbb{Z}\}\}$ and $\mathcal{P}_3 := \mathcal{P}_1 \cup \mathcal{P}_2$. The map id_G is not a relative quasi-isometry $(G, \mathcal{P}_1) \to (G, \mathcal{P}_2)$, however the isomorphism $f : G \to G$ given by $a^i b^j \mapsto a^j b^i$ is a relative quasi-isometry $(G, \mathcal{P}_1) \to (G, \mathcal{P}_2)$. We note that $(G, \mathcal{P}_3)$ is not relatively quasi-isometric to either $(G, \mathcal{P}_1)$ or $(G, \mathcal{P}_2)$. Although $\mathcal{P}_4 := \cup_{i\in\mathbb{Z}}\{\{b^i\langle a\rangle\}\}$ is a well-defined peripheral structure on G, it should not be confused with $\mathcal{P}_1$, since $\mathcal{P}_4$ contains infinitely many elements, whereas $\mathcal{P}_1$ contains a single element.

Being relatively quasi-isometric is an equivalence relation among pairs $(X, \mathcal{P}_X)$. We define the *relative quasi-isometry class* $[[(X, \mathcal{P}_X)]]$ to be the equivalence class consisting of all pairs $(Y, \mathcal{P}_Y)$ such that there exists a relative quasi-isometry $(X, \mathcal{P}_X) \to (Y, \mathcal{P}_Y)$. We let $\text{QI}(X, \mathcal{P}_X)$ denote that set of all relative quasi-isometries $(X, \mathcal{P}_X) \to (X, \mathcal{P}_X)$.

Standing Assumptions. Let G be a group and T a minimal G-tree such that all vertex and edge stabilizers of T are finitely generated. We assume no edge stabilizer is contained in another edge stabilizer with infinite index. We also assume that for any finitely generated group G' and any (K, A)-quasi-isometry $f : G \to G'$, there is a G'-tree T' and a constant $B = B(K, A)$ such that the following holds:

- for any $e \in ET$, there is an $e' \in ET'$ such that $d_{Haus}(f(X_e), X_{e'}) \leq B$;
- for any $e' \in ET'$, there is an $e \in ET$ such that $d_{Haus}(f(X_e), X_{e'}) \leq B$.

Moreover, we assume that when $G' = G$, the above holds with $T' = T$.

Example 4.72. If G is a finitely presented group with JSJ tree T over two-ended subgroups, work of Papasoglu [Pap05] shows that the above assumptions are satisfied. Another example is the following: let G be a group of type $FP_{n+1}^{\mathbb{Z}_2}$ and let T be a JSJ tree over the class of coarse $PD_n^{\mathbb{Z}_2}$ subgroups that coarsely 3-separate G. The methods of Chapter 3 implies these assumptions are satisfied.

We now explain how to give vertex spaces of the tree of cylinders T_{cyl} a peripheral structure. Each edge $e \in ET_{\text{cyl}}$ is incident to some V_0-vertex v and V_1-vertex w. We note that w is naturally identified with a vertex of T, which we also call w. We thus define

$$\mathcal{S}_e := \{X_f \mid f \in ET, f \text{ is incident to } w \text{ and } f \text{ is contained in the cylinder } v\}$$

and for each vertex $v \in T_{\text{cyl}}$, we give X_v the peripheral structure

$$\mathcal{P}_v := \{\mathcal{S}_e \mid \text{for all } e \in ET_{\text{cyl}} \text{ such that } e \text{ is adjacent to } v\}.$$

Remark 4.73. It would be more accurate, though notationally inconvenient, to refer to elements of $\mathcal{P}_v$ as pairs $(\mathcal{S}_e, e)$. This is because it may be the case that $S_e = S_{e'}$ for some $e \neq e'$, in which case we would like to have two distinct elements S_e and $S_{e'}$ of $\mathcal{P}_v$. This situation only occurs in the exceptional case where some edge $f = (w, w') \in ET_{\text{JSJ}}$ is the only edge contained in its cylinder v. Then $e = (v, w)$ and $e' = (v, w')$ are distinct edges of ET_{cyl} with $\mathcal{S}_e = \mathcal{S}_{e'} = \{X_f\}$.

This defines $\mathcal{P}_v$ for every vertex v unless T_{cyl} consists of precisely one vertex v. In this case, we set $\mathcal{P}_v = \{\{X_f \mid f \in ET\}\}$ if the vertex v is cylindrical, and $\mathcal{P}_v = \emptyset$ otherwise.

Theorem 4.74. *Let G and T satisfy the standing assumptions and suppose $f : G \to G'$ is a (K, A)-quasi-isometry; let T' be the corresponding G'-tree as in the standing assumptions. Let T_{cyl} and T'_{cyl} be trees of cylinders of T and T'. Then there is a constant C and a unique tree isomorphism $f_* : T_{\text{cyl}} \to T'_{\text{cyl}}$ such that the following holds.*

For every vertex v of T, $f(X_v)$ has Hausdorff distance at most C from $X'_{f_*(v)}$, and v and $f_*(v)$ are of the same vertex type, i.e. either they are both cylinders or neither are cylinders. Moreover, f restricted to each vertex space X_v is close to a relative quasi-isometry $(X_v, \mathcal{P}_v) \to (X_{f_*(v)}, \mathcal{P}_{f_*(v)})$.

In the case where T is the JSJ tree of cylinder of G over 2-ended subgroups, the standing assumptions are satisfied by work of [Pap05]. However, Theorem 4.74 is stronger than comparable results in the literature, see Theorem 2.8 of [CM17]. This is because the peripheral structure we use is finer than that considered in [CM17]. Theorem 4.74 will be applied in Chapter 5 to show one can (often) determine whether RAAGs are quasi-isometric using the JSJ decomposition over abelian (not necessarily 2-ended) subgroups.

We note that if T has quadratically hanging vertices and these are preserved by quasi-isometries in an appropriate sense, then one can strengthen Theorem 4.74 by deducing that f_* must preserve quadratically hanging vertices. By work of Papasoglu [Pap05], this is the case when considering JSJ decompositions of one-ended finitely presented groups over their two-ended subgroups.

We first prove the following Lemma.

Lemma 4.75. *Let $\mathcal{G}$ be a finite graph of groups with finitely generated edge and vertex groups, let T be the associated Bass-Serre tree and let X be the tree of spaces. There is a proper non-decreasing function $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that if X_e and $X_{e'}$ are edge spaces such that $[X_e] = [X_{e'}]$ and $X'_e \subseteq N_r(X_e)$ for some $r \geq 0$, then $\mathrm{d}_{\text{Haus}}(X'_e, X_e) \leq \phi(r)$.*

Proof. For some fixed $r \geq 0$ and edge space X_e, there are only finitely many edge spaces such that $[X_e] = [X_{e'}]$ and $X'_e \subseteq N_r(X_e)$. This follows from Proposition 4.8 and the observation that there are only finitely many cosets of edge groups of $\mathcal{G}$ that are contained in some $N_r(gG_e)$. We pick $R_e \geq 0$ such that $\mathrm{d}_{\text{Haus}}(X'_e, X_e) \leq R_e$ for all such edge spaces $X_{e'}$. The result holds as there are only finitely many $\pi_1(\mathcal{G})$-orbits of edge spaces in T. □

Proof of Theorem 4.74. Let $\mathcal{W}$ and $\mathcal{W}'$ be the collection of all edge spaces of T and T' respectively. Note that there are only finitely many G and G' orbits of elements

of $\mathcal{W}$ and $\mathcal{W}'$ respectively, so we can consider the regular neighbourhoods $T(\mathcal{W})$ and $T(\mathcal{W}')$, as in Theorem 4.42. By Theorem 3.3 of [GL10], $T(\mathcal{W})$ and $T(\mathcal{W}')$ may be identified with a subdivision of the trees of cylinders T_{cyl} and T'_{cyl}. This works because all elements of $\mathcal{W}$ and $\mathcal{W}'$ are actually finite Hausdorff distance from subgroups, so $T(\mathcal{W})$ and $T(\mathcal{W}')$ are regular neighbourhoods in the sense of [SS03].

If C is a coarse complementary component of some $W \in \mathcal{W}$, then our standing assumptions and Lemma 3.5 imply that $f(C)$ is a coarse complementary component of some $W' \in \mathcal{W}$. Similarly, if C' is coarse complementary component of some $W' \in \mathcal{W}'$, then $[C'] = [f(C)]$, where C is some coarse complementary component of some $W \in \mathcal{W}$ It follows from the construction of $T(\mathcal{W})$ and $T(\mathcal{W}')$ that f induces a tree isomorphism f_*.

Theorem 4.57 ensures that f sends edge and vertex spaces of $T(\mathcal{W})$ to within uniform finite Hausdorff distance of edge and vertex spaces of $T(\mathcal{W}')$ and vice versa. All that remains is to show that f preserves the peripheral structure. Pick some edge $e = (v, w) \in ET_{\text{cyl}}$ and let $\mathcal{S}_e$ be as defined above. Let $e' := (v', w') = (f_*(v), f_*(w))$. Without loss of generality, we may assume v is a cylinder and w is a V_1-vertex. Let $X_t \in \mathcal{S}_e$, where $t \in ET$ is an edge adjacent to w and contained in the cylinder v.

Our standing assumptions imply $f(X_t)$ has uniform finite Hausdorff distance from $X_{t'}$ for some $t' \in ET'$. This is enough if t' happens to be adjacent to w', since then $X_{t'} \in \mathcal{S}_{e'}$; however, this is not necessarily the case. Since $X_t \subseteq N_1(X_w)$, we deduce that $X_{t'} \subseteq N_D(X_{w'})$, where D can be chosen uniformly. Taking an edge path in T' from t' to w', we deduce that there exists an edge $t'' \in ET'$, adjacent to w', such that $X_{t'} \subseteq N_D(X_{t''})$. We note that $X_{t''} \in \mathcal{S}_{e'}$. By Lemma 4.75, we deduce that $f(X_t)$ and $X_{t''}$) have uniformly finite Hausdorff distance, which demonstrates the first condition in Definition 4.70. The other condition can be proven similarly. □

4.4.2 Decorated trees

The material in this subsection is adapted from [CM17]. The question we wish to answer is the following: when is there a tree isomorphism $\chi : T_{\text{cyl}} \to T'_{\text{cyl}}$ that satisfies the conclusions of Theorem 4.74? Theorem 4.74 says the existence of such a tree isomorphism is a necessary condition for G and G' to be quasi-isometric when G and T satisfy the standing assumptions. This gives a powerful method that is frequently able to show two groups are not quasi-isometric, and will be used in Chapter 5.

Definition 4.76. Let G be a finitely presented group and let T be a minimal G-tree. Let $\mathcal{O}$ be a set and $\delta : VT \sqcup ET \to \mathcal{O}$ be a map such that $\delta(e) = \delta(\bar{e})$ for every edge

$e \in ET$. We call the map δ a *decoration* and call $\mathcal{O}$ the set of *ornaments*. We call the pair (T, δ) a *decorated tree*.

We may think of a decorated tree as a tree in which each vertex and unoriented edge is labelled with an element of $\mathcal{O}$. For instance, if $\mathcal{O}$ is a set of colours, then a decorated tree is just a coloured tree. We assume throughout that no edge has the same ornament as a vertex, thus $\mathcal{O}$ can be thought of as a disjoint union of vertex and edge ornaments.

Definition 4.77. Let (T, δ) and (T, δ') be two decorated trees with the same set of ornaments. Then we say that (T, δ) and (T, δ') are *equivalent* if there exists a decoration preserving tree isomorphism $\chi : T \to T'$.

We fix a group G and T satisfying the assumptions in the previous section. Let $T_c := T_{\text{cyl}}$ be the tree of cylinders of G. Recall every quasi-isometry $f : G \to G$ induces a unique tree isomorphism f_* as in Theorem 4.74.

Since we interested in determining the quasi-isometry type of a group, we will work with decorations that are invariant under quasi-isometry in the following sense:

Definition 4.78. A decoration $\delta : T_c \to \mathcal{O}$ is said to be *geometric* if for any quasi-isometry $f : G \to G$ with associated tree isomorphism $f_* : T_c \to T_c$, we have $\delta \circ f_* = \delta$.

We can decorate T_c with the following geometric decoration. We decorate each vertex v with its type (cylinder or V_1) and its relative quasi-isometry class $[[X_v, \mathcal{P}_v]]$. We call this the *naïve decoration*. Theorem 4.74 ensures that this is indeed a geometric decoration.

The vertices and edges of a decorated tree are naturally endowed with an equivalence relation as follows: two vertices or edges are equivalent if and only if they have the same ornament. We say a decoration $\delta' : T_c \to \mathcal{O}'$ is a *refinement* of $\delta : T_c \to \mathcal{O}$ if the partition of vertices and edges into $\mathcal{O}'$-classes is finer than the partition into $\mathcal{O}$-classes.

We now explain two procedures described in [CM17] that, given any geometric decoration $\delta : T_c \to \mathcal{O}$, output a new geometric decoration δ' that refines δ. The underlying principle is as follows: we show that there are some vertices or edges x, y such that $\delta(x) = \delta(y)$ but there is some obvious obstruction to the existence of a quasi-isometry $f : G \to G$ such that $f_*(x) = y$. In such a case, we refine δ to eliminate this obstruction.

The first of these procedures is called *neighbour refinement*. The idea is as follows. We suppose $\mathcal{O}$ contains two elements that we call `red` and `blue`, and that T contains two `red` vertices v and w. Suppose also that no `blue` edge is incident to v, but

some `blue` edge is incident to w. Since the decoration is geometric, there can be no quasi-isometry $\phi : G \to G$ such that $\phi_*(v) = w$, since ϕ_* must send `blue` edges to `blue` edges.

We refine δ by assigning different ornaments to `red` vertices incident to no `blue` edge, `red` vertices incident to countably infinitely many `blue` edges, `red` vertices incident to one `blue` edge, `red` vertices incident to two `blue` edges etc. Neighbour refinement applies the above procedure to all distinct pairs $\{\texttt{red}, \texttt{blue}\} \subseteq \mathcal{O}$ simultaneously. We refer the reader to [CM17, Section 3.2] for a formal exposition of this.

The second procedure is called *vertex refinement*. We first observe that a decorated tree of cylinders (T, δ) induces a decoration $\delta_v : \mathcal{P}_v \to \mathcal{O}$ on the peripheral structure via the map $\mathcal{S}_e \mapsto \delta(e)$. If X_v is a vertex space with peripheral structure $\mathcal{P}_v$, then there is a bijective correspondence between edges adjacent to v and elements of $\mathcal{P}_v$. Suppose $(\phi_v, (\phi_v)_*) : (X_v, \mathcal{P}_v) \to (X_w, \mathcal{P}_w)$ is a relative quasi-isometry. For ease of notation, we identify $(\phi_v)_*$ with the corresponding map that sends an edge adjacent to v to an edge adjacent to w.

Suppose that v and w are vertices such that $\delta(v) = \delta(w)$. Suppose also that there is no relative quasi-isometry $(X_v, \mathcal{P}_v) \to (X_w, \mathcal{P}_w)$ that preserves the decoration on the peripheral structure. Then there cannot be a quasi-isometry $\phi : G \to G$ such that $\phi_*(v) = w$. We thus refine δ by assigning different ornaments to v and w.

Now suppose there are edges $e = (v, w)$ and $f = (v', w')$ such that $\delta(e) = \delta(f)$. Suppose also that either there is no relative quasi-isometry $\phi_v : X_v \to X_{v'}$ such that $(\phi_v)_*(e) = f$, or there is no relative quasi-isometry $\phi_w : X_w \to X_{w'}$ such that $(\phi_w)_*(\overline{e}) = \overline{f}$. In either case, there cannot be a quasi-isometry $\phi : G \to G$ such that $\phi_*(e) = f$. We thus refine δ by assigning different ornaments to e and f.

Vertex refinement is treated in [CM17, Section 5.3]. Our approach needs minor modifications, so we explain the details here. We define $\mathcal{Q} := \{[[X_v, \mathcal{P}_v]] \mid v \in VT_c\}$. For each $Q \in \mathcal{Q}$, we pick some metric space Z_Q with peripheral structure $\mathcal{P}_Q$ such that $[[Z_Q, \mathcal{P}_Q]] = Q$.

We recall every $\phi \in \mathrm{QI}(Z_Q, \mathcal{P}_Q)$ induces a bijection $\phi_* : \mathcal{P}_Q \to \mathcal{P}_Q$. We define

$$\mathcal{O}'_V := \mathcal{O} \times \left(\bigsqcup_{Q \in \mathcal{Q}} \mathrm{QI}(Z_Q, \mathcal{P}_Q) \backslash \mathcal{O}^{\mathcal{P}_Q} \right),$$

where some $\phi \in \mathrm{QI}(Z_Q, \mathcal{P}_Q)$ acts on $\rho \in \mathcal{O}^{\mathcal{P}_Q}$ via $\phi \cdot \rho := \rho \circ \phi_*^{-1}$. Similarly, we define

$$\mathcal{O}'_E := \mathcal{O} \times \left(\bigsqcup_{Q \in \mathcal{Q}} \mathrm{QI}(Z_Q, \mathcal{P}_Q) \backslash (\mathcal{O}^{\mathcal{P}_Q}, \mathcal{P}_Q) \right)^2,$$

where some $\phi \in \mathrm{QI}(Z_Q, \mathcal{P}_Q)$ acts on $(\rho, \mathcal{S}_e) \in (\mathcal{O}^{\mathcal{P}_Q}, \mathcal{P}_Q)$ via

$$\phi \cdot (\rho, \mathcal{S}_e) := \Big(\rho \circ \phi_*^{-1}, \phi_*(\mathcal{S}_e)\Big).$$

The vertex refinement $\delta' : VT_c \sqcup ET_c \to \mathcal{O}'_V \sqcup \mathcal{O}'_E$ of δ is defined as follows. For every vertex $v \in VT_c$, we choose a relative quasi-isometry $(\mu_v, (\mu_v)_*) : (X_v, \mathcal{P}_v) \to (Z_Q, \mathcal{P}_Q)$ for $Q = [[(X_v, \mathcal{P}_v)]]$. We recall δ_v is the decoration on $\mathcal{P}_v$ induced by δ. We thus define

$$\delta'(v) := \left(\delta(v), \mathrm{QI}(Z_Q, \mathcal{P}_Q) \cdot (\delta_v \circ (\mu_v)_*^{-1})\right) \in \mathcal{O}'_V.$$

We now define $\delta'(e)$ for each edge $e = (v, w)$ where v is a cylinder. This defines δ' on all edges, since we know that $\delta'(e)$ and $\delta'(\bar{e})$ must be equal. We define

$$\delta'(e) := \Bigg(\delta(e), \mathrm{QI}(Z_Q, \mathcal{P}_Q) \cdot \Big(\delta_v \circ (\mu_v)_*^{-1}, (\mu_v)_*(\mathcal{S}_e)\Big), \\ \mathrm{QI}(Z_Q, \mathcal{P}_Q) \cdot \Big(\delta_w \circ (\mu_w)_*^{-1}, (\mu_w)_*(\mathcal{S}_{\bar{e}})\Big)\Bigg) \in \mathcal{O}'_E.$$

The decoration δ' is independent of the choice of maps $\{\mu_v\}_{v \in VT_c}$.

Unlike vertex refinement in [CM17], we don't keep track of partial orientations, which make no sense if edge spaces aren't two-ended. In the definition of $\delta'(e)$, we also need to ensure that there are decoration preserving relative quasi-isometries at both the initial and terminal vertices of e. This is not needed in [CM17] because cylinders are two-ended, thus every edge is incident to exactly one non-elementary vertex space.

Starting with the naïve decoration, we iteratively perform neighbour and vertex refinement till this process stabilizes, i.e. the partitions induced by successive refinements do not get strictly finer. This will always happen since the partition of vertices and edges induced by any geometric decoration will necessarily be coarser than the partition into G-orbits.

We suppose $\delta : T_c \to \mathcal{O}$ is a decoration that is stable under vertex refinement. Suppose there are edges $e, f \in ET_c$ such that $\delta(e) = \delta(f)$, and that $v = \iota e$ and $w = \iota f$ are both cylinders. Let $Q = [[X_v, \mathcal{P}_v]]$. Since the decoration is stable under vertex refinement, there is some relative quasi-isometry $\psi \in QI(Z_Q, \mathcal{P}_Q)$ such that

$$\Big(\delta_w \circ (\mu_w)_*^{-1}, (\mu_w)_*(f)\Big) = \Big(\delta_v \circ (\mu_v)_*^{-1} \circ \psi_*^{-1}, (\psi_* \circ (\mu_v)_*)(e)\Big).$$

Let $\phi_v := \mu_w^{-1} \circ \psi \circ \mu_v : (X_v, \mathcal{P}_v) \to (X_w, \mathcal{P}_w)$, where μ_w^{-1} is a coarse inverse to μ_w. Then $f = (\phi_v)_*(e)$ and $\delta_w \circ (\phi_v)_* = \delta_v$. This observation illustrates part of the proof of the following proposition.

Proposition 4.79 (See Propositions 3.5, 5.9 and 5.12 of [CM17]). *Let T be the JSJ tree of cylinders of a group G. Suppose $\delta : T_c \to \mathcal{O}$ is a geometric decoration that is stable under neighbour and vertex refinement. Then the following hold:*

(1) Whenever $\delta(v) = \delta(w)$ for vertices of T_c, there is a decoration preserving tree automorphism χ such that $\chi(v) = w$. Moreover, there is a relative quasi-isometry $\phi_v : (X_v, \mathcal{P}_v) \to (X_w, \mathcal{P}_w)$ such that $(\phi_v)_* = \chi|_{lk(v)}$.

(2) Suppose also that $\delta(e) = \delta(f)$ for edges e, f of T_c, with $\iota e = v$ and $\iota f = w$. Then we may choose χ as above such that $\chi(e) = f$ and $(\phi_v)_*(e) = f$.

We wish to compare two JSJ trees of cylinders decorated with different sets of ornaments. To do this, Cashen and Martin [CM17] define a matrix known as the *structure invariant* which encodes the adjacency data of different ornaments.

Definition 4.80. Given a decoration $\delta : T_c \to \mathcal{O}$ that is stable under neighbour refinement, the structure invariant $S(T_c, \delta, \mathcal{O})$ is an $|\mathrm{im}(\delta)| \times |\mathrm{im}(\delta)|$ matrix with entries in $\mathbb{N} \cup \{\infty\}$. Choosing an ordering of $\mathrm{im}(\delta) = \{o_1, \ldots, o_{|\mathrm{im}(\delta)|}\}$, the (i, j)-entry of $S(T_c, \delta, \mathcal{O})$ counts the number of copies of vertices or edges with ornament o_j that are adjacent to a vertex or edge with ornament o_i.

This is well-defined because δ is stable under neighbour refinement. (There is a more general definition in [CM17] for a decoration that is not stable under neighbour refinement.)

Proposition 4.81 ([CM17, Theorem 5.13]). *Suppose we have two trees T_c and T'_c, obtained by taking trees of cylinders of trees T and T' satisfying our standing assumptions. Assume T_c and T'_c are endowed with geometric decorations $\delta_0 : T_c \to \mathcal{O}_0$ and $\delta'_0 : T'_c \to \mathcal{O}_0$. We perform neighbour and vertex refinement iteratively until they stabilize, obtaining decorations $\delta : T_c \to \mathcal{O}$ and $\delta' : T'_c \to \mathcal{O}'$. Then (T_c, δ_0) and (T'_c, δ'_0) are equivalent if and only if the following hold.*

(1) There is a bijection $\beta : \mathrm{im}(\delta) \to \mathrm{im}(\delta')$ such that

$$\delta_0 \circ \delta^{-1} = \delta'_0 \circ (\delta')^{-1} \circ \beta.$$

(2) After ordering $\mathrm{im}(\delta)$ and $\mathrm{im}(\delta')$ so that β is order preserving, the structure invariants $S(T_c, \delta, \mathcal{O})$ and $S(T'_c, \delta', \mathcal{O}')$ are equal.

(3) For every ornament $o \in \mathcal{O}$ with a vertex $v \in \delta^{-1}(o)$, there exists a vertex $v' \in (\delta')^{-1}(\beta(o))$ and a relative quasi-isometry $\phi_v : (X_v, \mathcal{P}_v) \to (X_{v'}, \mathcal{P}_{v'})$ such that the decorations $\delta'_v \circ (\phi_v)_*$ and $\beta \circ \delta_v$ on the peripheral structure $\mathcal{P}_v$ are equal.

Using the notation of Proposition 4.81, we may identify $\mathrm{im}(\delta)$ with $\mathrm{im}(\delta')$ via the bijection β. By combining Propositions 4.79 and 4.81, we thus deduce the following:

Corollary 4.82. *Let (T_c, δ_0) and (T'_c, δ'_0) be decorated JSJ trees of cylinders as above. Then there exist decorations $\delta : T_c \to \mathcal{O}$ and $\delta' : T'_c \to \mathcal{O}$ such that $\mathrm{im}(\delta) = \mathrm{im}(\delta')$ and the following hold:*

(1) If $\delta(v) = \delta'(w)$ for vertices $v \in VT_c$ and $w \in VT'_c$, then there exists a decoration preserving tree isomorphism $\chi : T_c \to T'_c$ such that $\chi(v) = w$. Moreover, there is a relative quasi-isometry $\phi_v : (X_v, \mathcal{P}_v) \to (X'_w, \mathcal{P}_w)$ such that $(\phi_v)_* = \chi|_{lk(v)}$.

(2) Suppose also that $\delta(e) = \delta(f)$ for edges $e \in ET_c$ and $f \in ET'_c$, with $\iota e = v$ and $\iota f = w$. Then we may choose χ as above such that $\chi(e) = f$ and $(\phi_v)_*(e) = f$.

Chapter 5

Quasi-isometry classification of certain RAAGs

In this chapter, we apply some of the methods developed in Chapter 4 to obtain a complete quasi-isometry classification of certain classes of RAAGs, generalising the classification of tree RAAGs by Behrstock and Neumann [BN08]. In Theorem 5.17, we show that the JSJ tree of cylinders of a RAAG over free abelian subgroups of arbitrary fixed rank is a quasi-isometry invariant of the RAAG. This allows one to determine the quasi-isometry class of RAAGs that cannot be determined by any other method.

We then restrict to the case of cyclic splittings and obtain a complete quasi-isometry classification for a certain class of RAAGs. The results stated here are a special case of the results I prove in [Mar18a]. The proof of Theorem 5.28 in Section 5.5 showcases the main ingredients needed to prove the results of [Mar18a], but is significantly simpler and less technical. This is because the stabilizer of rigid vertices in the JSJ tree of cylinders are assumed to be abelian, thus one doesn't need to keep track of *stretch factors*.

5.1 Preliminaries on RAAGs

A subgraph $\Gamma' \subseteq \Gamma$ is *full* if whenever vertices $v, w \in V\Gamma'$ are joined by an edge in Γ, they are joined by an edge of Γ'. The *link* $\mathrm{lk}(v)$ of a vertex v is defined to be the full subgraph whose vertex set consists of vertices adjacent to v. The *star* $\mathrm{star}(v)$ of a vertex v is the full subgraph whose vertex set consists of v and all adjacent vertices. When ambiguous, we write $\mathrm{lk}^\Gamma(v)$ and $\mathrm{star}^\Gamma(v)$ to denote the link and star of some $v \in V\Gamma$.

The Salvetti complex $S(\Gamma)$ of a RAAG $A(\Gamma)$ is defined as follows: there is a single vertex w and a one-edge loop e_v is attached to w for every $v \in V\Gamma$. We attach an n-torus to the edges $e_{v_1}, \dots, e_{v_n}$ whenever $\{v_1, \dots, v_n\}$ is a *clique*, i.e. whenever $\{v_1, \dots, v_n\}$ is the vertex set of a complete subgraph of Γ. The fundamental group of $S(\Gamma)$ is $A(\Gamma)$. In fact, $S(\Gamma)$ is a classifying space for $A(\Gamma)$.

Convention 5.1. If $\Gamma' \subseteq \Gamma$ is a full subgraph, we identify $A(\Gamma')$ with the subgroup of $A(\Gamma)$ generated by vertices of Γ', and identify $S(\Gamma')$ with the corresponding subcomplex of $S(\Gamma)$. Such a subgroup is said to be *special*.

Let $X(\Gamma)$ be the universal cover of $S(\Gamma)$ with covering map $p : X(\Gamma) \to S(\Gamma)$. Note that $X(\Gamma)$ is a $CAT(0)$ cube complex. A *standard k-flat* in $X(\Gamma)$ is a k-flat that covers a standard k-torus in $S(\Gamma)$, i.e. a k-torus corresponding to a k-clique of Γ. A *standard geodesic* in $X(\Gamma)$ is a standard 1-flat. A *standard subcomplex* of $X(\Gamma)$ is a connected component of $p^{-1}(S(\Gamma'))$ for some full subgraph $\Gamma' \subseteq \Gamma$.

We may identify the 1-skeleton of $X(\Gamma)$ with the Cayley graph of $A(\Gamma)$ with respect to the generating set $V\Gamma$. Then edges of $X(\Gamma)$ are labelled by vertices of Γ. Let $V_e \in V\Gamma$ be the label associated to some edge $e \in X(\Gamma)$. We define the *support* V_Y of a subcomplex $Y \subseteq X_\Gamma$ to be $\{V_e \mid e \text{ is an edge in } Y\}$. If Y is a standard subcomplex associated to a full subgraph $\Gamma' \subseteq \Gamma$, then $V_Y = V\Gamma'$. Given $v \in V\Gamma$ and a standard geodesic $l \subseteq X(\Gamma)$, we say that l is a v-geodesic if $V_l = \{v\}$. More generally, given a clique $K \subseteq \Gamma$, we say that Y is a *K-flat*, if Y is a standard flat with $V_Y = K$.

It is shown in the proof of Proposition 3.1 of [GW16] that every finite simplicial graph Γ has maximal join decomposition $\Gamma_1 \circ \cdots \circ \Gamma_k$, where Γ_1 is a maximal clique factor and no Γ_i splits as a join for $k \geq 2$. This gives direct product decompositions $X(\Gamma) = \mathbb{E}^n \times X(\Gamma_2) \times \cdots \times X(\Gamma_k)$ and $A(\Gamma) = \mathbb{Z}^n \times A(\Gamma_2) \cdots \times A(\Gamma_n)$. We call this the *de Rahm decomposition* of a RAAG. As noted by Huang [Hua17a], work of [BC12],[ABD$^+$13] and [KKL98] imply the de Rahm decomposition is stable under quasi-isometry in the following sense:

Theorem 5.2 ([Hua17a, Theorem 2.9])**.** *Let Γ and Λ be finite simplicial graphs, and let $X(\Gamma) = \mathbb{E}^n \times X(\Gamma_2) \times \cdots \times X(\Gamma_k)$ and $X(\Lambda) = \mathbb{E}^m \times X(\Lambda_2) \times \cdots \times X(\Lambda_l)$ be the de Rahm decompositions. Suppose there is a quasi-isometry $f : X(\Gamma) \to X(\Lambda)$. Then $n = m$, $k = l$, and after reordering $\Lambda_2, \dots, \Lambda_l$, there are quasi-isometries $f_i : X(\Gamma_i) \to X(\Lambda_i)$ such that the following diagram, whose vertical maps are projections, commutes up to some uniform error D.*

$$\begin{array}{ccc} X(\Gamma) & \xrightarrow{f} & X(\Lambda) \\ \downarrow & & \downarrow \\ X(\Gamma_i) & \xrightarrow{f_i} & X(\Lambda_i) \end{array}$$

Moreover, if $n = m = 0$, then f is uniformly close to the product of quasi-isometries $f_2 \times \cdots \times f_k$.

This theorem reduces the quasi-isometry classification of RAAGs to the quasi-isometry classification of RAAGs that don't split as a direct product.

In Section 4.4.1, we defined peripheral structures. In full generality, these peripheral structures might be quite hard to describe and work with. However in the case of RAAGs, they often have a simple combinatorial description given in terms of the defining graph.

Example 5.3. Let Γ be a finite simplicial graph. For each standard geodesic l in $X(\Gamma)$, we define

$$\mathcal{S}_l = \{l' \subseteq X(\Gamma) \mid l' \text{ is a standard geodesic parallel to } l\}.$$

Let $B \subseteq V\Gamma$ be an arbitrary set of vertices. We define a peripheral structure $\mathcal{P}_B$ of $X(\Gamma)$ to be

$$\mathcal{P}_B := \{\mathcal{S}_l \mid \ l \text{ is a } v\text{-geodesic for some } v \in B\}.$$

In other words, $\mathcal{P}_B$ consists of the set of parallelism classes of standard geodesics labelled by elements of B. We also let $\mathcal{P}_B$ denote the corresponding peripheral structure on $A(\Gamma)$ obtained by restricting to vertex sets of standard geodesics as above.

Example 5.4. Let Γ be the graph consisting of two vertices a and b joined by a single edge. Then $A(\Gamma) = \langle a, b \mid [a, b]\rangle \cong \mathbb{Z}^2$. Let $\mathcal{P}_1$, $\mathcal{P}_2$ and $\mathcal{P}_3$ be the peripheral structures defined in Example 4.71. Then $\mathcal{P}_1 = \mathcal{P}_{\{a\}}$, $\mathcal{P}_2 = \mathcal{P}_{\{b\}}$ and $\mathcal{P}_3 = \mathcal{P}_{\{a,b\}}$

5.2 The JSJ tree of cylinders of a RAAG

In this section, we construct the JSJ tree of cylinder over abelian subgroups of fixed rank and show that it is *visual*, i.e. it can be read from the defining graph. We then show it is a quasi-isometry invariant. Unlike in [GH17], who describe a JSJ decomposition of a RAAG over *all* abelian subgroups, we restrict to splittings over $\mathbb{Z}^n$ for some n, under the assumption the RAAG doesn't split over $\mathbb{Z}^k$ for any $k < n$. In doing so, it is easier to deduce quasi-isometry invariance of the JSJ tree of cylinders.

There is no loss of generality in doing this as we can then work inductively: we first examine the JSJ tree of cylinders over $\mathbb{Z}$, then the JSJ tree of cylinders over $\mathbb{Z}^2$ etc.

The following is a reformulation of the main result of [GH17]:

Theorem 5.5 ([GH17]). *We fix $n \geq 1$ and let Γ be a finite connected simplicial graph. Suppose that $A(\Gamma)$ does not split over $\mathbb{Z}^k$ for any $k < n$. Then $A(\Gamma)$ splits over $\mathbb{Z}^n$ if and only if either:*

- *Γ is a complete graph on $n+1$ vertices;*
- *Γ is separated by a clique of size n, i.e. there is a complete subgraph K with n vertices such that $\Gamma \backslash K$ is not connected.*

Let $n \geq 1$ and let Γ be a finite connected simplicial graph that is not complete, and not separated by a k-clique for any $k < n$. We now describe three visual graph of group decompositions of $A(\Gamma)$: $\mathcal{G}_n(\Gamma)$, $\mathcal{G}'_n(\Gamma)$ and $C_n(\Gamma)$; the latter two of these are a JSJ decomposition and the JSJ tree of cylinders decomposition over $\mathbb{Z}^n$ subgroups respectively. Figure 5.1 gives examples of these graphs of groups in the $n = 1$ case.

It is easy to see that a single separating clique K naturally induces a splitting of $A(\Gamma)$ as an amalgamated free product: if $R_1, \ldots, R_m$ are connected components of

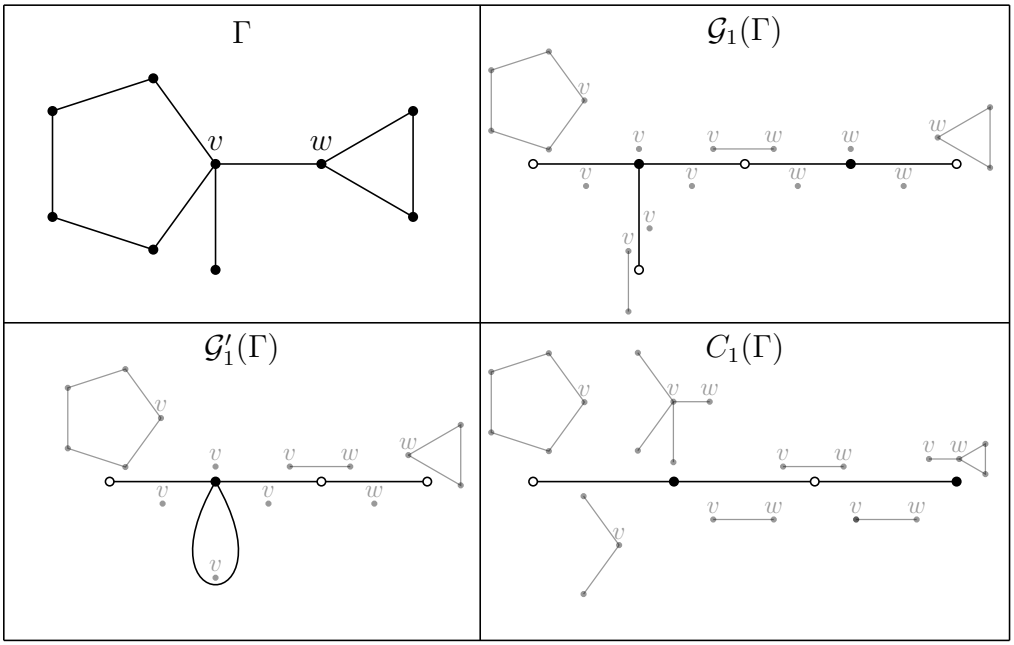


Figure 5.1: A graph Γ and the corresponding visual graph of groups decompositions $\mathcal{G}_1(\Gamma)$, $\mathcal{G}'_1(\Gamma)$ and $C_1(\Gamma)$. We label each vertex and edge with the full subgraph of Γ that generates the vertex or edge group, shown in grey.

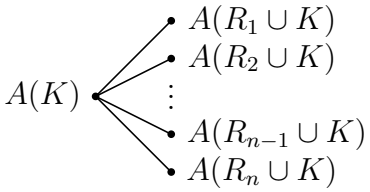


Figure 5.2: The graph of groups decomposition of $A(\Gamma)$ induced by a separating clique K, all of whose edge groups are $A(K)$.

$\Gamma \backslash K$, then there is a graph of groups as shown in Figure 5.2. Each subgraph $R_i \cup K$ cannot contain a separating k-clique for any $k < n$.

Let C_n be the collection of separating n-cliques in Γ. If $K, K' \in C_n$ are distinct cliques, then there is a unique component R_i of $\Gamma \backslash K$ such that $K' \subseteq K \cup R_i$, so separating cliques can't cross. Thus K' must be a separating clique of $K \cup R_i$. As before, we see that there is a component S of $(K \cup R_i) \backslash K'$ such that K is contained in $S \cup K'$. We are thus able to refine the graph of groups described above. We continue in this way and so obtain the following graph of groups.

Definition 5.6. Let $n \geq 1$ and let Γ be a finite connected simplicial graph that is not complete, and not separated by a k-clique for any $k < n$. Let C_n denote the set of separating n-cliques of Γ. We say a subgraph of Γ is *n-thick* if it is not separated by a k-clique for any $k \leq n$. Let B_n denote the set of maximal n-thick subgraphs of Γ. We define a bipartite graph of groups $\mathcal{G}_n(\Gamma)$ on the vertex set $B_n \sqcup C_n$ as follows:

- black vertices are cliques $K \in C_n$ with vertex group $A(K) \cong \mathbb{Z}^n$;
- white vertices are subgraphs $\Gamma' \in B_n$ with vertex group $A(\Gamma')$;
- $K \in C_n$ and $\Gamma' \in B_n$ are joined by an edge if and only if K is a subgraph of Γ'. The edge group of this edge is $A(K)$.

The inclusions of edge groups into adjacent vertex groups are as described in Convention 5.1, and $\pi_1(\mathcal{G}_n(\Gamma)) = A(\Gamma)$.

Lemma 5.7. *Let Γ be as above. Then the underlying graph of $\mathcal{G}_n(\Gamma)$ is a tree.*

Proof. Let Γ_n denote the underlying graph of the graph of groups $\mathcal{G}_n(\Gamma)$. Since Γ is connected, it is clear that Γ_n is also connected. Suppose for contradiction that Γ_n contains a cycle. This corresponds to a sequence $Y_1, K_1, Y_2, K_2, \ldots, Y_m, K_m$, all of whose elements are distinct, where $m > 1$, each $Y_i \in B_n$, each $K_i \in C_n$, and each K_i is contained in Y_i and Y_{i+1} with indices taken modulo m.

We now claim that $B := Y_1 \cup Y_2 \cup \cdots \cup Y_m$ is n-thick, contradicting our assumption that Y_1 is a maximal n-thick subgraph. Let K be a k-clique for some $k \leq n$. We show that $B \backslash K$ is connected. We may pick i such that $K \neq K_j$ for all $j \neq i$. We write $B \backslash K$ as

$$(Y_{i+1} \backslash K) \cup (Y_{i+2} \backslash K) \cup \cdots \cup (Y_m \backslash K) \cup (Y_1 \backslash K) \cup \cdots \cup (Y_i \backslash K). \tag{5.1}$$

Since each Y_j is n-thick, each $Y_j \backslash K$ is connected. Moreover, as $K \neq K_j$ for $j \neq i$, it follows that $Y_j \backslash K$ and $Y_{j+1} \backslash K$ must share a vertex, where indices are taken modulo m. It follows from (5.1) that $B \backslash K$ is connected. □

We now explain how to determine the JSJ tree of cylinders over $\mathbb{Z}^n$. Let $\mathcal{E}$ be the set of all subgroups of $A(\Gamma)$ that are isomorphic to $\mathbb{Z}^n$. An $\mathcal{E}$-tree is an $A(\Gamma)$-tree with edge stabilizers in $\mathcal{E}$. Unlike [GH17], we do not restrict to the relative setting of vertex-elliptic splittings. We first recall the following lemma from [GH17]:

Lemma 5.8 ([GH17, Lemma 1.2])**.** *Suppose a RAAG $A(\Gamma)$ acts on a tree T with abelian edge stabilizers. If some $v \in V\Gamma$ acts hyperbolically, then* $\mathrm{star}(v)$ *is a clique and some clique $K \subseteq \mathrm{lk}(v)$ separates v and $\Gamma \backslash \mathrm{star}(v)$.*

Corollary 5.9. *(1) For every $K \in C_n$, $A(K)$ is universally elliptic, i.e. it acts elliptically on any $\mathcal{E}$-tree.*

(2) If for some $\Gamma' \in B_n$, $A(\Gamma')$ is not universally elliptic, then Γ' is itself a clique. Moreover, if Γ' is a proper subgraph of Γ, then Γ' is the complete graph on the $n+1$ vertices $\{v\} \cup K$, where $K \in C_n$ and v is not a vertex of any $K' \in C_n$.

Proof. (1): Suppose $K \in C_n$ and T is an $\mathcal{E}$-tree. As no k-clique separates Γ for $k < n$, we deduce that for every $v \in VK$, there exist two vertices that are adjacent to v and lie in different components of $\Gamma \backslash K$; thus $\mathrm{star}(v)$ cannot be a clique. Hence for every $v \in VK$, we deduce from Lemma 5.8 that v cannot act hyperbolically on T. Thus for each $v, w \in VK$ we deduce $\mathrm{Fix}(v) \cap \mathrm{Fix}(w) \neq \emptyset$ (see Lemma 1.1 of [GH17]), so Helly's theorem ensures that $A(K)$ acts elliptically on T.

(2): If $A(\Gamma')$ doesn't act elliptically on an $\mathcal{E}$-tree T, it splits over $\mathbb{Z}^k$ for some $k \leq n$. As $\Gamma' \in B_n$, it has no separating k-clique for $k \leq n$, thus by Theorem 5.5, it is itself a clique. If Γ' is a proper subgraph of Γ, it properly contains some $K \in C_n$, so has at least $n+1$ vertices. Since $\mathbb{Z}^r$ splits over $\mathbb{Z}^s$ only if $r = s+1$, Γ' must have exactly $n+1$ vertices, thus has vertex set $K \cup \{v\}$ for some $v \notin K$. As v acts hyperbolically on T, it follows from (1) that v cannot be a vertex of any $K' \in C_n$. □

If some white vertex group $A(\Gamma')$ is not universally elliptic, Γ' has vertex set $\{v\} \cup K$ as above, so the $\mathcal{G}_n(\Gamma)$ can be refined by adding a 1-edge loop with edge group is $A(K)$, whose incident vertex group is also $A(K)$. Doing this for all vertices that are not universally elliptic, we obtain a graph of groups $\mathcal{G}'_n(\Gamma)$, all of whose vertex groups are universally elliptic. Thus $\mathcal{G}'_n(\Gamma)$ is a JSJ decomposition by [GL17, Lemma 3.1].

Lemma 5.10. *Let $T_n(\Gamma)$ and $T'_n(\Gamma)$ denote the Bass-Serre trees associated to $\mathcal{G}_n(\Gamma)$ and $\mathcal{G}'_n(\Gamma)$. Then $T'_n(\Gamma)_{\text{cyl}} = T_n(\Gamma)_{\text{cyl}}$. In particular, $T_n(\Gamma)_{\text{cyl}}$ is the JSJ tree of cylinders of $A(\Gamma)$ over $\mathcal{E}$.*

Proof. As $\mathcal{G}'_n(\Gamma)$ refines $\mathcal{G}_n(\Gamma)$, we can obtain $T'_n(\Gamma)$ from $T_n(\Gamma)$ by "blowing up" certain vertices $v \in VT_n(\Gamma)$ by adding in trees T_v. All edge stabilizers of the new tree T_v are commensurable to edge stabilizers incident to the original vertex v, so the entire blowing up procedure takes place within a cylinder. This doesn't affect cylinder stabilizers and their intersection patterns, so the tree of cylinders is unchanged. □

To understand cylinders in the JSJ tree, we use the following lemma.

Lemma 5.11. *Let $K \in C_n$ be separating n-clique in Γ. Let e be an edge of $T_n(\Gamma)$ with stabilizer $A(K) \cong \mathbb{Z}^n$. Then an edge f of $T_n(\Gamma)$ lies in the same cylinder as e if and only if $f = ge$ for some $g \in A(\text{star}(K))$.*

Proof. By the definition of $\mathcal{G}_n(\Gamma)$ and Proposition 4.8, edge spaces of $T_n(\Gamma)$ may be identified with cosets of $\langle A(K) \rangle_{K \in C_n}$. Identifying the 1-skeleton of $X(\Gamma)$ with the Cayley graph of G, we identify edge spaces with standard n-flats. Two edges lie in the same cylinder if and only if the associated standard flats are coarsely equivalent.

Let Y be a standard K-flat in $X(\Gamma)$, i.e. a standard flat labelled by K. Lemma 2.10 of [Hua17b] ensures that two standard flats are coarsely equivalent if and only if they are parallel in the CAT(0) metric. Moreover, Corollary 3.2 of [Hua17a] tells us that two standard flats are parallel if and only if they have the same support. We therefore deduce that a standard flat is parallel to Y if and only if it is a translate of Y. Lemma 3.4 of [Hua17a] says the parallel set of Y canonically splits as $Y \times X(\text{lk}(K))$. We thus deduce that gY is parallel to Y if and only if $g \in A(\text{star}(K))$. □

We now describe the visual construction of the JSJ tree of cylinders of $A(\Gamma)$, mirroring the definition of $\mathcal{G}_n(\Gamma)$.

Definition 5.12 (cf. Definition 5.6)**.** Let $n \geq 1$ and let Γ be a finite connected simplicial graph that is not complete, and not separated by a k-clique for any $k < n$. Let C_n denote the set of separating n-cliques of Γ. Let $D_n \subseteq B_n$ denote the set of

maximal n-thick subgraphs of Γ that either contain at least two elements of C_n, or are not contained in $\mathrm{star}(K)$ for any $K \in C_n$. We define a bipartite graph of groups $C_n(\Gamma)$ as follows.

(1) Black vertices are cliques $K \in C_n$ with vertex group $A(\mathrm{star}^{\Gamma}(K))$.

(2) White vertices are subgraphs $\Gamma' \in D_n$ with vertex group $A(\Gamma')$.

(3) A black vertex $K \in C_n$ and white vertex $\Gamma' \in D_n$ are joined by an edge if and only if K is a subgraph of Γ'. The corresponding edge group is $A(\mathrm{star}^{\Gamma'}(K)) = A(\Gamma') \cap A(\mathrm{star}^{\Gamma}(K)) = A(\{K\}) \times A(\mathrm{lk}^{\Gamma'}(K))$.

The inclusions of edge groups into adjacent vertex groups are as described in Convention 5.1, and $\pi_1(C_n(\Gamma)) = A(\Gamma)$.

The argument used in the proof of Lemma 5.7 demonstrates that the underlying graph of $C_n(\Gamma)$ is a tree. We also remark that if $A(\Gamma')$ is a white vertex group of $\mathcal{G}_n(\Gamma)$ that lifts to a vertex $v \in VT_n(\Gamma)$, then v is contained in at least two cylinders of $T_n(\Gamma)$ precisely when $\Gamma' \in D_n$. It follows from Lemma 5.11 that the Bass-Serre tree of $C_n(\Gamma)$ is indeed the tree of cylinders of $A(\Gamma)$. We summarise as follows:

Proposition 5.13. *Let Γ be a finite, connected simplicial graph. The Bass-Serre tree of $C_n(\Gamma)$ is the JSJ tree of cylinders of $A(\Gamma)$ over $\mathcal{E}$. Black and white vertices of $C_n(\Gamma)$ correspond to cylindrical and rigid vertices respectively.*

Convention 5.14. We recall that for a G-tree T and any vertex $v \in VT$ and edge $e \in ET$, we can define the vertex space X_v and edge space X_e as in Section 4.1. Suppose we are given a RAAG $A(\Gamma)$, and that T is the Bass-Serre tree of $C_n(\Gamma)$. Each edge vertex and edge space is thus naturally associated to a coset of some $A(\Gamma')$ where Γ' is a full subgraph of Γ. Therefore, we equip each vertex and edge space with the CAT(0) metric. This doesn't change the coarse equivalence class or quasi-isometry type of vertex and edge spaces.

Definition 5.15. Each vertex stabiliser associated to the cylinder $gA(\mathrm{star}(K))$ splits as a product $gA(K)g^{-1} \times gA(\mathrm{lk}(K))g^{-1} \cong \mathbb{Z}^n \times gA(\mathrm{lk}(K))g^{-1}$. Similarly, each edge stabilizer $gA(\mathrm{star}^{\Gamma'}(K))g^{-1}$ splits as $gA(K)g^{-1} \times gA(\mathrm{lk}^{\Gamma'}(K))g^{-1} \cong \mathbb{Z}^n \times gA(\mathrm{lk}^{\Gamma'}(K))g^{-1}$. There is an associated decomposition of the vertex and edge spaces: cylindrical vertex spaces split as $X_v = \mathbb{E}^n \times Y_v = gX(K) \times gX(\mathrm{lk}(K))$ and edge spaces split as $X_e = \mathbb{E}^n \times Y_e = gX(K) \times gX(\mathrm{lk}^{\Gamma'}(K))$. We call this the *cylindrical decomposition* of a cylindrical vertex space or an edge space.

A consequence of working with the CAT(0) metric is that each edge map $\alpha_e : X_e \to X_{\overline{e}}$ splits as a product of cubical isomorphisms $(\mathbb{E}^n, Y_e) \to (\mathbb{E}^n, Y_{\overline{e}})$.

Remark 5.16. Using Definitions 5.6 and 5.12, we can give an explicit description of elements in the peripheral structure of a vertex space in the JSJ tree of cylinders. Let $e = gA(\mathrm{star}^{\Gamma'}(K))$ be an edge in the Bass-Serre tree of $C_n(\Gamma)$ adjacent to a rigid vertex $gA(\Gamma')$ and a cylindrical vertex $g(\mathrm{star}^{\Gamma}(K))$. Let $T_n(\Gamma)$ be the Bass-Serre tree of $\mathcal{G}_n(\Gamma)$, and let Y be the cylinder in $T_n(\Gamma)$ corresponding to the vertex $g(\mathrm{star}^{\Gamma}(K))$ in the tree of cylinders. The edges of $T_n(\Gamma)$ that are adjacent to the white vertex $gA(\Gamma')$ and are contained in the cylinder Y are precisely those of the form $gg'A(K)$ for some $g' \in A(\mathrm{lk}^{\Gamma'}(K))$. The vertex set of Y_e is $gA(\mathrm{lk}^{\Gamma'}(K))$ — see Definition 5.15. Thus elements of $\mathcal{S}_e$ correspond precisely to subsets of X_e of the form $\mathbb{E}^1 \times \{y\}$ for some vertex $y \in Y_e$.

Having described the JSJ tree of cylinders and the peripheral structures of vertex and edge spaces, we will show quasi-isometry invariance of the JSJ tree of cylinders. This gives lots of combinatorial data that is preserved by quasi-isometries, and one is often able to distinguish quasi-isometry classes of RAAGs using this theorem.

Theorem 5.17. *Fix $n \geq 1$. Let Γ, Γ' be finite simplicial graphs such that the associated RAAGs do not split over $\mathbb{Z}^k$ for any $k < n$ (see Theorem 5.5). Let T_{cyl} and T'_{cyl} be the JSJ trees of cylinders of $A(\Gamma)$ and $A(\Gamma')$ over the class of $\mathbb{Z}^n$ subgroups. Suppose that $f : A(\Gamma) \to A(\Gamma')$ is a quasi-isometry. Then there is a constant C and a unique tree isomorphism $f_* : T_{\mathrm{cyl}} \to T'_{\mathrm{cyl}}$ such that the following holds.*

For every vertex v of T, $f(X_v)$ has Hausdorff distance at most C from $X'_{f_*(v)}$, and v and $f_*(v)$ are of the same vertex type, i.e. either they are both cylinders or neither are cylinders. Moreover, f restricted to each vertex space X_v is close to a relative quasi-isometry $(X_v, \mathcal{P}_v) \to (X_{f_*(v)}, \mathcal{P}_{f_*(v)})$.

Much of the work needed to prove this has been done by Huang:

Theorem 5.18 ([Hua17a, Theorem 3.21])**.** *Let Γ be a finite simplicial graph with $K \subseteq \Gamma$ an n-clique. Suppose there do not exist vertices $v, w \in V\Gamma$ such that $w \in K$, $v \in \Gamma \backslash K$ and $\mathrm{lk}(w) \subseteq \mathrm{star}(v)$. Then for any finite simplicial graph Λ and quasi-isometry $f : X(\Gamma) \to X(\Lambda)$, there exists a constant A such that for any K-flat $Y \subseteq X(\Gamma)$, there is a standard n-flat $Y' \subseteq X(\Lambda)$ with $d_{\mathrm{Haus}}(f(Y), Y')) \leq A$.*

Proof of Theorem 5.17. It follows from the JSJ decomposition as stated in [DS00] that every codimension one subgroup of $A(\Gamma)$ (resp. $A(\Lambda)$) that is isomorphic to $\mathbb{Z}^n$ is conjugate into a black vertex group of $\mathcal{G}'_n(\Gamma)$ (resp. $\mathcal{G}'_n(\Lambda)$). We thus deduce that a

standard n-flat coarsely separates the ambient RAAG if and only if its support set is a separating n-clique of the defining graph. (This can also be seen geometrically by examining the universal cover of the Salvetti complex.)

Suppose K is a separating n-clique of Γ with $w \in K$ and $v \in \Gamma \backslash K$. Suppose R_1 and R_2 are distinct components of $\Gamma \backslash K$ and that $v \in R_1$. Note that $\mathrm{star}(v) \subseteq K \cup R_1$. Since we assumed that Γ has no separating k-cliques for $k < n$, there exists a vertex $u \in \mathrm{lk}(w) \cap R_2$, verifying that $\mathrm{lk}(w) \not\subseteq \mathrm{star}(v)$. Thus Theorem 5.18 tells us that f sends every edge space of $\mathcal{G}_1(\Gamma)$ to within uniform finite Hausdorff distance of an edge space of $\mathcal{G}_1(\Lambda)$ and vice versa. The result now follows from Theorem 4.74. □

5.3 Quasi-isometries of free products

In order to determine when cylindrical vertex spaces in a JSJ tree of cylinders are relatively quasi-isometric, we use methods of Papasoglu-Whyte [PW02]. Specifically, we need an understanding of why the free product $G_1 * G_2$ of two finitely generated, infinite groups is quasi-isometric to:

(I) $G_1 * G_1 * G_2$;

(II) $G_1' * G_2$, whenever G_1' is quasi-isometric to G_1;

(III) $G_1 * G_2 * \mathbb{Z}$.

This can be used to reduce the quasi-isometry classification of RAAGs to the quasi-isometry classification of one-ended factors in its Gruško decomposition.

We not only need to know that such quasi-isometries exist, but we need a constructive way of building these quasi-isometries. While the main arguments are based on [PW02], the terminology used here (protographs, special subgraphs, isolated vertices, depth etc.) is not found in [PW02].

Definition 5.19. A *based graph* (A, a_0) is a graph A with a distinguished vertex $a_0 \in A$ called the *basepoint*. Let $\mathcal{R} := \{(A, a_0), (B, b_0), \ldots, (E, e_0)\}$ be a set of finitely many connected, infinite based graphs of uniformly bounded valance, which we call *protographs*. The free product is built out of copies of the protographs called *special subgraphs*. Vertices that are not contained in a special subgraph are called *isolated vertices*. Given $\mathcal{R}$ as above, we define the *free product* $A * \cdots * E$ of graphs inductively as follows.

We define Γ^1 by taking the disjoint union of a vertex v_0 and copies of each based protograph $(A, a_0) \in \mathcal{R}$. We then join v_0 to each of the basepoints a_0 of the protographs

via an edge. We call each copy of a protograph a special subgraph, and call the vertex v_0 an isolated vertex. We define Γ^{i+1} as follows. For every vertex of Γ^i, say $a \in A$, that is not incident to an isolated vertex, we attach a copy of each of the based protographs in $\mathcal{R}\backslash\{(A, a_0)\}$ via their basepoints to an isolated vertex, say v_a, which is itself attached via an edge to a. This is done so that for every isolated vertex v_a of Γ^{i+1}, there is exactly one copy of every protograph attached to v_a.

Noting that $\Gamma^1 \subseteq \Gamma^2 \subseteq \dots$ is a nested sequence of graphs, we define the graph $\Gamma := \cup_{i\in\mathbb{N}}\Gamma^i$. We call each Γ^i the *depth i subgraph of* Γ. We endow Γ with the path metric in which every edge has length one.

The free product $A*B*C$ is illustrated schematically in Figure 5.3. This definition differs very slightly from the procedure given in [PW02], which doesn't include isolated vertices. It is shown in [PW02] that the choice of basepoints doesn't affect the construction (up to bi-Lipschitz equivalence), provided the special subgraphs satisfy a mild quasi-homogeneous property; this will be evident from the description of type II moves.

Remark 5.20. The above construction is motivated by the following. Suppose $G = G_1 * \dots * G_n$ is the free product of finitely many finitely generated infinite groups. For each i, let A_i denote the Cayley graph of G_i with respect to some finite generating set $S_i \subseteq G_i\backslash\{e\}$. Let $\mathcal{G}$ be a finite connected graph of groups as shown in Figure 5.4. Clearly $\pi_1(\mathcal{G}) = G$. We equip G with the word metric with respect to $S_1 \sqcup \dots \sqcup S_n$.

Let X be the tree of spaces built by blowing up the Bass-Serre tree of $\mathcal{G}$ as in Proposition 4.8. We observe that X is isomorphic (as a graph) to the free product of graphs $A_1 * \dots * A_n$. Each isolated vertex corresponds to a coset of the form $g\{e\}$.

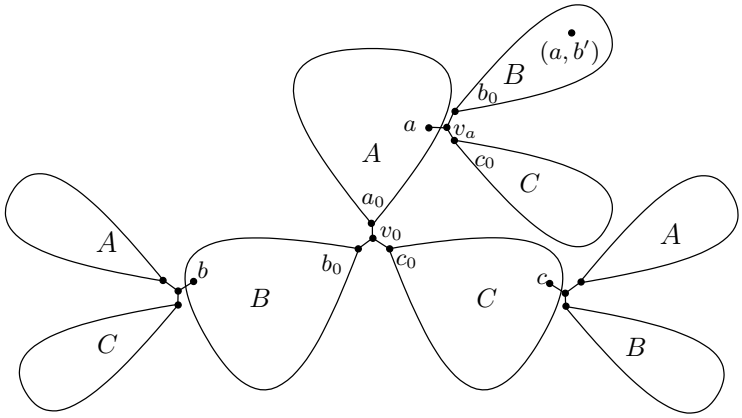


Figure 5.3: Schematic diagram of a part of the graph $A * B * C$. If we remove the component of $(A * B * C)\backslash\{v_a\}$ containing v_0, we obtain branch(v_a).

Each special subgraph that is a copy of the protograph A_i corresponds to a coset of the form gG_i. By Proposition 4.8, there is a quasi-isometry $f : G \to X$ such that $f(gG_i)$ gets mapped to within uniform finite Hausdorff distance of the corresponding special subgraph obtained from the protograph A_i.

We may represent each non-isolated vertex $x \in A * \cdots * E$ by an ordered tuple $(v_1, \ldots, v_n)$ of vertices of protographs, chosen such that every adjacent pair v_i and v_{i+1} lie in different protographs, and that v_i is not a basepoint for $i < n$. The first $(n-1)$ elements of the tuple encodes the unique sequence of points at which a path from v_0 to x with no backtracking exits a special subgraph. For example, (a, b') denotes the vertex b' contained in the copy of the protograph B attached to (a) as shown in Figure 5.3. We denote this special subgraph by (a, B).

Using this notation, we use $*$ as a wildcard character to define maps between free products of graphs. For instance, $(a, *) \mapsto (a', *)$ implies that $(a, b) \mapsto (a', b)$, that $(a, b, c) \mapsto (a', b, c)$, and so on. We write $(a, B, *) \mapsto (a', B, *)$ to denote $(a, b, *) \mapsto (a', b, *)$ for all $b \in B$. We may also write $(-, a, *) \mapsto (-, a', *)$, where the $-$ symbol denotes some fixed initial string.

The *depth* of a vertex or special subgraph is defined to be the least i such that Γ^i contains that vertex or subgraph. We define the *ascending link* $\mathrm{lk}^+(v)$ of an isolated vertex v to be the set of special subgraphs attached to v, whose depth is equal to that of v. We define $branch(v)$ to be the complement in Γ of the component of $\Gamma \backslash \{v\}$ containing v_0.

A free product of graphs is a specific example of a more general construction known as a *tree of graphs*. This is defined like a free product of graphs, except we no longer assume there is exactly one copy of each protograph attached to each isolated vertex, and we no longer require the set $\mathcal{R}$ of protographs to be finite. For instance, an isolated vertex could have two copies of some protograph $(A, a_0) \in \mathcal{R}$ attached, or could have no copies attached. We do however, assume that there is a uniform bound to the number of special subgraphs attached to each isolated vertex. We extend the definition of depth and lk^+ to trees of graphs. However, we ensure that lk^+ includes

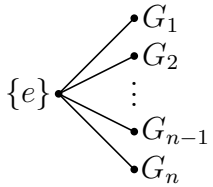


Figure 5.4: The graph of groups $\mathcal{G}$

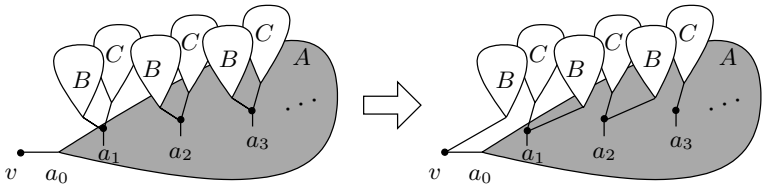


Figure 5.5: A type I move that takes place at v: sliding B onto v.

multiplicities. For example, if two copies of a protograph A are attached to v, these represent distinct elements of $\mathrm{lk}^+(v)$.

We also use the tuple notation to describe points on a tree of graphs. However, care must be taken when doing this, since it is not immediate when a tuple gives a well-defined vertex. For instance, the notation (a, b) is ambiguous if there are two copies of the protograph B attached to a.

We now describe three moves, which we call moves of type I, II and III, that map a tree of graphs to another tree of graphs. Each move will be a quasi-isometry of the underlying space. Moves of type I and II will induce a bijection between special subgraphs in the domain and target tree of graphs. By applying these moves iteratively, we can construct quasi-isometries of the form I, II and III described above.

Each of these moves is said to *take place at an isolated vertex* v. A move that takes place at v changes at least one special subgraph of $\mathrm{lk}^+(v)$, but acts trivially outside $\mathrm{branch}(v)$. In particular, if v is an isolated vertex of depth i, then a move that takes place at v will fix all vertices of depth at most $i - 1$.

Type I moves

A move of type I, also known as a *sliding move*, is shown pictorially in Figure 5.5. We describe a type I move that takes place at an isolated vertex v. We fix some special subgraph $(A, a_0) \in \mathrm{lk}^+(v_0)$ and choose an infinite geodesic ray $(a_0, a_1, \dots)$ in A, starting at the basepoint $a_0 \in A$.

Let v_{a_i} be the isolated vertex attached to a_i for each $i > 0$. We fix a based protograph (B, b_0). A type I move involves 'sliding' the edge attaching b_0 from the isolated vertex $v_{a_{i+1}}$ to v_{a_i} if $i > 1$, and to v if $i = 1$. This can be described in the

tuple notation by the following transformations:

$$\begin{aligned}(-, a_1, B, *) &\mapsto (-, \bar{B}, *)\\ (-, a_2, B, *) &\mapsto (-, a_1, B, *)\\ (-, a_3, B, *) &\mapsto (-, a_2, B, *)\\ &\vdots\end{aligned}$$

We write B and $\bar{B}$ to distinguish between the two copies of the photograph B attached to the isolated vertex v.

We denote the procedure described above as *sliding B onto v*. We can also perform the inverse move, which we call *sliding B off v* (providing there are copies of B attached to v and each v_{a_i} for $i > 0$). Either of these is called a type I move taking place at v. As shown in Lemma 1.2 of [PW02], this is a bi-Lipschitz equivalence and induces a bijection of special subgraphs.

Using sliding moves, one can show that $A * A * B$ is bi-Lipschitz equivalent to $A * B$. This is done by sliding a copy of A onto every isolated vertex. Provided we only perform a uniformly finite number of slides at each isolated vertex, the composition of infinitely many sliding moves will always be a bi-Lipschitz equivalence. This will be discussed in more detail at the end of this section.

Sliding moves are much more flexible than one might first think. For instance, it would appear in Figure 5.5 that one cannot slide A onto v. Although this cannot be done directly, it can be done stages: we first slide A onto v_{a_i} for every $i > 0$, then slide A onto v, and then slide A off v_{a_i} for every $i > 0$.[1] Although this trick will not work for an arbitrary tree of graphs, it will work for all trees of graphs that arise in practice. We will therefore be able to slide any protograph onto or off any isolated vertex.

Type II moves

We first observe the following: suppose (A, a_0) and (A', a_0') are based protographs and $\phi : A \to A'$ is a bi-Lipschitz equivalence with $\phi(a_0) = a_0'$. Then the transformation

$$(-, a, *) \mapsto (-, \phi(a), *)$$

is a bi-Lipschitz equivalence. Quasi-isometries are significantly less rigid than bi-Lipschitz equivalences; for instance, they are not necessarily bijections. Consequently, if ϕ is a quasi-isometry, the above transformation is not necessarily a quasi-isometry.

[1] We are implicitly assuming that the initial tree of graphs in Figure 5.5 is the free product of grahs $A * B * C$.

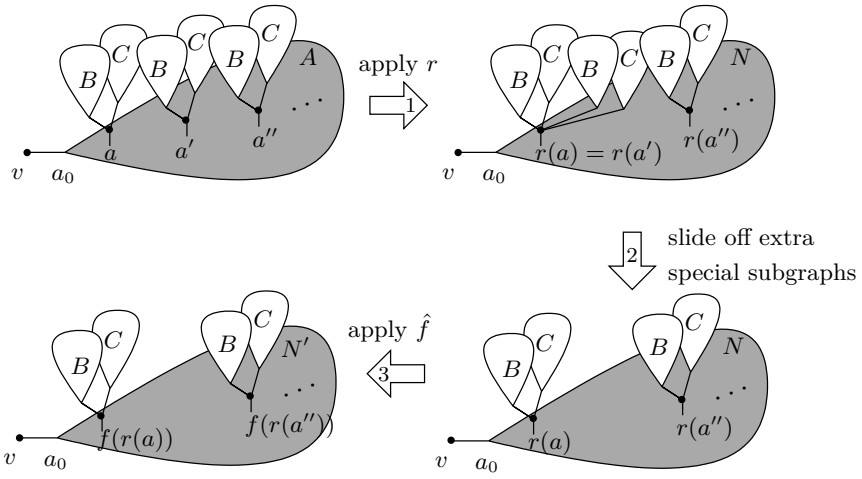


Figure 5.6: A type II move that takes place at v: 1) We apply r to the special subgraph A identifying isolated vertices v_a and $v_{a'}$ if $r(a) = r(a')$. 2) We slide special subgraphs off so that every point of N not equal to the basepoint has exactly one copy of B and C attached. 3) We apply $\hat{f}$ to the special subgraph N. 4) We apply steps 1) and 2) in reverse.

However, by combining with type I moves, we can nonetheless replace the special subgraph A in a tree of spaces with any graph A' quasi-isometric to A.

We now describe a type II moves that takes place at an isolated vertex v. Suppose $(A, a_0) \in \mathrm{lk}^+(v)$ and that $f : A \to A'$ is a quasi-isometry with $f(a_0) = a_0'$. We may choose nets $N \subseteq A$ and $N' \subseteq A'$ such that $f(N) = N'$ and $\hat{f} := f|N : N \to N'$ is a bi-Lipschitz equivalence. Let $r : A \to N$ and $s : A' \to N'$ be surjective maps that are close to the identity, chosen such that $r^{-1}(r(a_0))$ and $s^{-1}(s(a_0'))$ each contain exactly one point, namely a_0 and a_0' respectively.

We follow the procedure shown in Figure 5.6, which is also described in the proof of Theorem 0.1 of [PW02]. We first apply the transformations $(-, a, *) \mapsto (-, r(a), *)$. We do this so that whenever there are two points $a, a' \in A \backslash \{a_0\}$ such that $r(a) = r(a')$, special subgraphs in $\mathrm{lk}^+(v_a)$ and $\mathrm{lk}^+(v_{a'})$ get sent to special subgraphs in $\mathrm{lk}^+(v_{r(a)})$.

At each such isolated vertex $v_{r(a)}$, we slide off extra special subgraphs as necessary, ensuring that $\mathrm{lk}^+(v_{r(a)}) = \mathrm{lk}^+(v_a)$ for all $a \in A$. By our choice of r, we do not need to apply these sliding moves to the basepoint $a_0 \in A$. After doing this, we apply the transformation $(-, a, *) \mapsto (-, \hat{f}(a), *)$. This does give a well-defined map as $\hat{f}$ is a bi-Lipschitz equivalence. We have thus 'replaced' the special subgraph A with N'.

Finally, we perform the inverse of the operations we applied when replacing A with N, allowing us to replace N' with A'. For each $n' \in N$, we slide extra special subgraphs onto $v_{n'}$, allowing us to identify $\mathrm{lk}^+(v_{n'})$ with $\sqcup_{x \in s^{-1}(n')} \mathrm{lk}^+(v_x)$. For instance, suppose

$s^{-1}(n') = \{x, \overline{x}\}$ with $\mathrm{lk}^+(v_x) = \{B, C\}$ and $\mathrm{lk}^+(v_y) = \{B, C\}$. Then we slide each of B and C onto $v_{n'}$ so that $\mathrm{lk}^+(v_{n'}) = \{B, C, \overline{B}, \overline{C}\}$, where we write $\overline{B}$ and $\overline{C}$ to distinguish between the two copies of the protographs B and C in $\mathrm{lk}^+(v_{n'})$. We now apply the following transformations:

$$(-, n', B, *) \mapsto (-, x, B, *)$$
$$(-, n', C, *) \mapsto (-, x, C, *)$$
$$(-, n', \overline{B}, *) \mapsto (-, \overline{x}, B, *)$$
$$(-, n', \overline{C}, *) \mapsto (-, \overline{x}, C, *)$$

We send n' and v'_n to x and v_x respectively for some arbitrarily chosen $x \in s^{-1}(n)$. We apply this procedure for each $n' \in N$.

We call this a *type II move*. Type II moves allow one to show that $A * B$ is quasi-isometric to $A' * B$ if A is quasi-isometric to A'. We note the number of sliding moves we have to perform in the procedure is determined by the quasi-isometry constants of f.

This procedure is quite powerful, because it not only tells us that two trees of graphs are quasi-isometric, but allows us build a prescribed quasi-isometry of protographs into the quasi-isometry of the tree of graphs. More precisely, we have shown above that if $f : A \to A'$ is a quasi-isometry and A is a special subgraph of a tree of graphs Γ, we may apply a type II move to Γ, obtaining a tree of graphs Γ' and a quasi-isometry $\phi : \Gamma \to \Gamma'$ such that $\phi|_A$ is close to f.

This observation is essential in proving Theorem 5.28, since it allows us to define quasi-isometries such that the diagram of Proposition 4.9 commutes up to uniformly bounded error. We summarise with the following Lemma, which can be proved by applying type II moves to all vertices in $\mathrm{lk}^+(v)$.

Lemma 5.21. *Let Γ be a tree of spaces and $v \in \Gamma$ be an isolated vertex. Suppose there is a collection $\mathcal{R}$ of based protographs and a bijection $\chi : \mathrm{lk}^+(v) \to \mathcal{R}$ such that for every $A \in \mathrm{lk}^+(v)$, there is a basepoint preserving quasi-isometry $f_A : A \to \chi(A)$. Then we may apply moves of type II at v to obtain a quasi-isometry $f : \Gamma \to \Gamma'$ such that:*

(1) f is a graph isomorphism outside $\mathrm{branch}(v)$*;*

(2) $\mathrm{lk}^+(f(v)) = \mathcal{R}$*;*

(3) for every $A \in \mathrm{lk}^+(v)$, $f|_A$ is close to f_A.

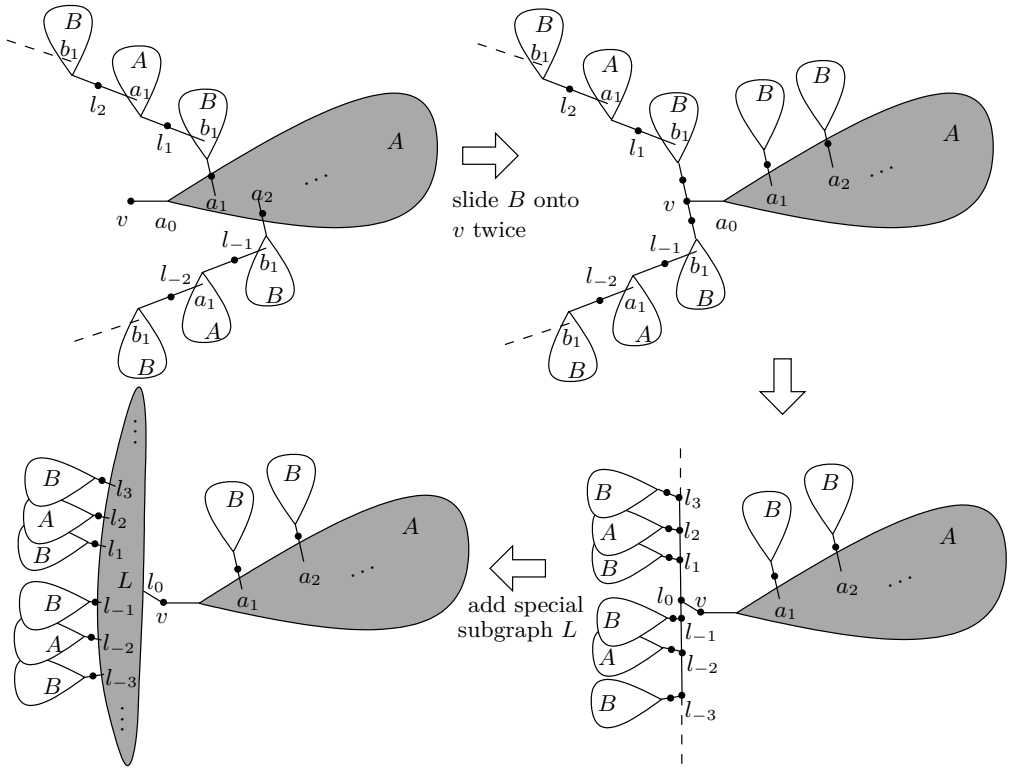


Figure 5.7: A type III move that takes place at v: we add the two-ended special subgraph L.

Type III moves

A *type III move* allows us to add or remove special subgraphs to trees of graphs. More specifically, only bi-infinite lines can be added or removed.

An example of a type III move that takes place at v is shown in Figure 5.7. We choose vertices $a_1, a_2 \in A$ and $b_1 \in B$ such that $d(a_0, a_1) = d(a_1, a_2) = d(b_0, b_1) = 1$. As shown in Figure 5.7, we define each l_i to be the isolated vertex adjacent to either $(-, a_1, b_1, a_1, b_1, \dots)$ or $(-, a_2, b_1, a_1, b_1, \dots)$. We slide B onto v twice, and then observe that

$$d(l_i, l_j) = 3|i - j|. \tag{5.2}$$

As shown in the third step of Figure 5.7, we can rearrange the l_i vertices so they lie along a single line $\{\dots, l_{-2}, l_{-1}, l_0, l_1, l_2, \dots\}$, adding vertices as in Figure 5.7 by subdividing edges. We now artificially declare L to be a special subgraph with basepoint l_0. This gives a tree of graphs with a two-ended special subgraph that doesn't correspond to a special subgraph of the starting tree of graphs. A type III move is defined to be either this procedure or its inverse.

A combination of type I and type III moves allows one to see that $A * B$ is quasi-isometric to $A * B * \mathbb{Z}$ (or more generally, to $A * B * F_n$). This is shown in Theorem 2.1 of [PW02]. It follows from Lemma 3.2 of [PW02] that one-ended special subgraphs can never be added or removed.

Combining moves

Each application of a move of type I, II or III provides us with a quasi-isometry $\gamma : \Gamma \to \Gamma'$ between trees of graphs. If for example, we want to show that $A * B$ is quasi-isometric to $A * B * B$, we need to apply infinitely many moves of type I, which means composing infinitely many quasi-isometries. We explain why this can be done.

We first consider the slide move shown in Figure 5.5, which we denote $\gamma : \Gamma \to \Gamma'$. We have already remarked that γ acts as an isometry outside brach(v). In fact, we observe that in some sense γ only alters distances near the special subgraph A. Although γ is not an isometry when restricted to $\Gamma \backslash A$, it is an isometry when restricted to components of $\Gamma \backslash A$. Similarly, type II and III moves also only alter distances 'locally'. This allows us to compose infinitely many moves, provided they satisfy the following condition:

Definition 5.22. A sequence of moves applied to a tree of graphs is said to be *uniformly locally finite* if there is a uniform bound on the number of moves that take place at an isolated vertex, and there is a uniform bound on the quasi-isometry constants of quasi-isometries used to define moves of type II.

Proposition 5.23. *The composition of a uniformly locally finite sequence of moves is a quasi-isometry.*

Proof. By the uniformly locally finite condition, we can assume each move is a (K, A)-quasi-isometry between trees of graphs for some uniform constants K and A. We claim that the composition of moves is then a (K', A')-quasi-isometry for suitable constants K' and A'. Suppose x and y are vertices in the tree of graphs that lie in special subgraphs (other cases can be dealt with similarly). Then there exists a unique sequence of points

$$x = x_0, y_0, v_1, x_1, y_1, v_2, x_2, \ldots, v_n, x_n, y_n = y \tag{5.3}$$

such that x_i and y_i are vertices in the same special subgraph, v_i is the isolated vertex joining y_{i-1} and x_i, and there is no backtracking, i.e. no v_i appears twice. The distance $d(x, y)$ can be computed by summing the distance between successive points of (5.3).

We first suppose that every move acts as a K-bi-Lipschitz equivalence on the tree of spaces. Indeed suppose For all x, y, each move only alters distances in a finite subsequence of (5.3). As only uniformly finitely many moves take place at each isolated vertex, we see that for any fixed finite substring of (5.3) of length R, there are at most $S = S(R)$ moves that alter distances in that substring. This is independent of x and y. Thus when we compose infinitely many moves, the resulting map is a K^N-bi-Lipschitz equivalence for some sufficiently large N.

We now explain how to deal with the general case where moves act as (K, A)-quasi-isometries. Quasi-isometries have additive errors (the 'A' factors). We need to show these cannot accumulate when applying a uniformly locally finite sequence of moves, e.g. if we apply 1000 moves, we need to show there is no $1000A$ factor. This does not follow directly from the above argument, which only shows multiplicative errors (the 'K' factors) don't accumulate.

It is sufficient to show this for maps such that

$$\dots, v_i, x_i, y_i, v_{i+1}, \dots$$

maps to

$$\dots, v_i, f(x_i), f(y_i), v_{i+1}, \dots,$$

where f is a (K, A)-quasi-isometry of the special subgraph containing x_i and y_i. Such maps occur in type II moves, e.g. step 1 of Figure 5.6. We can recover moves of type I, II and III by composing such maps with bi-Lipschitz equivalences.

There are two cases. First suppose $d(x_i, y_i) \geq 2KA$. Then

$$\frac{1}{2K} d(x_i, y_i) \leq \frac{1}{K} d(x_i, y_i) - A \leq d(f(x_i), f(y_i)) \leq K d(x_i, y_i) + A \leq (K + \frac{1}{2K}) d(x_i, y_i).$$

Now suppose $d(x_i, y_i) \leq 2KA$. Let $i > 0$. Since $d(v_i, x_i) = 1$, we see that

$$d(v_i, f(y_i)) = 1 + d(f(x_i), f(y_i)) \geq 1 \geq \frac{1}{2KA + 1} d(v_i, y_i).$$

If $x_i \neq y_i$, then $d(x_i, y_i) \geq 1$ (as x_i and y_i are vertices), so that $d(f(x_i), f(y_i)) \leq (K + A) d(x_i, y_i)$. The case $x_i = y_i$ is clear, since $0 = d(f(x_i), f(y_i)) = d(x_i, y_i) = 0$. Consequently, we see that additive errors cannot accumulate when we apply a uniformly locally finite sequence of moves. Indeed, f is essentially a biLipschitz equivalence provided x and y do not lie in the same special subgraph. We conclude by applying the above argument in the case where all moves are bi-Lipschitz equivalences. □

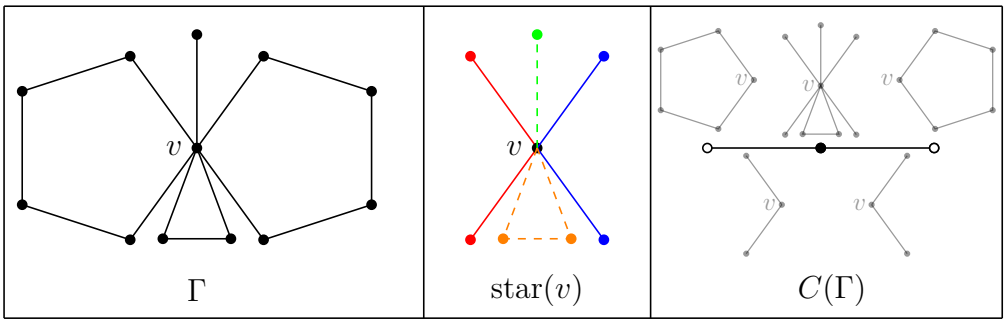


Figure 5.8: The decomposition of star(v) into coned-off blocks

5.4 Relative quasi-isometry type of cylinders

In this section, we analyse the relative quasi-isometry type of cylinder stabilizers in the JSJ tree of cylinders of a RAAG. For the rest of this thesis, we restrict to examining splittings over infinite cyclic subgroups. Thus when we talk about the JSJ tree of cylinders, it is assumed to be over the class of infinite cyclic subgroups.

As shown in Section 5.2, the cylindrical vertex groups of $C_1(\Gamma)$ are of the form $A(\text{star}(v))$ for some cut vertex $v \in V\Gamma$. We assume that Γ is connected and isn't of the form star(v) for some $v \in V\Gamma$, otherwise the graph $C_1(\Gamma)$ is just a point.

Recall the definition of D_1 and B_1 from Definitions 5.6 and 5.12. We let $\Gamma_1, \dots, \Gamma_n \in D_1$ and $\Gamma'_1, \dots \Gamma'_m \in B_1 \backslash D_1$ be all the subgraphs in B_1 that contain v. We define *blocks* $R_i := \text{lk}^{\Gamma_i}(v)$ and $S_i := \text{lk}^{\Gamma'_i}(v)$ for each i. We say the R_i are *peripheral blocks* and the S_i are *non-peripheral blocks*. A *coned-off block* is the join of a block with the cut vertex v. We call each $A(R_i)$ a *peripheral factor* and each $A(S_i)$ a *non-peripheral factor* of the cylinder.

The black (cylindrical) vertex groups of $C_1(\Gamma)$ are thus of the form

$$A(\text{star}(v)) \cong \mathbb{Z} \times (A(R_1) * \cdots * A(R_n) * A(S_1) * \cdots * A(S_m)). \tag{5.4}$$

Adjacent edge groups are precisely the $\mathbb{Z} \times A(R_i)$ terms for each peripheral factor $A(R_i)$. We note that $n \geq 1$ and $n + m \geq 2$.

Remark 5.24. As each non-peripheral block is connected, every $A(S_i)$ is either infinite cyclic or is one-ended.

We illustrate these ideas with Figure 5.8. There is a single cut vertex of Γ, labelled v. The graph of groups $C_1(\Gamma)$ has three vertices: two rigid (white) vertices that correspond to the left and right pentagons, and one cylindrical (black) vertex.

The cylinder has four coned-off blocks illustrated in different colours. The two coned-off blocks shown using filled lines are peripheral. These correspond to incident

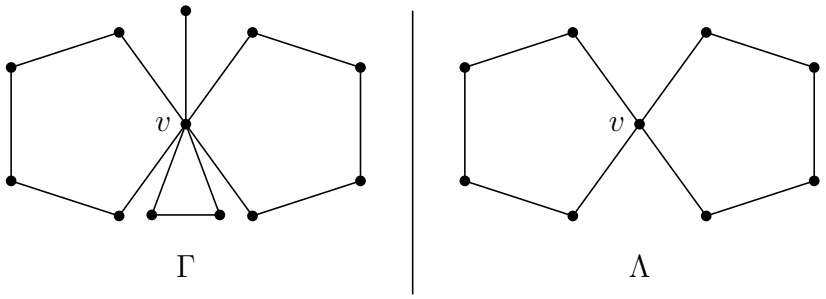


Figure 5.9: The defining graphs of RAAGs that are not quasi-isometric

edge groups in $C_1(\Gamma)$. The two coned-off blocks shown using dashed lines are non-peripheral. These do not correspond to incident edge groups in $C_1(\Gamma)$.

Theorem 4.74 tells us that a quasi-isometry between two RAAGs induces relative quasi-isometries between their cylinder stabilizers. As noted in Definition 5.15, cylinders have a direct product decomposition corresponding to the join decomposition $\{v\} \circ \mathrm{lk}(v)$. Since v is a cut vertex, $\mathrm{lk}(v)$ is disconnected and so doesn't split as a join. Theorem 5.2 thus tells us that a quasi-isometry between RAAGs induces quasi-isometries between the corresponding direct product factors. By applying [PW02], we find obstructions to RAAGs being quasi-isometric.

For example, consider the RAAGs $A(\Gamma)$ and $A(\Lambda)$, where Γ and Λ are as in Figure 5.9. The cylindrical vertex group in $C_1(\Gamma)$ is of the form $\mathbb{Z} \times (F_2 * F_2 * \mathbb{Z} * \mathbb{Z}^2)$, whereas the cylindrical vertex group in $C_1(\Lambda)$ is of the form $\mathbb{Z} \times (F_2 * F_2)$.

Since $F_2 * F_2 * \mathbb{Z} * \mathbb{Z}^2$ and $F_2 * F_2$ don't have the same one-ended factors in their Gruško decomposition, we deduce from Theorem 0.4 of [PW02] that they are not quasi-isometric. Therefore $A(\Gamma)$ and $A(\Lambda)$ do not have quasi-isometric cylinders, so Theorem 4.74 tells us they cannot themselves be quasi-isometric.

We contrast this with Example 3.22 of [Hua17a], which describes two RAAGs that are commensurable, hence quasi-isometric. The cylinders in this example are of the form $\mathbb{Z} \times (F_2 * F_2)$ and $\mathbb{Z} \times (F_3 * F_3 * \mathbb{Z})$, and these are quasi-isometric.

As remarked in the introduction, the RAAGs defined by the graphs shown in Figure 1 of [BHS17] can also be shown not be quasi-isometric by analysis of the quasi-isometry types of cylinders.

We can do much better than this by observing that a quasi-isometry doesn't just preserve the quasi-isometry type of cylinders, but also their relative quasi-isometry type. The following proposition tells us precisely when cylinders are relatively quasi-isometric.

Proposition 5.25. *Let Γ and Λ be connected subgraphs containing cut vertices $v \in V\Gamma$ and $w \in V\Lambda$. The cylindrical vertex groups $A(\mathrm{star}^{\Gamma}(v))$ and $A(\mathrm{star}^{\Lambda}(w))$ are relatively quasi-isometric if and only if they have the same quasi-isometry types of peripheral factors (one-ended or otherwise), and the same quasi-isometry types of one-ended non-peripheral factors.*

Before proving this, we need to modify the tree of spaces slightly so we are able to apply the machinery of Section 5.3. The idea is as follows: suppose that G_1 and G_2 are finitely generated groups with finite generating sets S_1 and S_2. We can *blow-up* the Cayley graph of $G_1 * G_2$ with respect to the generating set $S_1 \sqcup S_2$ by adding edges and isolated vertices, thus obtaining a free product of the Cayley graphs of G_1 and G_2. This is essentially a constructive reformulation of Remark 5.20.

We now observe that a cylindrical vertex space $X(\mathrm{star}^{\Gamma}(v))$ has a cylindrical decomposition $\mathbb{E}^1 \times X(\mathrm{link}^{\Gamma}(v))$. As v is a cut vertex, $A(\mathrm{link}^{\Gamma}(v))$ decomposes non-trivially as a free product $A(R_1) * \cdots * A(R_n) * A(S_1) * \cdots * A(S_m)$ as shown above. We apply the above blowing up procedure so that the cylindrical vertex space is of the form $\mathbb{E}^1 \times \Sigma_v$, where Σ_v is a free product of graphs.

Definition 5.26. Let G be a RAAG with JSJ tree of cylinders T and tree of spaces X. Applying the above blowing-up procedure yields the *blown-up tree of spaces* $\hat{X}$. It is clear that X is equivariantly quasi-isometric to $\hat{X}$.

For clarity, we recall the peripheral structure of the cylindrical vertex space $\mathbb{E}^1 \times \Sigma_v$. Each incident edge e of T is identified with a coset $gA(R_i)$, and hence with a special subgraph A_e in Σ_v. The associated edge space is then $\mathbb{E}^1 \times A_e$. Elements of $\mathcal{S}_e$ are precisely subsets of the form $\mathbb{E}^1 \times \{a\}$, where a is a vertex of the special subgraph A_e. This follows from Remark 5.16. Each special subgraph of Σ_v is either *peripheral* if it is equal to A_e for some $e \in ET$, or *non-peripheral* otherwise.

Proof of Proposition 5.25. Suppose that two cylindrical vertex groups $A(\mathrm{star}^{\Gamma}(v))$ and $A(\mathrm{star}^{\Lambda}(w))$ have the same quasi-isometry types of peripheral factors and the same quasi-isometry types of one-ended non-peripheral factors. We claim that their vertex spaces in the blown-up tree of spaces are relatively quasi-isometric. We do this by constructing a quasi-isometry $f : \Sigma_v \to \Sigma_w$ that bijectively sends peripheral special subgraphs to within uniform finite Hausdorff distance of peripheral special subgraphs. This implies the claim. We proceed as follows.

Let v_0 and w_0 be the depth one isolated vertices of Σ_v and Σ_w respectively. By Remark 5.24, non-peripheral special subgraphs are either lines or are one-ended. If

necessary, we apply a type III move at v_0 so that $\mathrm{lk}^+(v_0)$ and $\mathrm{lk}^+(w_0)$ consist of the same quasi-isometry types of non-peripheral special subgraphs.

We now apply type I moves at v_0 to obtain a bijection $\chi : \mathrm{lk}^+(v_0) \to \mathrm{lk}^+(w_0)$ that preserves quasi-isometry class and such that $\chi(A)$ is peripheral if and only if A is. We apply Lemma 5.21 at v_0, where for each $A \in \mathrm{lk}^+(v_0)$, we choose a quasi-isometry $f_A : A \to \chi(A)$ that preserves the base vertex; this can be done since each special subgraph is homogeneous. The depth one part of the resulting tree of spaces now agrees with Σ_w.

We proceed inductively, first applying the necessary type I and type III moves at each isolated vertex of depth i and then applying Lemma 5.21. The tree of graphs we end up with is Σ_w. As Σ_v and Σ_w each consist of finitely many isometry classes of special subgraphs, the quasi-isometry constants of all the f_A can be assumed to be uniformly bounded. This composition of moves is uniformly locally finite, so by Proposition 5.23 we have the required quasi-isometry $\Sigma_v \to \Sigma_w$.

For the converse, suppose that there is a relative quasi-isometry $\phi : A(\mathrm{star}^\Gamma(v)) \to A(\mathrm{star}^\Lambda(w))$. Theorem 5.2 then tells us that ϕ induces a quasi-isometry $f : \Sigma_v \to \Sigma_w$. As ϕ is a relative quasi-isometry, it must bijectively map peripheral special subgraphs to within uniformly finite Hausdorff distance of peripheral special subgraphs. Lemma 3.2 of [PW02] also says that f preserves one-ended special subgraphs up to uniformly finite Hausdorff distance. It follows that Σ_v and Σ_w must have the same quasi-isometry types of peripheral special subgraphs, and the same quasi-isometry types of one-ended non-peripheral special subgraphs. □

5.5 Constructing the quasi-isometry

The aim of this section is to provide a complete quasi-isometry classification of the following class of RAAGs.

Definition 5.27. We say a subgraph is *biconnected* if it is connected and has no cut vertex. A *block graph* is a finite, connected, simplicial graph such that every biconnected component is a clique. A RAAG whose defining graph is a block graph is called a *block graph RAAG*. Let $\mathcal{C}_{ab}$ denote the class of block graph RAAGs.

Recall every JSJ tree of cylinders can be equipped with the naïve decoration. We say two JSJ trees of cylinders T and T' are *equivalent* if, when T and T' are both equipped with the naïve decoration, there is a decoration-preserving tree isomorphism $\chi : T \to T'$.

Theorem 5.28. *Two RAAGs in $\mathcal{C}_{\mathrm{ab}}$ are quasi-isometric if and only if they have equivalent JSJ trees of cylinders.*

To prove this, we apply Proposition 4.9 to the JSJ tree of cylinders. We fix two RAAGs G and G' in $\mathcal{C}_{\mathrm{ab}}$ such that their JSJ trees of cylinders T and T' are equivalent. By Corollary 4.82, we can choose decorations (T, δ) and (T', δ') such that the conclusions of Corollary 4.82 hold. For the reader's convenience, we state them here:

(A) If $\delta(v) = \delta'(w)$ for vertices $v \in VT$ and $w \in VT'$, then there exists a decoration preserving tree isomorphism $\chi : T \to T'$ such that $\chi(v) = w$. Moreover, there is a relative quasi-isometry $\phi_v : (X_v, \mathcal{P}_v) \to (X'_w, \mathcal{P}_w)$ such that $(\phi_v)_* = \chi|_{lk(v)}$.

(B) Suppose also that $\delta(e) = \delta(f)$ for edges $e \in ET$ and $f \in ET'$, with $\iota e = v$ and $\iota f = w$. Then we may choose χ as above such that $\chi(e) = f$ and $(\phi_v)_*(e) = f$.

To each vertex space X_v, we have an induced peripheral structure $\mathcal{P}_v$ corresponding to incident edges in the JSJ tree of cylinders, and an induced decoration $\delta_v : \mathcal{P}_v \to \mathcal{O}$. Rigid vertex stabilizers are abelian, so are extremely easy to describe and work with. The defining graph is an n-clique Γ_v, and the peripheral structure $\mathcal{P}_v$ is an $A(\Gamma)$-invariant collection of parallelism classes of standard geodesics. Thus $\mathcal{P}_v$ can be described by a subset $B_v \subseteq V\Gamma_v$ as in Example 5.3. The induced decoration $\mathcal{P}_v \to \mathcal{O}$ is an $A(\Gamma)$-invariant map, so descends to a decoration $\hat{\delta}_v : B_v \to \mathcal{O}$ on B_v.

Suppose there are rigid vertices $v \in VT$ and $w \in VT'$ such that $\delta(v) = \delta'(w)$. Both vertex spaces X_v and X_w are cube complexes isometric to $\mathbb{E}^n$ for some n. Thus there is a decoration preserving relative quasi-isometry $\phi_v : (X_v, \mathcal{P}_v) \to (X'_w, \mathcal{P}_w)$. This induces a bijection $B_v \to B_w$, which we extend arbitrarily to a graph isomorphism $\Gamma_v \to \Gamma_w$.

This corresponds to map from a basis of $X_v \cong \mathbb{E}^n$ to a basis $X_w \cong \mathbb{E}^n$, which we use to "straighten" ϕ_v and obtain a decoration preserving relative cubical isometry $\hat{\phi}_v : (X_v, \mathcal{P}_v) \to (X'_w, \mathcal{P}_w)$ that sends standard geodesics to standard geodesics. We call such a map a *rectified relative quasi-isometry*.

Remark 5.29. There is a lot of freedom in how we choose the rectified map $\hat{\phi}_v$: it can be modified by composing with any label preserving cubical isomorphism of $X_v \cong \mathbb{E}^n$.

We now prove the main lemma required in the proof of Theorem 5.28. The key idea is to apply the machinery of Section 5.3 to build quasi-isometries between cylindrical vertex spaces as follows.

Lemma 5.30. *There are constants $K \geq 1$ and $A \geq 0$ such that for any cylinders $v \in VT$ and $w \in VT'$ such that $\delta(v) = \delta'(w)$, the following hold.*

(1) There is a (K, A) relative quasi-isometry $\phi_v : X_v \to X'_w$.

(2) For every edge $e \in ET$ with $\iota e = v$, there is a map $\phi_e : X_e \to X'_{(\phi_v)_*(e)}$ such that ϕ_e and $\phi_v|_{X_e}$ are A-close.

(3) For every edge $e \in ET$ with $\iota e = v$, there is a rectified (K, A)-quasi-isometry

$$\phi_{\tau e} : (X_{\tau e}, \mathcal{P}_{\tau e}, \delta_{\tau e}) \to (X_{\tau((\phi_v)_* e)}, \mathcal{P}_{\tau((\phi_v)_* e)}, \delta_{\tau((\phi_v)_* e)})$$

such that the maps $\alpha_{(\phi_v)_* e} \circ \phi_e$ and $\phi_{\tau e} \circ \alpha_e$ are A-close.

Moreover, suppose there exist edges $e \in ET$ and $f \in ET'$ such that $\iota e = v$, $\iota f = w$ and $\delta(e) = \delta'(f)$. If there is a rectified relative quasi-isometry $\rho : X_{\tau e} \to X'_{\tau f}$ such that $\rho_*(e) = f$, then (1)-(3) hold with the restriction $(\phi_v)_*(e) = f$ and $\phi_{\tau e} = \rho$.

Proof. Every cylindrical vertex space X_v admits a cylindrical decomposition $X_v = \mathbb{E}^1 \times Y_v$ as in Definition 5.15. By the argument used in Section 5.4, we replace the trees of spaces with blown up trees of spaces. We thus assume that each vertex space X_v is of the form $\mathbb{E}^1 \times \Sigma_v$, where Σ_v is a free product of graphs.

If $e \in T$ is an edge with $\iota e = v$, the edge space $X_e \subseteq X_v$ is of the form $\mathbb{E}^1 \times A_e$, where A_e is some peripheral special subgraph of Σ_v. Moreover, since rigid vertices are abelian, the vertex space $X_{\tau e}$ is actually equal to $X_{\bar{e}}$. For simplicity, we omit the edge maps from expressions and simply identify $X_{\tau e}$ with $\mathbb{E}^1 \times A_e$. Let v_0 and w_0 be depth one isolated vertices of Γ_v and Γ_w respectively. We now show the following:

(i) There exists a (K, A)-quasi-isometry $\psi : \Sigma_v \to \Sigma_w$ and a map $\psi_* : \mathrm{lk}(v) \to \mathrm{lk}(w)$ such that $d_{\mathrm{Haus}}(\psi(A_e), A_{\psi_*(e)}) \leq A$ for every edge $e \in ET$ with $\iota e = v$.

(ii) For every edge $e \in ET$ with $\iota e = v$, there is a rectified decoration preserving relative (K, A)-quasi-isometry $\phi_{\tau e} : X_{\tau e} \to X_{\tau((\psi)_* e)}$ that splits as

$$\phi^{(1)}_{\tau e} \times \phi^{(2)}_{\tau e} : \mathbb{E}^1 \times A_e \to \mathbb{E}^1 \times A_{(\psi)_* e}.$$

Moreover, $\psi|_{A_e}$ is A-close to $\phi^{(2)}_{\tau e}$, and $\phi^{(1)}_{\tau e} = \phi^{(1)}_{\tau e'}$ for all $e, e' \in ET$.

Once we have shown this, the map $\phi_v := \phi^{(1)}_{\tau e} \times \psi : \mathbb{E}^1 \times \Sigma_v \to \mathbb{E}^1 \times \Sigma_w$ will be the required map.

The construction of ψ is the most involved part of the proof. It is a modification of the construction of the quasi-isometry between free products of graphs in the proof

of Proposition 5.25. The reader is advised to understand this simpler construction first before proceeding with the current proof.

We decorate special subgraphs of Σ_v and Σ_w with the following data:

- the type of each special subgraph: peripheral or non-peripheral;
- the quasi-isometry class of each special subgraph;
- for each peripheral special subgraph, the ornament assigned to the edge corresponding to that special subgraph.

As $\delta(v) = \delta'(w)$, (A) ensures there is a relative quasi-isometry $\rho_v : (X_v, \mathcal{P}_v) \to (X'_w, \mathcal{P}_w)$ that preserves the decoration on incident edge spaces. We thus deduce that special subgraphs in $\mathrm{lk}^+(v_0)$ and $\mathrm{lk}^+(w_0)$ that are either peripheral or are one-ended and non-peripheral, are decorated with the same sets of ornaments.

As in the proof of Proposition 5.25, we apply type I and III moves at v_0 to ensure that there is a decoration-preserving bijection $\chi : \mathrm{lk}^+(v_0) \to \mathrm{lk}^+(w_0)$. Let $A \in \mathrm{lk}^+(v_0)$ and $\chi(A) \in \mathrm{lk}^+(w_0)$ denote peripheral special subgraphs, and suppose e_A and e'_A denote the corresponding edges in $\mathrm{lk}(v)$ and $\mathrm{lk}(w)$ respectively. We orient these edges so that $\iota e_A = v$ and $\iota e'_A = w$.

Let $a_0 \in A$ and $a'_0 \in \chi(A)$ be the basepoints. As $\delta(\tau e_A) = \delta'(\tau e'_A)$, (B) ensures we may choose a rectified decoration preserving relative (K, A)-quasi-isometry $\phi_{\tau e_A} : X_{\tau e} \to X_{\tau e'_A}$. Note that $\phi_{\tau e_A}$ splits as $\phi^{(1)}_{\tau e_A} \times \phi^{(2)}_{\tau e_A}$. By Remark 5.29, we may assume that $(\phi^{(2)}_{\tau e_A})(a_0) = a'_0$ and that $\phi^{(1)}_{\tau e_A} = \mathrm{id}$. We set $f_A = \phi^{(2)}_{\tau e_A}$. We define $f_{A'}$ similarly for each peripheral special subgraph $A' \in \mathrm{lk}^+(v_0)$. For non-peripheral special subgraphs $A \in \mathrm{lk}^+(v_0)$, we set f_A to be any basepoint preserving isometry $f_A : A \to \chi(A)$.

We now proceed as in Proposition 5.25 and apply Lemma 5.21 with the maps $\{f_A\}_{A \in \mathrm{lk}^+(v_0)}$. Doing this, we obtain a new tree of graphs whose depth 1 part is equal to Σ_w. We proceed by induction on depth of special subgraphs and apply moves of type I, II and III till we obtain Σ_w. This collection of moves is uniformly locally finitely, so we obtain the required quasi-isometry $\psi : \Sigma_v \to \Sigma_w$.

We now prove the final part of the proposition and explain how to modify this construction if one already has e, f and ρ as stated. The rectified map ρ splits as $\rho^{(1)} \times \rho^{(2)} : \mathbb{E}^1 \times A_e \to \mathbb{E}^1 \times A_f$. Suppose a_0 is the basepoint of A_e. We choose isolated vertices $x \in \Sigma_v$ and $y \in \Sigma_w$ so that A_e and A'_f are attached to x and y at a_0 and $\rho^{(2)}(a_0)$ respectively. We "rebase" Σ_v and Σ_w so that x and y are the depth one isolated vertices. We now proceed as above, ensuring $\chi(A) = A'$ and $\phi_{\tau e} = \rho$, and setting $\rho^{(1)} = \phi^{(1)}_{\tau e_A}$ for all special subgraphs A. □

Proof of Theorem 5.28. We fix two RAAGs G and G' in $\mathcal{C}_{\text{ab}}$. Theorem 5.17 ensures that if G and G' are quasi-isometric, then they have equivalent JSJ trees of cylinders when decorated with the naïve decoration. We prove the converse. If T and T' are equivalent, we proceed as above and apply Corollary 4.82. We thus obtain decorations (T, δ) and (T', δ') that satisfy the properties stated in Corollary 4.82. We now construct an isomorphism $\chi : T \to T'$ and maps $\{\phi_v\}$ and $\{\phi_e\}$ that satisfy the hypotheses of Proposition 4.9.

We choose arbitrary cylindrical vertices $v \in VT$ and $w \in VT'$ decorated with the same ornament and set $\chi(v) = w$. As $\delta(v) = \delta'(w)$, there exists a quasi-isometry $\phi_v : X_v \to X'_w$ as in Lemma 5.30. We define $\chi|_{\text{lk}(v)} := (\phi_v)_*$. For each vertex v' adjacent to v, Lemma 5.30 gives a quasi-isometry $\phi_{v'} : X_{v'} \to X_{\chi(v')}$. We extend χ to $\text{lk}(v')$ by $(\phi_{v'})_*$. Lemma 5.30 ensures the the quasi-isometries $\{\phi_e\}$ and $\{\phi_v\}$, where defined thus far, do indeed satisfy the hypotheses of Proposition 4.9.

We work outwards from v, applying Lemma 5.30 at each cylindrical vertex, thus extending the domain of χ and defining new relative quasi-isometries. For instance, suppose v_1 is a rigid vertex adjacent to v and $v_2 \neq v$ is a cylindrical vertex adjacent to v_1. Let $e = (v_2, v_1)$. We use Lemma 5.30 to construct a relative quasi-isometry $\phi_{v_2} : X_{v_2} \to X'_{\chi(v_2)}$. By the final part of Lemma 5.30, ϕ_{v_2} can be chosen so that $(\phi_{v_2})_*(e) = \chi(e)$, and that ϕ_{v_2} and ϕ_{v_1} agree on the edge space X_e. We continue in this way, defining more maps ϕ_v and ϕ_e and extending the domain of χ.

We conclude with a uniformity argument: as there are only finitely many isometry classes of vertex and edge spaces in T and T', all maps ϕ_v and ϕ_e can be chosen from a finite set of model maps. We thus have uniform quasi-isometry constants for all maps $\{\phi_v\}$ and $\{\phi_e\}$. Thus Proposition 4.9 can be applied, showing that G and G' are indeed quasi-isometric. □

5.6 The quasi-isometry problem

We now wish to investigate whether one can algorithmically decide if two RAAGs in $\mathcal{C}_{\text{ab}}$ are quasi-isometric. Let $\mathcal{D}$ be an arbitrary class of RAAGs and $\mathcal{O}$ be an arbitrary finite set.

Suppose we are given the tuple $(\Gamma, B_\Gamma, \hat{\delta}_\Gamma, v)$, where Γ is a finite simplicial graph, $B_\Gamma \subseteq V\Gamma$, $\hat{\delta}$ is a function $B_\Gamma \to \mathcal{O}$ and $v \in B_\Gamma$. As in Example 5.3, the vertex set B_Γ induces a peripheral structure $\mathcal{P}_{B_\Gamma}$ on $X(\Gamma)$. The map $\hat{\delta}$ induces a function $\delta_\Gamma : \mathcal{P}_{B_\Gamma} \to \mathcal{O}$. We call elements of $\mathcal{O}$ *ornaments* and call the functions $\hat{\delta}$ and δ_Γ *decorations*.

Definition 5.31. We say that $\mathcal{D}$ has *solvable quasi-isometry problem* if there is an algorithm that takes as input finite simplicial graphs Γ and Λ and outputs whether there is a quasi-isometry $f : X(\Gamma) \to X(\Lambda)$.

We say that $\mathcal{D}$ has *solvable strong quasi-isometry problem* if there is an algorithm that takes as input the tuples $(\Gamma, B_\Gamma, \hat{\delta}_\Gamma, v)$ and $(\Lambda, B_\Lambda, \hat{\delta}_\Lambda, w)$ as above, such that $A(\Gamma), A(\Lambda) \in \mathcal{D}$, and outputs whether there is a decoration preserving relative quasi-isometry $f : (X(\Gamma), \mathcal{P}_{B_\Gamma}) \to (X(\Lambda), \mathcal{P}_{B_\Lambda})$ with $[f(l)] = [l']$, where l and l' are standard geodesics labelled by v and w respectively.

Proposition 5.32. *Abelian RAAGs have solvable strong relative quasi-isometry problem.*

Proof. Suppose we are given the data $(\Gamma, B_\Gamma, \hat{\delta}_\Gamma, v)$ and $(\Lambda, B_\Lambda, \hat{\delta}_\Lambda, w)$, where Γ and Λ are complete graphs on n and m vertices respectively. Suppose also that there exists a decoration preserving relative quasi-isometry $f : (X(\Gamma), \mathcal{P}_{B_\Gamma}) \to (X(\Lambda), \mathcal{P}_{B_\Lambda})$ such that $[f(l)] = [l']$, where l and l' are standard geodesics labelled by v and w respectively. Then $n = m$, and there exists a decoration-preserving bijection $r_* : B_\Gamma \to B_\Lambda$ such that $r_*(v) = w$.

Conversely, if such an r_* exists, it can be arbitrarily extended to a decoration preserving graph isomorphism $\hat{r}_* : \Gamma \to \Lambda$, which then induces the required decoration-preserving based relative quasi-isometry. Thus there exists a suitable relative quasi-isometry if and only if there exists a bijection $r_* : B_\Gamma \to B_\Lambda$ as above. Since B_Γ and B_Λ are finite, it is easy to verify whether or not such an r_* exists. □

Proposition 5.33. *We fix classes of RAAGs $\mathcal{C}_1$ and $\mathcal{C}_2$ that have solvable strong quasi-isometry problem and solvable quasi-isometry problem respectively. We consider the class $\mathcal{C}_0$ of RAAGs that satisfy the following.*

(1) For each $A(\Gamma) \in \mathcal{C}_0$, all rigid vertex stabilizers of the JSJ tree of cylinders decomposition $C_1(\Gamma)$ lie in $\mathcal{C}_1$.

(2) For each $A(\Gamma) \in \mathcal{C}_0$, all non-peripheral factors of cylinder stabilizers of $C_1(\Gamma)$ lie in $\mathcal{C}_2$.

Then there exists an algorithm that takes as input $A(\Gamma), A(\Lambda) \in \mathcal{C}_0$ and outputs whether the trees of cylinders $(T(\Gamma), \delta)$ and $(T(\Lambda), \delta')$, equipped with the naïve decoration, are equivalent.

Proof. We first show that cylinder stabilizers of RAAGs in $\mathcal{C}_0$ have solvable quasi-isometry problem. Indeed, by applying the strong solvable quasi-isometry problem to adjacent rigid vertex stabilizers, we are able to solve the quasi-isometry problem for peripheral factors of cylinder stabilizers of $C_1(\Gamma)$. By assumption, we are also able to solve the quasi-isometry problem for non-peripheral factors of cylinder stabilizer. By Proposition 5.25, we thus have an algorithm to decide whether cylinder stabilizers are relatively quasi-isometric. In fact, a slight modification of the proof of Proposition 5.25 allows us to solve the strong quasi-isometry problem for cylinder stabilizers of RAAGs.

We now decorate vertex and edges in $C_1(\Gamma)$ with the naïve decoration. More precisely, we decorate each vertex and edge with a symbol from the set $\{o_1, o_2, \dots\}$ such that two vertices share the symbol if and only if they have the same ornament when assigned the naïve decoration. This can be done using the fact that all cylindrical and rigid vertex stabilizers have solvable relative quasi-isometry problem.

We now perform neighbour refinement, which can be done by simply counting adjacency data of ornaments, when lifted to the Bass-Serre tree. We pick unused letters from the set $\{o_1, o_2, \dots\}$ when refinement occurs. We may also perform vertex refinement, using the fact that the class of all cylindrical and rigid vertex stabilizers have solvable strong relative quasi-isometry problem. We apply this procedure till it stabilizes, giving a decoration δ.

We apply the same procedure to $C_1(\Lambda)$, using a different set of ornaments $\{o'_1, \dots\}$, and obtaining a decoration δ'. By examining all possible bijections $\mathrm{im}(\delta) \to \mathrm{im}(\delta')$ and seeing if the conditions of Proposition 4.81 hold, we are able to determine whether the decorated trees of cylinders are equivalent. □

Combining this with Theorem 5.28, we deduce the following:

Theorem 5.34. *There is an algorithm that, when given as input finite simplicial graphs Γ and Λ such that $A(\Gamma) \in \mathcal{C}_{\mathrm{ab}}$, determines whether $A(\Gamma)$ is quasi-isometric to $A(\Lambda)$.*

5.7 n-clique tree-graded RAAGs

We give a sample application of Theorem 5.28. We recall the definition of n-clique tree-graded RAAGs:

Definition 5.35. Let Γ be a finite connected simplicial graph, and let C_0 denote the set of cut vertices of Γ. For $n \geq 2$, we say that Γ is an *n-clique tree-graded graph* if:

- every vertex of Γ is either contained in C_0 or has valence one;
- Γ has diameter at least three;
- each maximal biconnected subgraph is either an n-clique, all of whose vertices are in C_0, or a 2-clique containing exactly one vertex in C_0.

We say that $A(\Gamma)$ is an *n-clique tree-graded RAAG* whenever Γ is an n-clique tree-graded graph.

Theorem 5.36. *Fix some $n \geq 2$. All n-clique tree-graded RAAGs are quasi-isometric to one another, and every RAAG quasi-isometric to an n-clique tree-graded RAAG is itself an n-clique tree-graded RAAG.*

Proof. Let Γ be an n-clique tree-graded graph. We analyse cylindrical and rigid vertex groups in $C(\Gamma)$. Cylindrical vertex groups are of the form

$$\mathbb{Z} \times (\mathbb{Z}^{n-1} * \cdots * \mathbb{Z}^{n-1} * \mathbb{Z} * \cdots * \mathbb{Z})$$

where the $\mathbb{Z}^{n-1}$ are peripheral factors and the $\mathbb{Z}$ are non-peripheral factors. In particular, there are no one-ended non-peripheral factors. Using Proposition 5.25, we deduce that $C(\Gamma)$ has a single relative quasi-isometry class of cylindrical vertices.

The rigid vertex groups are all of the form $A(\Gamma') \cong \mathbb{Z}^n$. The peripheral structure consists of n elements, each of which is a parallelism class of standard geodesics in $X(\Gamma')$. Thus there is a single relative quasi-isometry class of rigid vertices.

The naïve decoration consists of two ornaments, one given to all rigid vertices and one given to all cylinder vertices. Each cylinder (resp. rigid) vertex is adjacent to countably infinitely many rigid (resp. cylinder) vertices, so the naïve decoration is stable under neighbour refinement. All vertices with the same ornament are relatively quasi-isometric via a relative quasi-isometry that preserves the decoration on incident edges. Thus the naïve decoration is also stable under vertex refinement. Proposition 4.81 ensures that for some fixed $n \geq 2$, any two n-clique tree-graded RAAGs have equivalent JSJ trees of cylinders when decorated with the initial decoration. Thus by Theorem 5.28, any two n-clique tree-graded RAAGs are quasi-isometric.

Now suppose $A(\Gamma)$ is a RAAG that is a quasi-isometric to an n-clique tree-graded RAAG. Thus by Theorem 5.17 and Proposition 5.25, cylindrical and rigid vertex groups of $C_1(\Gamma)$ are of the above form described above. In particular, the subgraph associated to each rigid vertex is an n-clique, all of whose vertices are cut vertices of Γ. Every cylindrical vertex has peripheral $\mathbb{Z}^{n-1}$ factors corresponding to adjacent

rigid vertices, and non-peripheral $\mathbb{Z}$ factors corresponding to adjacent leaves. Thus Γ must be an n-clique tree-graded graph. $\square$

Bibliography

[ABD+13] Aaron Abrams, Noel Brady, Pallavi Dani, Moon Duchin, and Robert Young. Pushing fillings in right-angled Artin groups. *J. Lond. Math. Soc. (2)*, 87(3):663–688, 2013.

[Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.

[AN98] S. A. Adeleke and Peter M. Neumann. Relations related to betweenness: their structure and automorphisms. *Mem. Amer. Math. Soc.*, 131(623):viii+125, 1998.

[Bar18] Benjamin Barrett. Computing JSJ decompositions of hyperbolic groups. *Journal of Topology*, 11(2):527–558, 2018.

[BB97] Mladen Bestvina and Noel Brady. Morse theory and finiteness properties of groups. *Invent. Math.*, 129(3):445–470, 1997.

[BC12] Jason Behrstock and Ruth Charney. Divergence and quasimorphisms of right-angled Artin groups. *Math. Ann.*, 352(2):339–356, 2012.

[BE78] Robert Bieri and Beno Eckmann. Relative homology and Poincaré duality for group pairs. *J. Pure Appl. Algebra*, 13(3):277–319, 1978.

[Bes96] Mladen Bestvina. Local homology properties of boundaries of groups. *Michigan Math. J.*, 43(1):123–139, 1996.

[BF91] Mladen Bestvina and Mark Feighn. Bounding the complexity of simplicial group actions on trees. *Invent. Math.*, 103(3):449–469, 1991.

[BHS17] Jason Behrstock, Mark F Hagen, and Alessandro Sisto. Quasiflats in hierarchically hyperbolic spaces. *arXiv preprint arXiv:1704.04271*, 2017.

[Bie81] Robert Bieri. *Homological dimension of discrete groups.* Queen Mary College Mathematical Notes. Queen Mary College, Department of Pure Mathematics, London, second edition, 1981.

[BJN10] Jason A. Behrstock, Tadeusz Januszkiewicz, and Walter D. Neumann. Quasi-isometric classification of some high dimensional right-angled Artin groups. *Groups Geom. Dyn.*, 4(4):681–692, 2010.

[BKS08] Mladen Bestvina, Bruce Kleiner, and Michah Sageev. The asymptotic geometry of right-angled Artin groups. I. *Geom. Topol.*, 12(3):1653–1699, 2008.

[BL18] Benjamin Beeker and Nir Lazarovich. Detecting sphere boundaries of hyperbolic groups. *Geom. Topol.*, 22(1):439–470, 2018.

[BM91] Mladen Bestvina and Geoffrey Mess. The boundary of negatively curved groups. *J. Amer. Math. Soc.*, 4(3):469–481, 1991.

[BM00] Kenneth S. Brown and John Meier. Improper actions and higher connectivity at infinity. *Comment. Math. Helv.*, 75(1):171–188, 2000.

[BM01] Noel Brady and John Meier. Connectivity at infinity for right angled Artin groups. *Trans. Amer. Math. Soc.*, 353(1):117–132, 2001.

[BN08] Jason A. Behrstock and Walter D. Neumann. Quasi-isometric classification of graph manifold groups. *Duke Math. J.*, 141(2):217–240, 2008.

[Bow98] Brian H. Bowditch. Cut points and canonical splittings of hyperbolic groups. *Acta Math.*, 180(2):145–186, 1998.

[Bow99] B. H. Bowditch. Treelike structures arising from continua and convergence groups. *Mem. Amer. Math. Soc.*, 139(662):viii+86, 1999.

[Bow02] Brian H. Bowditch. Splittings of finitely generated groups over two-ended subgroups. *Trans. Amer. Math. Soc.*, 354(3):1049–1078, 2002.

[Bow04] Brian H. Bowditch. Planar groups and the Seifert conjecture. *J. Reine Angew. Math.*, 576:11–62, 2004.

[Bro82] Kenneth S. Brown. *Cohomology of groups*, volume 87 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1982.

[Bro87] Kenneth S. Brown. Finiteness properties of groups. *J. Pure Appl. Algebra*, 44(1-3):45–75, 1987.

[Cha07] Ruth Charney. An introduction to right-angled Artin groups. *Geom. Dedicata*, 125:141–158, 2007.

[CJ94] Andrew Casson and Douglas Jungreis. Convergence groups and Seifert fibered 3-manifolds. *Invent. Math.*, 118(3):441–456, 1994.

[Cla14] Matt Clay. When does a right-angled Artin group split over $\mathbb{Z}$? *Internat. J. Algebra Comput.*, 24(6):815–825, 2014.

[CM17] Christopher H. Cashen and Alexandre Martin. Quasi-isometries between groups with two-ended splittings. *Math. Proc. Cambridge Philos. Soc.*, 162(2):249–291, 2017.

[CRKZ16] Montserrat Casals-Ruiz, Ilya Kazachkov, and Alexander Zakharov. On commensurability of some right-angled artin groups. *arXiv preprint arXiv:1611.01741*, 2016.

[DD89] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs*, volume 17 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1989.

[DJ00] Michael W. Davis and Tadeusz Januszkiewicz. Right-angled Artin groups are commensurable with right-angled Coxeter groups. *J. Pure Appl. Algebra*, 153(3):229–235, 2000.

[DK18] Cornelia Druţu and Michael Kapovich. *Geometric group theory*, volume 63 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2018. With an appendix by Bogdan Nica.

[DM02] Michael W. Davis and John Meier. The topology at infinity of Coxeter groups and buildings. *Comment. Math. Helv.*, 77(4):746–766, 2002.

[DR93] M. J. Dunwoody and M. A. Roller. Splitting groups over polycyclic-by-finite subgroups. *Bull. London Math. Soc.*, 25(1):29–36, 1993.

[DS99] M. J. Dunwoody and M. E. Sageev. JSJ-splittings for finitely presented groups over slender groups. *Invent. Math.*, 135(1):25–44, 1999.

[DS00] M. J. Dunwoody and E. L. Swenson. The algebraic torus theorem. *Invent. Math.*, 140(3):605–637, 2000.

[DT17] Pallavi Dani and Anne Thomas. Bowditch's JSJ tree and the quasi-isometry classification of certain Coxeter groups. *J. Topol.*, 10(4):1066–1106, 2017.

[Dun79] M. J. Dunwoody. Accessibility and groups of cohomological dimension one. *Proc. London Math. Soc. (3)*, 38(2):193–215, 1979.

[Dun85] M. J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.

[Dun93] Martin J. Dunwoody. An inaccessible group. In *Geometric group theory, Vol. 1 (Sussex, 1991)*, volume 181 of *London Math. Soc. Lecture Note Ser.*, pages 75–78. Cambridge Univ. Press, Cambridge, 1993.

[Dun17] Martin J. Dunwoody. Structure trees, networks and almost invariant sets. In *Groups, graphs and random walks*, volume 436 of *London Math. Soc. Lecture Note Ser.*, pages 137–175. Cambridge Univ. Press, Cambridge, 2017.

[EP14] Anastasia Evangelidou and Panos Papasoglu. A cactus theorem for end cuts. *Internat. J. Algebra Comput.*, 24(1):95–112, 2014.

[Far75] F. Thomas Farrell. Poincaré duality and groups of type (FP). *Comment. Math. Helv.*, 50:187–195, 1975.

[FP06] K. Fujiwara and P. Papasoglu. JSJ-decompositions of finitely presented groups and complexes of groups. *Geom. Funct. Anal.*, 16(1):70–125, 2006.

[FS96] Benson Farb and Richard Schwartz. The large-scale geometry of Hilbert modular groups. *J. Differential Geom.*, 44(3):435–478, 1996.

[Gab92] David Gabai. Convergence groups are Fuchsian groups. *Ann. of Math. (2)*, 136(3):447–510, 1992.

[Geo08] Ross Geoghegan. *Topological methods in group theory*, volume 243 of *Graduate Texts in Mathematics*. Springer, New York, 2008.

[GH17] Daniel Groves and Michael Hull. Abelian splittings of right-angled Artin groups. In *Hyperbolic geometry and geometric group theory*, volume 73 of *Adv. Stud. Pure Math.*, pages 159–165. Math. Soc. Japan, Tokyo, 2017.

[GL10] Vincent Guirardel and Gilbert Levitt. Scott and Swarup's regular neighborhood as a tree of cylinders. *Pacific J. Math.*, 245(1):79–98, 2010.

[GL11] Vincent Guirardel and Gilbert Levitt. Trees of cylinders and canonical splittings. *Geom. Topol.*, 15(2):977–1012, 2011.

[GL17] Vincent Guirardel and Gilbert Levitt. JSJ decompositions of groups. *Astérisque*, (395):viii + 165, 2017.

[GM85] Ross Geoghegan and Michael L. Mihalik. Free abelian cohomology of groups and ends of universal covers. *J. Pure Appl. Algebra*, 36(2):123–137, 1985.

[GM86] Ross Geoghegan and Michael L. Mihalik. A note on the vanishing of $H^n(G, \mathbf{Z}G)$. *J. Pure Appl. Algebra*, 39(3):301–304, 1986.

[Gro81] Mikhael Gromov. Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math.*, (53):53–73, 1981.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[Gro93] M. Gromov. Asymptotic invariants of infinite groups. In *Geometric group theory, Vol. 2 (Sussex, 1991)*, volume 182 of *London Math. Soc. Lecture Note Ser.*, pages 1–295. Cambridge Univ. Press, Cambridge, 1993.

[GW16] Giovanni Gandini and Nathalie Wahl. Homological stability for automorphism groups of RAAGs. *Algebr. Geom. Topol.*, 16(4):2421–2441, 2016.

[HK18] Jingyin Huang and Bruce Kleiner. Groups quasi-isometric to right-angled Artin groups. *Duke Math. J.*, 167(3):537–602, 2018.

[HO17] Jingyin Huang and Damian Osajda. Quasi-Euclidean tilings over 2–dimensional Artin groups and their applications. *arXiv:1711.00122*, 2017.

[Hop44] Heinz Hopf. Enden offener Räume und unendliche diskontinuierliche Gruppen. *Comment. Math. Helv.*, 16:81–100, 1944.

[Hou74] 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.

[HRY93] Nigel Higson, John Roe, and Guoliang Yu. A coarse Mayer-Vietoris principle. *Math. Proc. Cambridge Philos. Soc.*, 114(1):85–97, 1993.

[Hua15] Jingyin Huang. Quasi-isometry of right-angled Artin groups, II: Several infinite out case. *preprint*, 2015.

[Hua17a] Jingyin Huang. Quasi-isometric classification of right-angled Artin groups, I: The finite out case. *Geom. Topol.*, 21(6):3467–3537, 2017.

[Hua17b] Jingyin Huang. Top-dimensional quasiflats in CAT(0) cube complexes. *Geom. Topol.*, 21(4):2281–2352, 2017.

[Hua18] Jingyin Huang. Commensurability of groups quasi-isometric to RAAGs. *Inventiones mathematicae*, May 2018.

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[KK05] Michael Kapovich and Bruce Kleiner. Coarse Alexander duality and duality groups. *J. Differential Geom.*, 69(2):279–352, 2005.

[KKL98] Michael Kapovich, Bruce Kleiner, and Bernhard Leeb. Quasi-isometries and the de Rham decomposition. *Topology*, 37(6):1193–1211, 1998.

[Kle02] Bruce Kleiner. The cohomology group of G with RG-coefficients, 2002. Unpublished.

[KM12] Jeremy Kahn and Vladimir Markovic. Immersing almost geodesic surfaces in a closed hyperbolic three manifold. *Ann. of Math. (2)*, 175(3):1127–1190, 2012.

[KR89] P. H. Kropholler and M. A. Roller. Relative ends and duality groups. *J. Pure Appl. Algebra*, 61(2):197–210, 1989.

[Mar13] Vladimir Markovic. Criterion for Cannon's conjecture. *Geom. Funct. Anal.*, 23(3):1035–1061, 2013.

[Mar18a] Alexander Margolis. Quasi-isometry classification of RAAGs that split over cyclic subgroups. *arXiv preprint arXiv:1803.05493*, 2018.

[Mar18b] Alexander J Margolis. Quasi-isometry invariance of group splittings over coarse Poincaré duality groups. *Proceedings of the London Mathematical Society*, 116(6):1406–1456, 2018.

[MSW03] Lee Mosher, Michah Sageev, and Kevin Whyte. Quasi-actions on trees. I. Bounded valence. *Ann. of Math. (2)*, 158(1):115–164, 2003.

[MSW11] Lee Mosher, Michah Sageev, and Kevin Whyte. Quasi-actions on trees II: Finite depth Bass-Serre trees. *Mem. Amer. Math. Soc.*, 214(1008):vi+105, 2011.

[Nib02] G. A. Niblo. The singularity obstruction for group splittings. *Topology Appl.*, 119(1):17–31, 2002.

[Pap05] Panos Papasoglu. Quasi-isometry invariance of group splittings. *Ann. of Math. (2)*, 161(2):759–830, 2005.

[Pap07] Panos Papasoglu. Group splittings and asymptotic topology. *J. Reine Angew. Math.*, 602:1–16, 2007.

[Pap12] Panos Papasoglu. Splittings and the asymptotic topology of the lamplighter group. *Trans. Amer. Math. Soc.*, 364(7):3861–3873, 2012.

[PS06] Panos Papasoglu and Eric Swenson. From continua to $\mathbb{R}$-trees. *Algebr. Geom. Topol.*, 6:1759–1784, 2006.

[PS09] Panos Papasoglu and Eric Swenson. Boundaries and JSJ decompositions of CAT(0)-groups. *Geom. Funct. Anal.*, 19(2):559–590, 2009.

[PW02] Panos Papasoglu and Kevin Whyte. Quasi-isometries between groups with infinitely many ends. *Comment. Math. Helv.*, 77(1):133–144, 2002.

[Roe93] John Roe. Coarse cohomology and index theory on complete Riemannian manifolds. *Mem. Amer. Math. Soc.*, 104(497):x+90, 1993.

[Roe03] John Roe. *Lectures on coarse geometry*, volume 31 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2003.

[RS97] 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.

[Sco98] Peter Scott. The symmetry of intersection numbers in group theory. *Geom. Topol.*, 2:11–29, 1998.

[Sco78] Peter Scott. Ends of pairs of groups. *J. Pure Appl. Algebra*, 11(1-3):179–198, 1977/78.

[Ser77] Jean-Pierre 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.

[SS00] Peter Scott and Gadde A. Swarup. Splittings of groups and intersection numbers. *Geom. Topol.*, 4:179–218 (electronic), 2000.

[SS03] Peter Scott and Gadde A. Swarup. Regular neighbourhoods and canonical decompositions for groups. *Astérisque*, (289):vi+233, 2003. Corrections available at `http://www.math.lsa.umich.edu/~pscott/`.

[Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Ann. of Math. (2)*, 88:312–334, 1968.

[Sta71] John Stallings. *Group theory and three-dimensional manifolds.* Yale University Press, New Haven, Conn.-London, 1971. A James K. Whittemore Lecture in Mathematics given at Yale University, 1969, Yale Mathematical Monographs, 4.

[SW79] Peter Scott and Terry 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-New York, 1979.

[Swa96] G. A. Swarup. On the cut point conjecture. *Electron. Res. Announc. Amer. Math. Soc.*, 2(2):98–100, 1996.

[Tuk88] Pekka Tukia. Homeomorphic conjugates of Fuchsian groups. *J. Reine Angew. Math.*, 391:1–54, 1988.

[TW93] Carsten Thomassen and Wolfgang Woess. Vertex-transitive graphs and accessibility. *J. Combin. Theory Ser. B*, 58(2):248–268, 1993.

[Vav12] Diane M. Vavrichek. On the coarse-geometric detection of subgroups. *Geom. Dedicata*, 156:93–126, 2012.

[Vav13] Diane M. Vavrichek. The quasi-isometry invariance of commensurizer subgroups. *Groups Geom. Dyn.*, 7(1):205–261, 2013.

[War81] L. E. Ward, Jr. Axioms for cutpoints. In *General topology and modern analysis (Proc. Conf., Univ. California, Riverside, Calif., 1980)*, pages 327–336. Academic Press, New York-London, 1981.

[Wei94] Charles A. Weibel. *An introduction to homological algebra*, volume 38 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1994.

[Wis09] Daniel T. Wise. Research announcement: the structure of groups with a quasiconvex hierarchy. *Electron. Res. Announc. Math. Sci.*, 16:44–55, 2009.

Index

(ϵ, δ)-net, 15
1-geodesic metric space, 18
B-dense, 16
PD_n^R group, 42
    orientable, 42
∂_r, 49
$\tilde{e}(G, H)$, 55
n-thick, 130
r-close, 16
t-chain, 17

acyclic at infinity, 39
almost invariant set, 55
anti-Čech approximation, 36

bi-Lipschitz equivalence, 17
biconnected, 148
bounded geometry, 21

chain complex over X, 40
clique, 127
coarse n-manifold group, 79
coarse PD_n^R complex, 41
    group action, 41
coarse PD_n^R space, 41
    half-space, 65
coarse PD_n^R group, 42
coarse $PD_{n+1}^{\mathbb{Z}_2}$ pair, 111
coarse cohomology, 35
coarse complementary component, 49
    almost-essential, 67
    associated to a given splitting, 90
    bisect, 100
    coarse disjoint, 53
    crossing of, 88
    deep, 53
    essential, 64
    irreducible, 51
    non-trivial, 53
    shallow, 53
    strong crossing of, 89
coarse embedding, 16
coarse equivalence, 15
coarse fibration, 75
coarse geodesic metric space, 17
coarse inverse, 16
coarse isometry, 16
coarse Mayer-Vietoris sequence, 58
coarse separation, 53
coarsely uniformly acyclic metric space, 20
codimension one subgroup, 56
control space, 24
cross-connected component, 92
    isolated, 92
    strong, 92
    weak, 92
crossing criterion, 98
cut, 94

decorated tree, 121
    equivalence, 121
decoration, 121

geometric, 121
naïve, 121
refinement, 121
deformation space, 84
domination, 84

edge space, 81
elliptic, 85
universally, 85
enclosing, 97
essentially embedded, 67
almost, 67

flexible vertex, 85
free R-module over X, 24
displacement of a map, 25
finite type, 24
support, 24
free product of graphs, 135

group cohomology, 38

Hausdorff distance, 15

JSJ
deformation space, 85
tree, 85
tree of cylinders, 86

metric complex, 24
n-skeleton, 26
group action, 29
uniformly acyclic, 29
mobility set, 69
modified cohomology, 27
move on a tree of graphs
sliding, 138
type I, 138
type II, 139
type III, 142
uniformly locally finite sequence, 143

ornaments, 121

peripheral structure, 117
pretree, 93
discrete, 94
star, 93
proper chain homotopy, 28
proper chain map, 28

quadratically hanging, 7
quasi-geodesic metric space, 17
quasi-isometric embedding, 16
quasi-isometry, 17
rectified relative, 149
relative, 117
quasi-isometry problem, 153
strong, 153

refinement, 121
neighbour, 121
vertex, 122
relative ends of groups, 54
Houghton–Scott, 56
Kropholler–Roller, 55
right-angled Artin group (RAAG), 9
n-clique tree-graded, 155
block graph, 148
rigid vertex, 85
Rips complex, 17

small, 88
standard k-flat, 127
standard geodesic, 127
standard subcomplex, 127
structure invariant, 124

support of a cocycle, 27
support of a standard subcomplex, 127

topologically homogeneous subtree, 112
tree of cylinders, 85
tree of graphs, 137
tree of spaces, 82
type F_n, 23
type FP^R, 42
type FP^R_n, 23

uniform $PD^{\mathbb{Z}_2}_n$-set, 99

vertex space, 81