Acylindrical and strong accessibility

Michael Edward Hill

Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Selwyn College

June 2021

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Michael Edward Hill
June 2021

Acylindrical and strong accessibility

Michael Edward Hill

Abstract

Weidmann has produced a bound on the number of edges of a graph of groups splitting for when a finitely generated group acts on a tree (k,C)-acylindrically [25]. In the same paper Weidmann conjectures a common generalisation between their result and a theorem of Bestvina and Feighn [2]; which provides a similar bound for finitely generated groups acting on a tree with small edge stabilisers. We will produce an example which shows this conjecture is false. We then extend Weidmann's result to actions which are k-acylindrical except on some set of subgroups with finite height. We then apply this result to a couple of specific cases. The first gives us a bound for actions of hyperbolic groups which are k-acylindrical on non virtually-cyclic subgroups. The second give a bound for a RAAG acting k-acylindrically on non-abelian subgroups. We also provide a sharp bound for finitely generated groups acting k-acylindrically.

We also touch on the subject of strong accessibility. In particular we give an account of a theorem by Louder and Touikan [19] which shows that many hierarchies consisting of slender JSJ-decompositions are finite; in particular JSJ-hierarchies of 2-torsion-free hyperbolic groups are always finite.

Acknowledgements

Firstly I would like thank my supervisor Henry Wilton; whose insight, guidance and encouragement have been vital throughout the entire process. This thesis wouldn't have possible without their help.

I owe my parents the world for their continued love and belief in me, even when I didn't believe in myself. I also wish to give extra thanks to my Dad for being an extra typo checker, despite barely understanding a word!

The company of other mathematicians has also been a crucial stabilizing factor for me. Most notably Michal Buran, Matthew Conder and Calum Ashcroft; but also the other geometric group theorists at Cambridge as well. Additionally the GGSE conferences have been a continual source of both mathematical interest and human connection.

I also wish to give thanks to my various non-academic friends. Particularly those from the group known as the "Selwyn Ninja Turtles" and from the local Super Smash Brothers community. Additionally I wish to thank my mentor Fiona Whelan for all their non-academic help, as well as chocolate cake!

Finally I wish to thank Andrew, who is the best brother ever simply by virtue of being himself.

Table of contents

Chapter 1

Introduction

Remark 1.0.1. *Throughout mathematical literature there is a split on the exact meaning of the symbol '$<$' when applied to groups. Throughout this thesis we use the notation $H < G$ to indicate that H is a* proper *subgroup of G. If we wish to include the possibility that $H = G$ we will instead use the notation $H \leqslant G$.*

Suppose we are given a pair of groups and wish to combine them to create a new larger group. A natural way to do this is to take a subgroup which is common to both of them and "glue" the two groups together along this subgroup to create an *amalgamated free product*. Similarly we can extend a group by taking two isomorphic subgroups and "attaching a cylinder along these subgroups" to create a *HNN-extension*. We can combine combinations of these operations and represent them by a graph, called a *graph of groups*.

Graphs of groups have a natural correspondence with the actions of groups on trees. More explicitly the quotient of a group action on a tree (which doesn't invert any edges) can naturally be turned into a graph of groups. Conversely for any graph of groups there is a corresponding tree, called the *Bass-Serre tree*, together with a group action of the *fundamental group* of the graph of groups. This correspondence forms the basis of *Bass-Serre theory*. Basic facts about Bass-Serre theory will henceforth be assumed; which can be found in [21]. Also unless otherwise specified all group actions will be on trees.

Given that we now have this useful way of combining groups it's natural to ask to what extent we can do the converse. In other words we would like to know how a given group G can be decomposed as a graph of groups. For example we can ask if there is a limit on how complex these decompositions of our group can become; typically with some extra restriction such as the permissible edge groups. It's these sorts of *accessibility* questions that we will be focused on here.

The earliest example of an accessibility result is Grushko's theorem [13]. This implies that a free decomposition of a finitely generated group G has at most $\mathrm{rank}(G)$ vertices with non-trivial label. (The *rank* of a finitely generated group is the size of its smallest generating set.) Further examples include a result by Dunwoody [11] which gives a bound for the number of edges of a (reduced) splitting over finite edge groups for a given finitely presented group. Later Bestvina and Feighn [2] would extend this to any decomposition with *small* edge groups. (A group is *small* if it doesn't act hyperbolically on any tree. An action on a tree is *hyperbolic* if it doesn't fix a point of the tree, a point on the boundary of the tree or a line of the tree.) These results do not extend to finitely generated groups; for instance Dunwoody [10] gives an example of a finitely generated group which has splittings over finite edge groups with an arbitrary number of edges.

It's also possible to obtain bounds by imposing restrictions other than restricting the class of edge groups. We call an action *k–acylindrical* if the fixed point set of any non-trivial element of the acting group has diameter at most k. Sela [20] showed that there is a bound for the size of a (minimal) splitting where the action on the corresponding Bass-Serre tree is k–acylindrical, assuming that G is freely indecomposable and finitely generated. Weidmann [26] later reproved this and gave a nice bound of at most $2k(\mathrm{rank}\,G-1)$ edges for the splitting. Afterwards Delzant [8] showed that a bound for (k,C)–acylindrical actions exists provided that the acting group is finitely presented. (An action of a group G is *(k,C)–acylindrical* if the fixed point set of any subgroup of G with cardinality greater than C has diameter at most k.) Weidmann [25] then extended this to finitely generated groups and suggested that there should be a common generalisation between this and Bestvina and Feighn's aforementioned result on actions with small edge stabilisers. We formalise this question into the following conjecture.

Conjecture 2.4.4. *[25, pg.213] Given a finitely presented group G and* $k>0$ *there some* $C(G,k)$ *such that any reduced action of G which is k–acylindrical on large subgroups has at most* $C(G,k)$ *orbits of edges. (A group is* large *if it's not small. Note that this definition is different to the standard one of mapping onto a non-abelian free group.)*

In Chapter 2 we will give an overview of the proofs of the above results of Dunwoody, Bestvina—Feighn and Weidmann with emphasis on ideas which will be useful moving forward. We will also construct original examples which give a negative answer to Conjecture 2.4.4. These examples also give significant restrictions on what partial versions of Conjecture 2.4.4 could hold.

Theorem 2.4.5. *There is a finitely presented group G which for any $N > 0$ acts on a reduced tree which is 1–acylindrical on infinite subgroups and has N orbits of edges.*

Theorem 2.4.6. *For any $N > 0$ there is an action of F_2 on a reduced tree which is 1–acylindrical on non-cyclic subgroups and has N orbits of edges.*

These constructions use the fact that the set of subgroups on which the acylindrical condition applies (finite subgroups in the case of Theorem 2.4.5 and infinite cyclic groups in the case of Theorem 2.4.6) can have chains of subgroups $H_1 < H_2 < \cdots < H_n$ with arbitrary length. We say a set of subgroups $\mathcal{P}$ has *height* equal to the supremum of the length of these chains for groups in $\mathcal{P}$. (See Definition 3.1.8 for a full definition of the height of $\mathcal{P}$.) We also say an action of a group G on a tree T is $\mathcal{P}$*-closed* if whenever a subgroup $K \leq G$ is contained in a member of $\mathcal{P}$ and fixes an edge e of T there is some $H \in \mathcal{P}$ with $K \leqslant H \leqslant \operatorname{Stab} e$. In light of the constructions in Theorem 2.4.5 and Theorem 2.4.6 we now recast Conjecture 2.4.4 into the following.

Question 1.0.2. *Let G be a finitely generated group and let $\mathcal{P}$ be a conjugation invariant set of subgroups of G with finite height. Is there a constant $C(G,k,\mathcal{P})$ such that any reduced action of G on a tree which is P-closed and k–acylindrical on groups larger than $\mathcal{P}$ has at most $C(G,k,\mathcal{P})$ orbits of edges? (A group is* larger than *$\mathcal{P}$ if it's not contained in any member of $\mathcal{P}$.)*

In Chapter 3 we will prove Theorem 3.1.12 which gives a positive answer to the above in many cases. As the exact statement is somewhat technical we will skip it for now. Instead we will give a couple of interesting applications of this result in specific cases. In the first we consider hyperbolic groups which act k-acylindrically on subgroups which aren't virtually cyclic. In the second we consider right-angled Artin groups which act k-acylindrically on non-abelian subgroups.

Definition 3.1.3. *Suppose $\mathcal{P}$ is a class of subgroups of G which is closed under conjugation. A minimal action is said to be* partially-reduced *over $\mathcal{P}$ if either*

- *T/G is a circle consisting of a single vertex and edge; or*
- *whenever a vertex v of T has stabiliser equal to that of an edge and which is contained in a subgroup of a member of $\mathcal{P}$ then v/G has valence at least 3 in T/G.*

Definition 3.1.15. *Let G be a hyperbolic group. We say a virtually $\mathbb{Z}$ subgroup $H \leqslant G$ is m-*almost maximal *if whenever we have a virtually cyclic $K \leqslant G$ with $H \leqslant K$ then $[H : K] \leq m$. Let $\mathcal{P}_m$ be the collection of subgroups of G which are either finite or m-almost maximal.*

Corollary 3.1.16. *Suppose G acts on a tree T which is partially reduced on $\mathcal{P}_m$ and k–acylindrical on groups larger than $\mathcal{P}_m$. Suppose also that T is $\mathcal{P}_m$-closed. Then the number of edges of T/G is bounded above by $(2k+1)C'(G)$ (where $C'(G)$ is a constant determined by G).*

Corollary 3.1.18. *Let $G = A(\Gamma)$ be a RAAG. An abelian subgroup $H \leqslant G$ is said to be* rank maximal *if whenever we have an abelian subgroup $K \leqslant G$ which contains H with finite index we have $K = H$. Let $\mathcal{P}$ be the collection of rank maximal abelian subgroups of G. Suppose G acts on a tree T which is partially reduced on abelian subgroups and k–acylindrical on non-abelian subgroups. Suppose also that T is $\mathcal{P}$-closed. Then the number of edges of T/G is bounded above by $(2k+1)2^nC(G)$ (where n is the size of the largest complete subgraph of Γ and $C(G)$ is the same constant determined by Dunwoody's Resolution Lemma (Theorem 2.1.12)).*

We also give some attention to the original k–acylindrical case where the condition applies to all non-trivial subgroups. We have already mentioned Weidmann's result from [26], which gives a good bound of $2k(\operatorname{rank} G - 1)$ edges for the splitting. Theorem 3.1.12 recovers this bound immediately in this case. We will show that it is possible to tighten this bound further to $\left\lfloor \left(2\operatorname{rank} G - \frac{5}{2}\right) k \right\rfloor$ edges and even further to $\lfloor (2\operatorname{rank} G - 3)\, k \rfloor$ if the underlying group is torsion-free. Moreover we construct examples to show that these new bounds are sharp.

Theorem 3.5.1. *Let G be a (non-cyclic) finitely generated group acting k–acylindrically on a minimal tree T (where $k \geq 1$.) Suppose that each edge of T has non-trivial stabiliser. Then T/G has at most $\left\lfloor \left(2\operatorname{rank} G - \frac{5}{2}\right) k \right\rfloor$ edges. If G is torsion-free then this bound can be improved to $(2\operatorname{rank} G - 3)\, k$.*

Theorem 3.5.2. *For any $k > 0$ and $r \geq 2$ there is a finitely presented group G with $\operatorname{rank} G = r$ which acts k–acylindrically on a minimal tree T where each edge of T has non-trivial stabiliser and T/G has exactly $\left\lfloor \left(2\operatorname{rank} G - \frac{5}{2}\right) k \right\rfloor$ edges.*

Similarly F_r admits a k–acylindrical action on a minimal tree T where each edge of T has non-trivial stabiliser and T/F_r has exactly $(2r-3)\, k$ edges.

A natural extension of our original question about the properties of a single splitting is to ask what happens if we allow ourselves to recursively consider splittings of the vertex groups. We naturally get the notion of a *hierarchy* as a rooted tree with a group associated to each vertex, where the immediate descendants of a vertex correspond to the vertex stabilisers of a splitting of its group. We would like to know if our group has finite hierarchies with terminal vertices which have indecomposable groups and so is *strongly accessible* in some sense. Over finite edge groups this question immediately reduces to the regular accessibility question as finite subgroup always fixes a point of a tree [21]; however we run into problems as soon as we begin looking at infinite groups. For example since $F_2 \cong (F_2)*_{\mathbb{Z}}$ (as $F_2 \cong \langle a, b, c \mid cac^{-1} = b \rangle$) we can easily build an infinite hierarchy for free groups over cyclic edge groups. As such we instead try and show that some particular hierarchy with indecomposable terminal vertices is finite. For example the Haken hierarchy of a 3–manifold is finite [16].

Delzant and Potyagailo [9] attempted to show that such a finite hierarchy always exists for finitely presented 2–torsion-free groups over any elementary family of subgroups. Unfortunately their paper contains a fatal error which has been pointed out by Louder and Touikan [19]. In the same paper Louder and Touikan prove a weaker version of this result where an ascending chain condition is required to hold as well as showing that many hierarchies of JSJ-decompositions over slender edge groups are finite. (Recall that a group is *slender* if all its subgroups are finitely generated.) In Chapter 4 we give full detailed account of Louder and Touikan's theorem on the finiteness of JSJ-hierarchies for virtually 2-torsion free hyperbolic groups.

Theorem 4.2.1. *(Louder, Touikan [19, Corollary 2.7]) Let G be a hyperbolic group which is virtually* 2*–torsion-free. Then any JSJ-hierarchy for G is finite.*

Chapter 2

Background

The purpose of this chapter is to give an overview of the existing literature on accessibility. As such almost all of the material presented here is the work of others with citations given as appropriate. The exceptions are Theorem 2.4.5 and Theorem 2.4.6 which give original counterexamples to a previously open conjecture.

2.1 Splittings over finite edge groups

Firstly note that we can trivially make arbitrarily complicated splittings by adding additional "hanging" edges to a tree which our group acts on. For example the splitting $G \cong G *_H H$ where $H \leqslant G$ tells us nothing about the structure of the group G. The following definition prevents these sorts of trivialities.

Definition 2.1.1. *An action of a group G on a tree T is said to be* minimal *if T has no G-invariant proper subtrees.*

Our starting point will be Stallings theorem about ends of groups.

Theorem 2.1.2 (Stallings [23])**.** *A finitely generated group G admits a (minimal) splitting $G \cong A *_C B$ or $G \cong A*_C$ over a finite edge group if and only if G has multiple ends.*

This naturally leads to the following question. Suppose we have a finitely generated group G. If G is multi-ended then Theorem 2.1.2 says that G admits a non-trivial splitting over a finite edge group. For each of the vertex group of this splitting we can ask if it's multi-ended and if so we can split it over a finite group as well. We want to know if we keep iterating this procedure whether there must eventually come a point where all of the vertex groups have either zero or one end. If this is the case then we can refine our initial splitting to get a new splitting for G where all of the edge groups are finite and the vertex groups have either zero or one end. This leads us nicely to the following definition.

Definition 2.1.3. *A group G is said to be* accessible *if it has a finite graph of groups decomposition where the edge groups are finite and the vertex groups all have zero or one end.*

The previous question can now be neatly expressed as the following.

Question 2.1.4. *Are all finitely generated groups accessible?*

Now another trivial way of creating arbitrarily complicated splittings is by repeatedly subdividing an edge of an existing one. For example if $G \cong A *_C B$ then we also have $G \cong A *_C C *_C \cdots *_C B$ which gives no new information about the structure of the group. We introduce a new restriction to prevent these sorts of pathologies.

Definition 2.1.5. *A minimal action of a group G on a tree T is said to be* reduced *if either:*

- *T/G is a circle consisting of a single edge and vertex; or*
- *whenever a vertex v has the same stabiliser as that of an edge e the image of v in T/G has valence of at least* 3.

The following easy fact gives a useful criterion for a group to be accessible.

Proposition 2.1.6. *A finitely generated group G is accessible if there is a constant $C(G)$ such that any reduced graph of groups decomposition for G with finite edge groups has at most $C(G)$ edges.*

In [18] Linnell gives a partial answer to Question 2.1.4. They show that a finitely generated group is accessible if its finite subgroups have bounded order. This includes many important classes of groups such as finitely generated linear groups (in characteristic 0), hyperbolic groups and any finitely generated torsion-free group. More precisely they prove the following.

Theorem 2.1.7 (Linnell Accessibility [18])**.** *Let G be a finitely generated group and $n \in \mathbb{N}$. Then there is a constant $C_n(G)$ such that any reduced graph of groups decomposition for G where the edge groups have order bounded above by n has at most $C_n(G)$ edges.*

Later Dunwoody [10] showed that Question 2.1.4 is false in its full generality. They produce an example of a finitely generated group which admits splittings over finite edge groups with an arbitrary number of edges.

Theorem 2.1.8 (Dunwoody [10])**.** *There's a finitely generated group which isn't accessible.*

In a separate paper Dunwoody also showed that all finitely presented groups are accessible [11]. In fact they showed this is true for a slightly larger class of groups.

Definition 2.1.9. *A finitely generated group G is said to be* almost finitely presented *if it acts freely and cocompactly on a simplicial complex with $H^1(X) = 0$.*

Theorem 2.1.10 (Dunwoody [11])**.** *Let G be an almost finitely presented group. There is a constant $C(G)$ such that any reduced graph of groups decomposition for G where the edge groups have finite order has at most $C(G)$ edges.*

Dunwoody's proof of Theorem 2.1.10 introduced a few important ideas. In particular they give us a very useful result showing the existence of certain resolutions.

Definition 2.1.11. *A* combinatorial map *$\Psi : S \to T$ is a G–equivariant map between trees where each vertex gets sent to a vertex and each edge $e = [u, v]$ gets sent to the reduced edge path from $\Psi(u)$ to $\Psi(v)$.*

Theorem 2.1.12 (Dunwoody's resolution lemma)**.** *For any almost finitely presented group G there's a constant $C(G)$ so that the following holds. Suppose G acts on a minimal tree T. Then G acts on another minimal tree T' with at most $C(G)$ orbits of edges and there's a combinatorial map $\alpha : T' \to T$. (Without loss of generality we can assume that no edge gets mapped to a point.)*

Remark 2.1.13. *Strictly speaking Theorem* 2.1.12 *doesn't appear in Dunwoody's paper [11]. However it's more or less an immediate consequence of the arguments it contains.*

We will now give an account of the proof of Dunwoody's resolution lemma (Theorem 2.1.12). We begin by introducing the notion of a track, which can be thought of as a generalisation of a 1-dimensional sub-manifold of a surface.

Definition 2.1.14. *Let X be a connected* 2*-dimensional simplicial complex. A* track *$R \subset X$ is a connected subset with the following intersection properties.*

- *R doesn't contain any vertex of X.*
- *For each* 2*-simplex σ of X the intersection $\sigma \cap R$ is the union of finitely many intervals joining distinct edges of σ.*
- *For each edge e of X which is not a face of any* 2*-simplex the intersection $e \cap R$ is either empty or consists of a single point in the interior of e.*

Definition 2.1.15. *Let X be a connected 2-dimensional simplicial complex. A* band $B \subset X$ *is a closed connected subset with the following intersection properties.*

- *B doesn't contain any vertex of X.*
- *For each 2-simplex σ of X the intersection $\sigma \cap B$ is the union of finitely many components consisting of quadrilaterals where two of the edges are closed subsets of distinct edges of σ. (See Figure 2.1.)*
- *For each edge e of X which is not a face of any 2-simplex the intersection $e \cap B$ is either empty or consists of a non-trivial closed interval in the interior of e.*

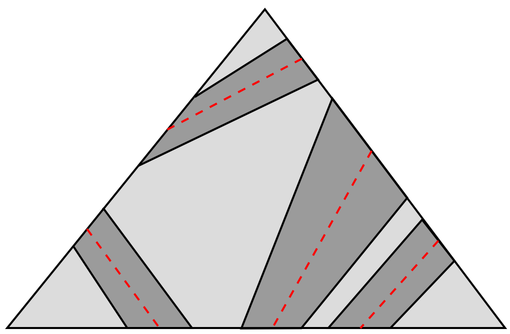

Fig. 2.1 An example of the intersection of a band with a triangle. The shaded area represents the band and the dashed line represents its midtrack.

Observe that a band is essentially just the product of a track with an interval.

Definition 2.1.16. *For B a band we define its* midtrack *to be the track obtained by taking the midpoints of each component of $e \cap B$ for each edge e in X and joining them by lines corresponding to the components of $\sigma \cap B$ for each 2-simplex σ in X. (See Figure 2.1.)*

Definition 2.1.17. *A band B is said to be* twisted *if its boundary ∂B is connected. Otherwise ∂B consists of two disjoint tracks and we say that B is* untwisted*. A track is* (un)twisted *if it's the midtrack of an (un)twisted band. Two disjoint tracks R_1 and R_2 are said to be* parallel *if there's an (untwisted) band B with $\partial B = R_1 \cup R_2$.*

We will now show that a complex with lots of tracks on it must split into many distinct pieces.

Lemma 2.1.18. *[11, Theorem 2.1] Let X be a connected simplicial complex and suppose $\beta = dim(H^1(X,\mathbb{Z}_2))$ is finite. Let $\{T_1,\cdots,T_n\}$ be a set of disjoint tracks on X. Then at most β of the T_i are twisted and $X \setminus (\coprod T_i)$ contains at least $n-\beta$ connected components.*

Proof. Suppose T is a track on X. We define $z_T \in H^1(X,\mathbb{Z}_2)$ to be the class containing the cocycle which sends an edge e to $|T \cap e| \pmod 2$. Observe that z_T is the class containing the zero cocycle if and only if T separates X; thus $z_T = 0$ implies that T is an untwisted track. Moreover any finite collection of twisted tracks cannot separate X; so the twisted tracks give rise to a linearly independent set in $H^1(X,\mathbb{Z}_2)$. In particular there are at most β twisted tracks.

Consider the map $\sigma : \mathbb{Z}_2 z_{T_1} \oplus \cdots \oplus \mathbb{Z}_2 z_{T_n} \to H^1(X,\mathbb{Z}_2)$. Suppose $X \setminus (\coprod T_i)$ has k connected components. For each component of $X \setminus (\coprod T_i)$ we get an element of $\ker\sigma$ defined as the sum of its boundary tracks; call these elements $\{\tilde{z}_1, \cdots, \tilde{z}_k\}$. Observe that any member of $\ker\sigma$ can be represented as a collection of the tracks T_i which separate X; hence can be represented by a sum of the $\tilde{z}_j$. So $\{\tilde{z}_1, \cdots, \tilde{z}_k\}$ is a spanning set for $\ker\sigma$ and hence by rank-nullity we see that $k \geq n - \beta$. □

If a complex has many disjoint tracks it seems likely that many pairs of these tracks will be parallel. As such if we split this complex into lots of components along these tracks many of resulting components will just be untwisted bands. These bands tell us nothing about the topology of the complex and so we can ignore them. We formalise this into the following theorem.

Lemma 2.1.19. *[11, Theorem 2.2] Let X be a finite simplicial complex and $\beta = dim(H^1(X))$. Let $C(X) = 2\beta + N_V + N_T$ where N_V is the number of vertices of X and N_L is the number of 2-simplices. Suppose $T_1, \cdots, T_{C(X)+1}$ is a collection of pairwise disjoint tracks on X. Then there are distinct i and j such that T_i and T_j are parallel.*

Proof. We will consider the structure of the connected components of $X \setminus (\coprod T_i)$ and show that there must be many bands. First let σ be any 2-simplex of X and observe that a connected component Y of $\sigma \setminus (\sigma \cap (\coprod T_i))$ must be of one of the following three forms. (See Figure 2.2.)

- Y contains one of the vertices of ρ. (There are at most 3 of these.)
- Y is a quadrilateral where two of the sides are parts of tracks and the other two sides are disjoint parts of different edges of σ.
- Y is a hexagon where three of the sides are parts of tracks and the other three sides are disjoint parts each of the edges of σ. Call such a component *bad* and note that at most a single component can be of this form.

Now we look at the components of $X \setminus (\coprod T_i)$. Lemma 2.1.18 implies that there must be at least $C(X) + 1 - \beta$ such components. Now at most N_V of these components contain a vertex

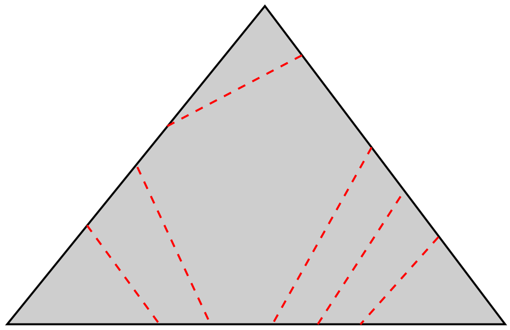

Fig. 2.2 An example of the intersection of σ with a triangle T. Observe that there is always exactly one component of $\sigma \backslash T$ which intersects each edge. (The "central" component.)

and at most N_T of them contain a 'bad' piece. The rest of the components are bands and so we have a collection of at least $(C(X)+1)-(\beta+N_V+N_T)=\beta+1$ bands with disjoint interiors and boundaries consisting of the tracks T_i. Applying Lemma 2.1.18 to the midtracks of these bands we see that at least one of them is untwisted; which exactly says that two of the T_i are parallel. □

We are now ready to prove Dunwoody's resolution lemma. (Theorem 2.1.12)

Proof of Theorem 2.1.12. Let X be any simplicial complex of dimension 2 with $H^1(X,\mathbb{Z}_2) = 0$ which G acts on freely and cocompactly by simplicial automorphisms. Pick any representative set of vertices V_0 for X (under the action of G) and pick any map $\rho : V_0 \to T$ which maps into the vertices of T. First extend this map equivariantly to every vertex of X, then to each edge of X by mapping to T linearly. It remains to extend ρ to each triangle of X. Pick a triangle Δ in X with vertices v_1, v_2 and v_3. If the v_i all get mapped to the same vertex then map all of Δ to this vertex. Now suppose that the v_i all get mapped to different vertices in T. We will define ρ on Δ by determining $\rho^{-1}(y)$ for each y in the image of $\partial\Delta$. Set $\rho^{-1}(y)$ to be the polygon in Δ where the vertices are the points on $\partial\Delta$ which map to y. (This polygon will be an interval for all choices of y except possibly for a single point x which may be a triangle and the images of the v_i which may be a single point. See Figure 2.3.) Each point of Δ lies in exactly one these polygons and so ρ is well defined. The case where exactly two of the v_i gets sent to the same point is essentially the same. (See Figure 2.4.)

Let m_e be the midpoint of an edge e of T. Observe that $\rho^{-1}(m_e)$ consists of a disjoint collection of tracks on X. Let λ be one of the tracks of $\rho^{-1}(m_e)$. By considering ρ we see that a neighbourhood of $\lambda \subset X$ is isomorphic to $\lambda \times I$ where I is an interval representing the

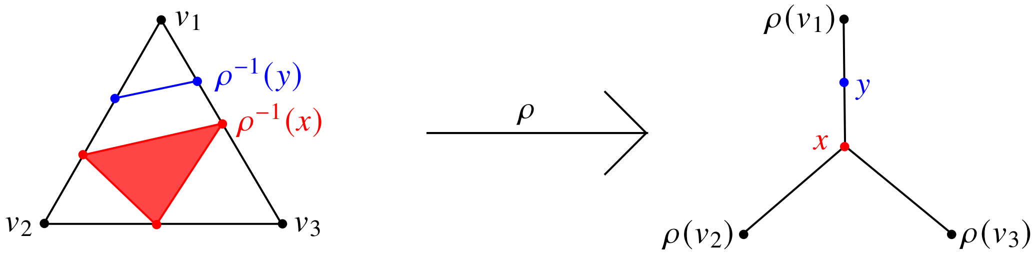

Fig. 2.3 An example of extending ρ over a triangle where its vertices get sent to different points in T.

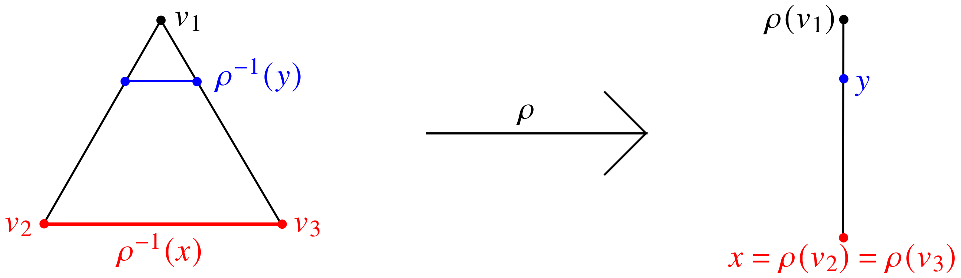

Fig. 2.4 An example of extending ρ over a triangle where exactly 2 vertices get sent to the same vertex in T.

middle third of e. In particular we see that λ is untwisted. Let $\Lambda := \coprod_e \rho^{-1}(m_e)$ and let $\Lambda^* \subseteq \Lambda$ be a maximal G-invariant subset of non-parallel tracks for X. Define a graph T' where the vertices correspond to the connected components of $X \setminus \Lambda^*$ and the edges correspond to the tracks in Λ^*. (The edges connect to the vertices in the obvious way.) Since $H^1(X, \mathbb{Z}_2) = 0$ each track in Λ^* splits X and so T' is a tree. Also the map ρ induces a combinatorial map $\alpha : T' \to T$.

It remains to check that T' has boundedly many orbits of edges. Observe that Λ^* induces a finite non-parallel collection of tracks on X/G. Hence by Lemma 2.1.19 we see that Λ^* consists of at most $C(X)$ orbits of tracks, which correspond to orbits of edges in T'. □

2.2 Stallings folds

The idea of a fold will be of vital importance. Recall the following.

Definition 2.2.1. *[2] Let G act on a tree T. Let $e_1 = [x, y_1]$, $e_2 = [x, y_2]$ be distinct edges with a common endpoint x and let $\phi : e_1 \to e_2$ be the linear map which leaves x fixed. Let $\sim$ be the minimal equivalence relation on T such that $z \sim \phi(z)$ for each $z \in e_1$ and such that*

$T/\sim$ is naturally a tree on which G acts. Equivalently $\sim$ is the minimal equivalence relation such that $g \cdot z \sim g \cdot \phi(z)$ for each $z \in e_1$ and $g \in G$. A fold *is the map $T \to T/\sim$.*

Folds were introduced by Stallings in [22]. In particular they showed that maps between trees with finitely generated edge stabilisers can be decomposed into a finite sequence of folds. We will need something similar for trees whose edge groups need not be finitely generated. Fortunately we can "add in" generators of each edge group one at a time, then take a limit to see that our maps are a composition of (potentially infinitely many) folds. This is formalised into the following theorem, the proof of which is largely the same as the one given in [2, p.455] with the aforementioned limiting process to deal with the fact that the edge stabilisers aren't necessarily finitely generated.

Theorem 2.2.2 (Stallings folding theorem)**.** *Let G be a countable group. Suppose $\Psi : S \to T$ is a surjective simplical equivariant map between trees which G acts on with S/G finite and where no edge of S gets mapped to a point by Ψ. Then Ψ can be viewed as a (possibly infinite) composition of folds. i.e. $\alpha = \cdots \alpha_2 \alpha_1$ where each α_i is a fold (without edge inversions). Moreover if each edge group of T is finitely generated then this composition consists of only finitely many folds.*

Remark 2.2.3. *The codomain of $\cdots \alpha_2 \alpha_1$ is $S/\sim$ where $\sim$ is the equivalence relation generated by all the α_i. This is a tree which G acts on in the obvious way with vertex and edge stabilisers equal to the natural direct limit of their preimages.*

Proof. Throughout we'll let α_i fold the tree S_{i-1} into the tree S_i. Also we let $\beta_i := \alpha_i \circ \cdots \circ \alpha_1$ and γ_i be the map such that $\Psi = \gamma_i \circ \beta_i$.

First suppose that we have a surjective simplical map $m : A \to B$ between finite trees where no edge gets mapped to a point. Claim that m can be considered to be a finite series of folds; a fact we shall refer to as $(\star)$. If m is injective then as T is minimal we see that m is an isomorphism and so the result is trivial. Otherwise we have distinct vertices $u, v \in A$ with $m(u) = m(v)$. Let p be the reduced edge path from u to v. Since every edge of A is mapped to an edge in B and B is a tree we see that there must be a vertex $z \in p$ such that $m|_p$ is not locally injective at z. Let e_1 and e_2 be the edges in p which contain z as an endpoint and observe that $m(e_1) = m(e_2)$. Thus m factors though the fold with edges e_1 and e_2. Repeat this process until the map is injective, which must happen as the number of edges is finite and decreasing at each stage. This completes the decomposition of m into folds.

Now suppose we have an equivariant simplical map $\delta : R \to R'$. Let A be a finite subtree of R. We can apply $(\star)$ to $\delta|_A$ to obtain a finite series of folds $\delta' : R \to R''$ which factors through δ and where the corresponding $\delta'' : R'' \to R'$ is injective on $\delta'(A)$. This is how we will apply $(\star)$ in practice.

Our initial folds of Ψ will be to set S_i/G isomorphic as a graph to T/G for all sufficiently large i. Let F be the closure of a fundamental domain of $S = S_0$. We now use $(\star)$ to find a series of folds $\alpha_1, \cdots, \alpha_N$ such that $\gamma|_{\beta_N(F)}$ is a homeomorphism onto its image. Thus γ_N/G is a homeomorphism of graphs.

Let $K = \Psi(F)$. Let v be a vertex in K and $g \in \mathrm{Stab}(v)$. Let v_i be a preimage of v in F_i and let p_i be the reduced edge path from v_i to gv_i. Now we apply $(\star)$ to p_i in order to get folds $\alpha_i, \cdots, \alpha_j$ which get g in the relevant vertex group in S_j/G. Now repeat this process for each vertex $v \in K$ and $g \in \mathrm{Stab}(v)$. (One can use a diagonalization process to give an order which insures no combination of v and g is missed.) This potentially gives an infinite sequence of folds as K has finitely many vertices and each $\mathrm{Stab}(v)$ is countable.

Now let $\tilde{\gamma}$ be the map such that $\Psi = \tilde{\gamma} \circ (\cdots \circ \alpha_2 \circ \alpha_1)$. Claim that $\tilde{\gamma}$ is a homeomorphism. Indeed by construction we see that $\tilde{\gamma}$ induces a bijection between the orbits of vertices; moreover the stabiliser of each vertex is the same as that of its image. Hence $\tilde{\gamma}$ induces a bijection between the vertices and hence is a homeomorphism between trees. Thus $\cdots \circ \alpha_2 \circ \alpha_1$ is a decomposition of Ψ into folds. □

Remark 2.2.4. *It should be straightforward to extend this result to the case where G is uncountable and where S/G is not necessarily finite using the well ordering principle. We do not do this here because it is unnecessary to prove our main results.*

Remark 2.2.5. *The condition that no edge of T gets collapsed is not a restrictive one in practice. In particular if Ψ maps an edge of S to a point then let π be the map which collapses each edge of S which is sent to a point by Ψ. Then there is a natural composition $\Psi = \Psi' \circ \pi$ where no edge in the domain of Ψ' is sent to a point in T.*

Remark 2.2.6. *Theorem* 2.2.2 *also works for combinatorial maps instead of just simplicial ones. By subdividing the edges of S we can turn a combinatorial map into a simplicial one.*

Remark 2.2.7. *We can extend the notion of a fold in the following way. Pick some $H \leqslant \mathrm{Stab}\,x$. We define a* generalised fold *the same way as a regular fold except we now identify e_1 with he_2 for every $h \in H$. The advantage of this is that we get a version of Stallings folding theorem which only needs finitely many generalised folds. We forgo this notion here as the present author believes there is value in both approaches being represented in the literature.*

Depending on which of the vertices and edges are in common G–orbits there are a few different cases that can arise from a fold. The following classification of folds is the same as the one found in [2].

We make the distinction between whether x is in the same G orbit as one of the y_i. If x is not in the same G–orbit as either of the y_i we say that the fold is of type A. Otherwise WLOG we have $gx = y_1$ for some $g \in G$ and we say that the fold is of type B. Note that such a g must act hyperbolically on T, (with translation length 1), as it moves a vertex an odd distance.

Additionally we split each of these cases into three additional categories. We say the fold is of type I if y_1 and y_2 are in distinct orbits of G. We say the fold is of type II if e_1 and e_2 are in a common orbit of G. Finally we say the fold is of type III if y_1 and y_2 are in a common G–orbit, but e_1 and e_2 are not. We will now go into the specifics of each type of fold. Throughout we let e_i be the vertex between the vertices x and y_i and use capital letters to denote the group associated to the corresponding vertex or edge.

Remark 2.2.8. *The following diagrams represent what happens to the relevant subgraph of a particular graph of groups decomposition. Crucially the pictures for type I and III folds only give the correct groups if both e_1 and e_2 are in the fundamental domain for this decomposition. In general we need to conjugate certain groups in the decomposition before these pictures become accurate.*

Type I We have y_1 and y_2 in distinct orbits of G. In this case the number of vertices and the number of edges of the graph of groups decomposition both decrease by one so the Euler characteristic of the underlying graph stays the same.

Type II We have e_1 and e_2 in a common orbit of G, suppose that $he_1 = e_2$. Observe that if h acts hyperbolically on T then the action of G after the fold inverts an edge and so we will ignore this case. Thus we can assume that $h \in X$. In this case the underlying graph of the graph of groups decomposition doesn't change. Instead the element h gets "pulled" along the edge in the graph of groups decomposition.

Type III We have y_1 and y_2 in a common orbit, but e_1 and e_2 are not. Suppose that $hy_1 = y_2$. Observe that h has to act hyperbolically on T with translation length 2. After the fold this h now fixes the image of y_1 and y_2 thus no longer acts hyperbolically. This type

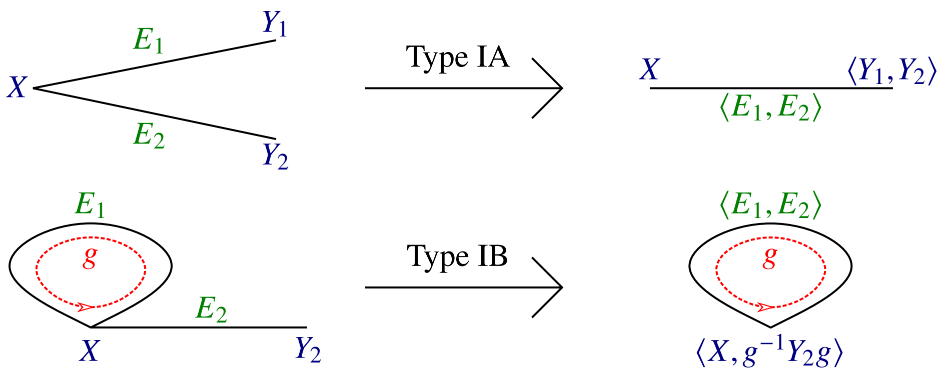

Fig. 2.5 A typical example of the effects of a type I fold on a graph of groups. The vertices y_1 and y_2 are inequivalent so the fold reduces the number of vertices by 1. Likewise for the edges e_1 and e_2.

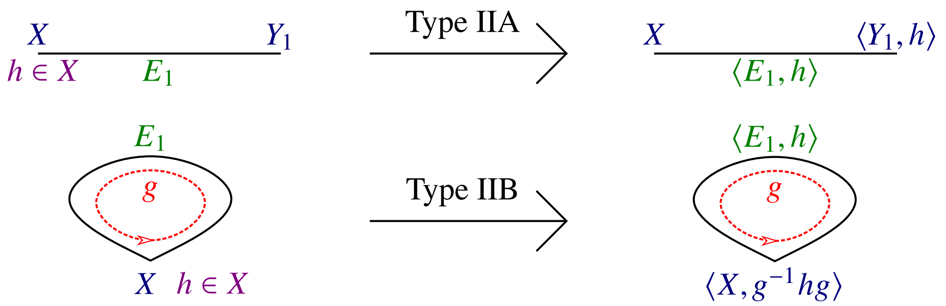

Fig. 2.6 A typical example of the effects of a type II fold on a graph of groups. The vertices y_1 and y_2 are equivalent so the fold keeps the number of vertices the same. Likewise for the edges e_1 and e_2.

of fold reduces the number of edges of the graph of groups by one while keeping the number of vertices fixed. Thus the Euler characteristic of the underlying graph increases by one.

Folds can also be used to construct interesting examples of splittings. For example the following appears in a paper of Bestvina and Feighn [2].

Theorem 2.2.9. *[2, pg.450] For any $N \in \mathbb{N}$ there is a reduced graph of groups decomposition for F_2 with N edges.*

Proof. Consider the one edge splitting corresponding to $F_2 \cong \mathbb{Z} * \mathbb{Z} \cong \langle a \rangle * \langle b \rangle$. First we subdivide this edge into N subedges. (In the diagrams we take $N = 4$.)

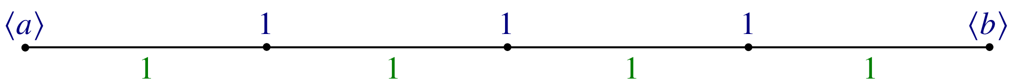

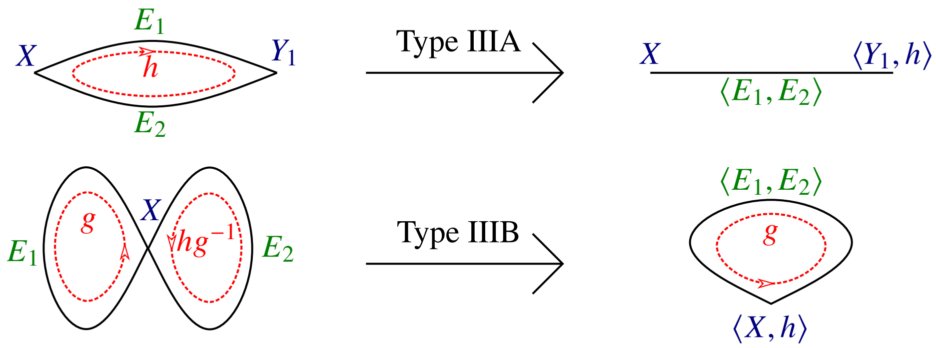

Fig. 2.7 A typical example of the effects of a type III fold on a graph of groups. The vertices y_1 and y_2 are equivalent so the fold keeps the number of vertices the same. However the edges e_1 and e_2 are inequivalent so the fold reduces the number of edges by 1.

Now we apply $N-1$ folds of type II. The first "pulls" a^2 across the first edge, the second "pulls" a^4 across the second edge and so on, so that the i^{th} fold "pulls" a^{2^i} across the i^{th} edge.

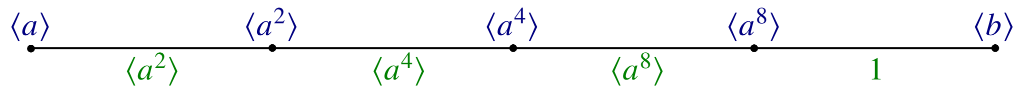

We now apply $N-1$ folds in the opposite direction. The first "pulls" b^2 across the first edge, the second "pulls" b^4 across the second edge and so on, so that the i^{th} fold "pulls" b^{2^i} across the i^{th} edge.

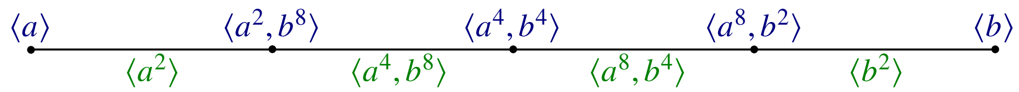

It's clear that this is a reduced decomposition of F_2 with N edges. □

2.3 Splittings over small edge groups

In light of Theorem 2.2.9 we see that it's impossible to extend Theorem 2.1.10 to arbitrary edge groups. In this section we will give a partial account of a theorem by Bestvina and Feighn [2], which gives a bound on the number of edges of a (reduced) splitting of an (almost) finitely presented group with "small" edge groups. We will only consider the case where the edge groups act elliptically on any tree as the core principles remain the same as the full theorem but with fewer technical details. Even this is still a generalisation of Dunwoody's result (Theorem 2.1.10) as finite groups always fix a vertex of a tree they act on [21].

We will begin by recalling the different actions a finitely generated subgroup $K \leqslant G$ can have on a tree T. A group element $g \in G$ either acts trivially on some subtree of T or has an invariant *axis* consisting of the points which are moved a minimal amount by g [21]. In the former case we call g *elliptic* and in the latter case we say g is *hyperbolic*. This leads to the following well known classification of actions, which is similar to the one given in [2] except with a distinction between linear and parabolic actions.

- If every element of K is elliptic then there is some point in T which is fixed by all of K [21]. We call such an action *elliptic*.

- Suppose every hyperbolic element of K has a common axis. We call this action *linear* if the ends of this axis are fixed and *dihedral* if they are not. A linearly acting group can be written in the form $K \cong E*_E$ where E is a subgroup of an edge group of T and the inclusion maps are isomorphic. Meanwhile a dihedral group can be written in the form $K \cong A *_E B$ where E is a subgroup of an edge group of T and both A and B contain E as an index 2 subgroup.

- Suppose that the axes of any two hyperbolic elements of K have infinite intersection, but that no line is fixed by K. We call such a K *parabolic*. Such a group fixes a single point of ∂T and is a strictly ascending HNN-extension $K \cong E*_E$ where E is a subgroup of an edge group of T. Observe that K has an infinitely generated subgroup which is generated by $\left\{a^{t^n} \mid n \in \mathbb{Z}\right\}$ where t is the stable letter of the HNN-extension and $a \in E$ is not contained in the non surjective end of the HNN-extension.

- Suppose that K contains two hyperbolic elements whose axes have compact (possibly empty) intersection. We call this action of K *hyperbolic*.

Definition 2.3.1. *A group G is* large *if it acts on some tree T hyperbolicly. A group is* small *if it's not large. A group has* Serre's property (FA) *if all its actions are elliptic.*

Remark 2.3.2. *The ping-pong lemma implies that any large group must contain F_2 as a subgroup. The converse is not true; for example $SL_3(\mathbb{Z})$ contains many subgroups isomorphic to F_2 and has Serre's property (FA) [21].*

Now that we have the definition of a small and large group we are ready to state the main result of this section.

Theorem 2.3.3. *(Bestvina, Feighn [2, Main theorem]) Let G be an almost finitely presented group. There is a constant $C'(G)$ such that any reduced graph of groups decomposition for G where the edge groups are small has at most $C'(G)$ edges.*

As mentioned in the introduction to this section we will restrict our attention to the case where the edge stabilisers have Serre's property (FA). As such we will only prove the following weaker statement here.

Theorem 2.3.4. *(Bestvina, Feighn [2, Section 4]) Let G be an almost finitely presented group. There is a constant $C'(G)$ such that any reduced graph of groups decomposition for G where the edge groups have Serre's property (FA) has at most $C'(G)$ edges.*

The main idea behind the proof of this theorem is to use Dunwoody's resolution lemma (Theorem 2.1.12) and Stallings' folding theorem (Theorem 2.2.2) to determine a set of vertices of the splitting which essentially gives us all the information about possible vertex groups. (The so called "live" vertices.) We will then show that there can't be too long of a gap between these "live" vertices, thus bounding total number of vertices in the splitting.

Definition 2.3.5. *Let G act on trees T and T' where there's a surjective combinatorial map $\alpha : T' \to T$ where no edge gets mapped to a point. Subdivide the edges of T' so that α becomes a simplicial map and decompose α into folds $\alpha = \alpha_n \circ \cdots \circ \alpha_2 \circ \alpha_1$. A vertex v of T is said to be* live *(with respect to our chosen decomposition of α) if either:*

- *there's a vertex v' of T' before subdivision with $\alpha(v') = v$; or*
- *there's some fold α_i of type III which folds together $[x, y_1]$ and $[x, y_2]$ to get $[x', y']$ where $\alpha_n \circ \cdots \circ \alpha_{i+1}(y') = v$.*

A vertex of T is said to be dead *if it's not live.*

The following simple results are rather useful.

Proposition 2.3.6. *Let G act on a tree T and suppose T/G is finite with v_1 vertices of valence 1. Then T/G has at most $2(\beta_1(T/G) - 1) + v_1$ vertices of valence at least 3.*

Proof. Suppose T/G has v_i vertices of valence i. By summing the valency of each vertex we see that $\sum_i i v_i$ is equal to twice the number of edges of T/G. Hence by the definition of the Euler characteristic we see that

$$\begin{aligned} \chi(T/G) = 1 - \beta_1(T/G) &= \sum_i \left(1 - \frac{i}{2}\right) v_i \\ &\leq \frac{1}{2}\left(v_1 - \sum_{i \geq 3} v_i\right) \end{aligned}$$

Elementary rearrangement of terms now gives the result. □

Proposition 2.3.7. *[2, Lemma 4 (i)] Let G act on a minimal tree T. Suppose $W \subseteq G$ is a set of elements which act elliptically on T and $W \cup \{g_1, \cdots, g_m\}$ normally generates G. Then $m \geq \beta_1(T/G)$. In particular by taking $\{g_1, \cdots, g_m\}$ to be a minimal generating set for G we get that $\beta_1(T/G) \leq \operatorname{rank} G$.*

Proof. Consider a graph of groups decomposition corresponding to T/G. By replacing each vertex and edge group with the trivial group we get a surjective group homomorphism $\sigma : G \to F_{\beta_1(T/G)}$. Consider the map $\mathrm{Ab}(\sigma) : \mathrm{Ab}(G) \to \mathbb{Z}^{\beta_1(T/G)}$ obtained by abelianization. Let $H = \langle g_1, \cdots, g_m \rangle$. Then since $W \cup H$ normally generates G and $W \subseteq \ker\sigma$ we see that $\mathrm{Ab}(\sigma)|_{\mathrm{Ab}(H)}$ is also surjective. Hence $\{\mathrm{Ab}(\sigma)(g_1), \cdots, \mathrm{Ab}(\sigma)(g_m)\}$ is a generating set for $\mathbb{Z}^{\beta_1(T/G)}$; which implies that $m \geq \beta_1(T/G)$ as $\mathrm{rank}\left(\mathbb{Z}^{\beta_1(T/G)}\right) = \beta_1(T/G)$. □

Remark 2.3.8. *Proposition* 2.3.7 *implies that there's a bound on the number of live vertices of T. More precisely the number of folds of type III in a folding decomposition of $T' \to T$ is bounded above $\beta_1(T'/G) \leq \mathrm{rank}\, G$. (This is easily seen by looking at the effect that each type of fold has on the Euler characteristic of the graph of groups.) Hence if T' is given to us by Dunwoody's resolution lemma (Theorem* 2.1.12*) we see that T has at most $C'(G) + \mathrm{rank}\, G$ (orbits of) live vertices, (where $C'(G)$ is the constant determined by Theorem* 2.1.12*). Thus in order to prove Theorem* 2.3.3 *we just need to bound the number of dead vertices.*

The following lemma justifies our previous assertion that the stabilisers of the live vertices contain all the useful information.

Lemma 2.3.9. *[2, Dead vertex lemma] Let v be a dead vertex in T (with respect to some decomposition of $\alpha : T' \to T$).*

(a) The stabiliser $\mathrm{Stab}(v)$ *is generated by the stabilisers of the edges containing v. So the vertex label for the corresponding vertex in the graph of groups is* normally *generated by the corresponding edge labels.*

(b) Whenever a finitely generated subgroup $W \leqslant$ $\mathrm{Stab}(v)$ *acts elliptically on T' we have W contained in the edge stabiliser of an edge containing v.*

Proof. We will prove (a) by induction on the number of folds in the decomposition of α; say $\alpha = \alpha_n \circ \cdots \circ \alpha_2 \circ \alpha_1$. First observe that a vertex of T' which is obtained by subdivision has the same stabiliser as its two edges. Now assume the result holds for $\alpha_{n-1} \circ \cdots \circ \alpha_2 \circ \alpha_1$. Suppose $\alpha_n : T'' \to T$ is a fold of type I or II. Then from our prior discussion on the properties of different folds from page 19 we see that the stabiliser of every vertex of T is generated by the union of the stabilisers of its edges and its preimages in T''. The result now follows as the preimage of a dead vertex is also a dead vertex. Instead suppose α_n is a type III fold which identifies the edges $e_1 = [x, y_1]$ and $e_2 = [x, y_2]$ to $e' = [x', y']$. The above discussion still applies to every vertex in T except vertices in the same orbit as y'. However by definition the vertex y' is live and so the result follows as before.

For (b) simply note that any finitely generated group acting elliptically on T' fixes a vertex of it; thus fixes a live vertex in T. So if $W \leqslant \mathrm{Stab}(v)$ acts elliptically on T' then it fixes a non-empty edge path from v to a live vertex of T. □

We will need to examine how the stabilisers of dead vertices act on T'. The following will give us a very controlled structure in this case when combined with the dead vertex lemma.

Lemma 2.3.10. *[2, Lemma 4 (ii)] Suppose G is a finitely generated group which acts with a reduced action on a tree S. Let $\mathfrak{V} := \{v_1, \cdots, v_n\}$ be a set of inequivalent vertices for S and let $V_i \leqslant \mathrm{Stab}\, v_i$. Suppose that every vertex stabiliser of S is conjugate into one of the V_i. Then there are subsets $\mathfrak{V}_1, \mathfrak{V}_2 \subseteq \mathfrak{V}$ which are sets of representatives for the valence 1 and valence 2 vertices of S/G respectively. Moreover $V_i = \mathrm{Stab}\, v_i$ whenever $v_i \in \mathfrak{V}_1 \cup \mathfrak{V}_2$.*

Proof. Observe that for any w in T we have a path from w to a vertex in the orbit of some v_i which is stabilised by $\mathrm{Stab}\, w$. If the image of w has valence 1 in S/G then this path must be trivial as T is minimal, hence w is in the same orbit as a v_i and hence there's a subset $\mathfrak{V}_1 \subseteq \mathfrak{V}$ which is a set of representatives for the valence 1 vertices of S/G.

Now if v_i is in $\mathfrak{V}_1$ then observe that there is some $v_j \in \mathfrak{V}$ and $g \in G$ such that $\mathrm{Stab}\, v_i \leqslant V_j^g$. The same argument as before implies that $i = j$ and $g \in V_i$. Hence we have $V_i \leqslant \mathrm{Stab}\, v_i \leqslant V_i$ and hence $\mathrm{Stab}\, v_i = V_i$.

An essentially identical argument works for the existence and properties of $\mathfrak{V}_2$. □

Proof of Theorem 2.3.4. Let T be a tree which G acts on where all the edges are stabilised by a group with property (FA). Dunwoody's resolution theorem (Theorem 2.1.12) says there is a tree T' with at most $C(G)$ orbits of edges together with a combinatorial map $\alpha : T' \to T$. Using Stallings' folding theorem (Theorem 2.2.2) we fix some decomposition of α into folds $\alpha_N \circ \cdots \circ \alpha_1$. We use this decomposition to label the vertices of T either dead or alive.

Recall that Remark 2.3.8 says that we have a bound on the number of live vertices and so we only need to bound the number of dead ones. Moreover Proposition 2.3.6 says that we get a bound on the number of vertices of valence 3 or greater as long as there is bound on the number of vertices of valence 1. In particular it's enough to bound the number of dead vertices of valence 1 or 2.

First we wish to show that any vertex $v \in T$ of valence 1 in T/G is live. Suppose instead that v is dead. Let e be a representative edge with endpoint v. Let $V = \mathrm{Stab}\, v$ and $E = \mathrm{Stab}\, e$.

By Lemma 2.3.9 (a) we see that V is normally generated by E. Let T'_V be a (reduced) minimal subtree of T' which is fixed by V. By Lemma 2.3.9 (b) we see that every vertex stabiliser of T'_V is conjugate into E. As E acts elliptically on T' we use Lemma 2.3.10 to see that T'_V has at most a single vertex of valence 1 and $\beta_1(T'_V/V) = 0$. Thus T'_V is a single vertex and $E = V$, contradicting the fact that T is reduced.

Now we examine a dead vertex $v \in T$ with valence 2 in T/G. As before let e_1 and e_2 be inequivalent representative edges with endpoint v. Let $V = \operatorname{Stab} v$ and $E_i = \operatorname{Stab} e_i$. By Lemma 2.3.9 (a) we see that V is normally generated by $E_1 \cup E_2$. Let T'_V be a (reduced) minimal subtree of T' which is fixed by V. By Lemma 2.3.9 (b) we see that every vertex stabiliser of T'_V is conjugate into either E_1 or E_2. As the E_i act elliptically on T' we use Lemma 2.3.10 to see that T'_V/V has at most two vertices of valence 1 and these have labels conjugate to the E_i. Moreover T'_V/V contains no vertices of valence greater than 1 and $\beta_1(T'_V/V) = 0$. If T'_V is a point then we get a contradiction as in the valence 1 case. Hence T'_V/V corresponds to the non-trivial splitting $V = E_1 *_{E_1 \cap E_2^g} E_2^g$ for some $g \in V$.

We now wish to show that there cannot be two adjacent dead vertices of valence 2 in T/G. Once we have shown this Theorem 2.3.4 will follow immediately. Indeed our argument so far gives us gives us a bound on the number of edges of a modified version of T/G where we "ignore" the dead valence 2 vertices by "combining" its two edges. If dead vertices of valence 2 cannot be adjacent then the actual number of edges of T/G will be at most twice the number of this modified graph.

Suppose $v_1 \in T$ and $v_2 \in T$ are adjacent dead vertices both with valence 2 in T/G. Suppose e is the edge connecting the v_i and let e_i be a representative edge with endpoint v_i which is not in the same orbit as e. Let $V_i = \operatorname{Stab} v_i$, $E_i = \operatorname{Stab} e_i$ and $E = \operatorname{Stab} e$. Now as in the proof of Lemma 2.3.9 (b) there is a path from a live vertex of T to either v_1 or v_2, (without loss of generality it contains v_1 and this path doesn't include e), which is fixed by E. Let e' be the final edge of this path. If e' is equivalent to e_1 then E is conjugate into E_1 and so by Lemma 2.3.10 we see that $V_1 = E_1$, which cannot happen as T is reduced. Instead suppose $e' = ge$ where $g \in V_1$. By the previous paragraph recall the tree T_{V_1} corresponding to the non-trivial splitting $V_1 = E_1 *_{E_1 \cap E} E$. Now if $g \notin E$ we see that E is contained in a conjugate of E_1 and so fixes an edge in T_V; which cannot happen as T_V is reduced. Hence $g \in E$ and so $e' = e$; but we already assumed this wasn't the case. □

Remark 2.3.11. *In order to extend the argument to Theorem* 2.3.3 *we consider each class of edges in turn. This is done by collapsing each edge which doesn't fit into the relevant* *class. We have shown the elliptically acting case already. The parabolically acting case (which includes linear actions) is similar to the elliptic case with a moderate amount of extra steps. The dihedral case by contrast is significantly harder. It is curious that dihedral actions pose a challenge both here as well as for strong accessibility (see Chapter 4), but there is no apparent reason as to why this is the case.*

2.4 Acylindrical actions

So far we've obtained bounds on the number of edges given a restriction on the class of allowed edge groups. We will now turn our attention to acylindrical actions, where there is no restriction on the class of edge groups and instead there is a bound on the diameter of the fixed point set of subgroups. More precisely we have the following definition.

Definition 2.4.1. *(Sela) Let G be a group and k be a non-negative integer. An action of G on a tree T is k–*acylindrical *if whenever some non-trivial $H \leqslant G$ fixes every edge in a reduced edge path p we have p containing at most k edges. Likewise an action is (k, C)–*acylindrical *(for some positive integer C) if whenever some $H \leqslant G$ fixes every edge in a reduced edge path p containing more than k edges we have $|H| \leq C$.*

The first example of a bound for acylindrical actions is due to Sela [20], who showed that there is a bound for the size of a (minimal) splitting where the action on the corresponding Bass-Serre tree is k–acylindrical, assuming that G is freely indecomposable and finitely generated. Later Delzant [8] showed that a bound for (k, C)–acylindrical actions exists provided that the acting group is finitely presented. Weidmann [25] then extended this to finitely generated groups and gives a nice bound in the process.

Theorem 2.4.2. *(Weidmann [25]) Let G be a finitely generated group and suppose G acts (k, C)–acylindrically (where $k \geq 1$) on a tree T and that this action is reduced. Then there is some $C(G)$ (which depends only on G) such the that number of edges of T/G is bounded above by $(2k+1)2^{\lfloor \log_2 C \rfloor} C(G)$.*

Remark 2.4.3. *Observe that Theorem 2.4.2 is a generalisation of Linnell accessibility (Theorem 2.1.7) by setting $k = 1$.*

The methods used in Chapter 3 are essentially a refined version of Weidmann's arguments and so will not be stated in detail here. However it involves using Stallings folding theorem (Theorem 2.2.2) in a way which is reminiscent of Bestvina and Feighn's result for actions with small edge stabilisers (Theorem 2.3.3). As a result Weidmann suggested that some sort of common generalisation between Theorem 2.4.2 and Theorem 2.3.3 should exist. We formalise this into the following conjecture.

Conjecture 2.4.4. *[25, pg.213] Given a finitely presented group G and $k > 0$ there some $C(G,k)$ such that any reduced action of G which is k–acylindrical on large subgroups (see Definition* 3.1.1 *for a precise definition) has at most $C(G,k)$ orbits of edges.*

The rest of this section is dedicated to constructing counterexamples to the above conjecture. In contrast to the rest of this chapter these constructions are original.

Theorem 2.4.5. *There is a finitely presented group G which for any $N > 0$ acts on a reduced tree which is* 1*–acylindrical on infinite subgroups and has N orbits of edges.*

Proof. Let $D := \left\langle a_1, a_2, \cdots \mid a_1^2 = 1,\, a_{i+1}^2 = a_i \ \forall i \geq 1 \right\rangle \cong \mathbb{Z}\left[\frac{1}{2}\right]/\mathbb{Z}$; the additive group of dyadic rationals modulo $\mathbb{Z}$ or the Prüfer 2-group. Let A be any finitely presented group into which D embeds; for example we can take A to be Thompson's group T [3]. Let $B := \langle b \rangle \cong \mathbb{Z}$. Take $G := A * B$ and pick any $N > 0$. Start by taking the one edge splitting corresponding to $G \cong A * B$ and subdividing this edge into N subedges. (In the diagrams we take $N = 4$.)

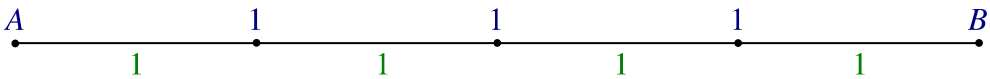

Now we apply N folds of type II. The first "pulls" a_N across the first edge, the second "pulls" a_{N-1} across the second edge and so on, so that the i^{th} fold "pulls" a_{N+1-i} across the i^{th} edge.

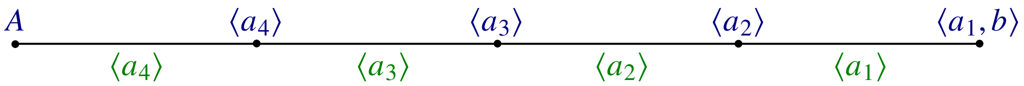

Let $b_0 := b$ and for $1 \leq i \leq N$ we define $b_i := b_{i-1} a_i b_{i-1}^2$. We now apply $N-1$ folds in the opposite direction. The first "pulls" b_1 across the first edge, the second "pulls" b_2 across the second edge and so on, so that the i^{th} fold "pulls" b_i across the i^{th} edge.

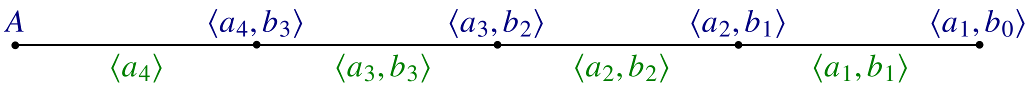

It's clear that this is a reduced decomposition. It remains to show that the action on the corresponding Bass-Serre tree is 1–acylindrical on infinite subgroups. In other words it suffices to show that the intersection of the stabilisers of any two distinct edges with a common end vertex are finite. Observe that a generic vertex of this decomposition has label $\left\langle a', b' \mid a'^{2^r} \right\rangle \cong \mathbb{Z} * (\mathbb{Z}/2^r\mathbb{Z})$ with two edges with labels $\left\langle a'^2, b' \right\rangle$ and $\left\langle a', b' a' b'^2 \right\rangle$ respectively.

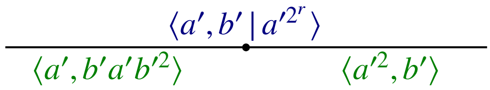

So there are three different pairs of edges we need to consider.

Case 1 The intersection of $\langle a'^2, b'\rangle$ and $\langle a'^2, b'\rangle^g$ for $g \in \langle a', b'\rangle \setminus \langle a'^2, b'\rangle$.

Let $\tilde{w} = \alpha_1\alpha_2\cdots\alpha_n$ and $\tilde{w}^h$ be cyclicly reduced words in $\{a'^{\pm 2}, b'^{\pm 1}\}$ for some $h \in \langle a', b'\rangle$. Observe by cycling letters in $\{a'^{\pm 1}, b'^{\pm 1}\}$ that either $h \in \langle a'^2, b'\rangle$ or $\alpha_i = a'^{\pm 2}$ for all i and h is an (odd) power of a'. It follows that every element in the intersection is conjugate to a power of a'. Moreover we see that $w \in \langle a'^2, b'\rangle \cap \langle a'^2, b'\rangle^g$ if and only if there are $g_1, g_2 \in \langle a'^2, b'\rangle$ and $n \in \mathbb{Z}$ such that $g = g_1 a' g_2$ and $w = (a'^{2n})^{g_2}$.

If we can show that g can only be expressed in the form $g_1 a' g_2$ in an "essentially unique" way then it follows that the intersection is cyclic and hence finite as a'^2 has finite order. More precisely it suffices to show that whenever $g_1 a' g_2 = g'_1 a' g'_2$ (where $g_1, g'_1, g_2, g'_2 \in \langle a'^2, b'\rangle$) then $a'^{g_2} = a'^{g'_2}$. By the rigidity of reduced words in $\langle a', b' \mid a'^{2^r}\rangle$ observe that this equality only happens if $g'_1 = g_1 a'^{-2r}$ and $g'_2 = a'^{2r} g_2$ for some $r \in \mathbb{Z}$. Thus

$$a'^{g'_2} = a'^{\left(a'^{2r} g_2\right)} = \left(a'^{a'^{2r}}\right)^{g_2} = a'^{g_2}$$

As required.

Case 2 The intersection of $\langle a', b'a'b'^2\rangle$ and $\langle a', b'a'b'^2\rangle^g$ for $g \in \langle a', b'\rangle \setminus \langle a', b'a'b'^2\rangle$.

Let $\tilde{w} = \alpha_1\alpha_2\cdots\alpha_n$ and $\tilde{w}^h$ be cyclicly reduced words in $\{a'^{\pm 1}, (b'a'b'^2)^{\pm 1}\}$ for some $h \in \langle a', b'\rangle$. Observe by cycling letters in $\{a'^{\pm 1}, b'^{\pm 1}\}$ that we must have $h \in \langle a', b'a'b'^2\rangle$. It follows that every element in the intersection is trivial unless $g \in \langle a', b'a'b'^2\rangle$.

Case 3 The intersection of $\langle a'^2, b'\rangle$ and $\langle a', b'a'b'^2\rangle^g$ for $g \in \langle a', b'\rangle$.

Let $\tilde{w}_1 = \alpha_1\alpha_2\cdots\alpha_n$ be a cyclicly reduced word in $\{a'^{\pm 2}, b'^{\pm 1}\}$and let $\tilde{w}_2 = \beta_1\beta_2\cdots\beta_m$ be a cyclicly reduced word in $\{a'^{\pm 1}, (b'a'b'^2)^{\pm 1}\}$. Suppose that $\tilde{w}_1$ and $\tilde{w}_2$ are conjugate to each other. Observe that $(b'a'b'^2)^{\pm}$ can't be a subword of any cyclic permutation (in $\{a'^{\pm 1}, b'^{\pm 1}\}$) of $\tilde{w}_1$ and so $\beta_j \neq (b'a'b'^2)^{\pm}$ for any j. Hence $\tilde{w}_1$ and $\tilde{w}_2$ are (even) powers of a'. Essentially the same argument also shows that $\langle a'^2, b'\rangle \cap \langle a', b'a'b'^2\rangle = \langle a'^2\rangle$.

The previous paragraph implies that if $w \in \langle a'^2, b' \rangle \cap \langle a', b'a'b'^2 \rangle^g$ then we must have $w = (a'^{2n})^h$ for some $h \in \langle a', b' \rangle$. Now arguments from case 1 imply that either h or $a'h$ must be in $\langle a'^2, b' \rangle$ as $w \in \langle a'^2, b' \rangle$. Likewise arguments from case 2 implies that $k^{-1} := hg^{-1} \in \langle a', b'a'b'^2 \rangle$ as $w^{g^{-1}} \in \langle a'^2, b' \rangle$. So $g = kh$ and without loss of generality we have $h \in \langle a'^2, b' \rangle$; as if $a'h \in \langle a'^2, b' \rangle$ we can just replace h with $a'^{-1}h$ and k with ka'.

We wish to show that if $kh = k'h'$ (where $k, k' \in \langle a', b'a'b'^2 \rangle$ and $h, h' \in \langle a'^2, b' \rangle$) then $a'^h = a'^{h'}$. Once we show this it follows that the intersection is cyclic and hence finite as a'^2 has finite order. Since $\langle a'^2, b' \rangle \cap \langle a', b'a'b'^2 \rangle = \langle a'^2 \rangle$ and by the rigidity of reduced words in $\langle a', b' \,|\, a'^{2^r} \rangle$ we see that this only happens if $k' = ka'^{-2r}$ and $h' = a'^{2r}h$ for some $r \in \mathbb{Z}$. Once again we get $a'^h = a'^{h'}$ in the same way as in case 1.

□

Note that a hyperbolic group G cannot satisfy Theorem 2.4.5. This is because there are only finitely many conjugacy classes of finite subgroups of a hyperbolic group [6]; thus there is some bound on the order of finite subgroups. We can then apply (k, C)-acylindrical accessibility (Theorem 2.4.2) to get a bound here. One may then wonder if Conjecture 2.4.4 holds for hyperbolic groups; however a slight tweak to our example shows that this isn't true either, even for free groups.

Theorem 2.4.6. *For any $N > 0$ there is an action of F_2 on a reduced tree which is 1–acylindrical on non-cyclic subgroups and has N orbits of edges.*

Proof. The construction is mostly the same as Theorem 2.4.5 and so we will only detail the changes. This time we define $A := \langle a \rangle \cong \mathbb{Z}$ so that $G = \langle a, b \rangle \cong F_2$. Pick any $N > 0$ and define $a_i = a^{2^i}$. We now define the tree and see that it satisfies the necessary conditions in the same way as before. □

Chapter 3

Partially acylindrical actions

In the constructions in Theorem 2.4.5 and Theorem 2.4.6 we exploit chains of subgroups with arbitrary length. More precisely we have the chain of subgroups $\langle a_0 \rangle > \langle a_1 \rangle > \cdots > \langle a_N \rangle$ and build the tree in such a way that each group in this chain fixes a vertex which isn't fixed by any of the larger ones. In this chapter we extend Theorem 2.4.2 from accessible except on subgroups of bounded size to accessible except on a set of subgroups without long chains.

3.1 Statement of main theorems

Before stating the results of this section we need a series of basic definitions.

Definition 3.1.1. *Let G be a group and k be a non-negative integer. Let Q be a class of subgroups of G which is closed under conjugation. An action of G on a tree T is* (partially) k–acylindrical on (or over) Q *if whenever some $H \in Q$ fixes every edge in a reduced edge path p then p contains at most k edges.*

Remark 3.1.2. *If Q isn't specified then we will assume it contains all the non-trivial subgroups of G; matching our previous definition of a k–acylindrical action. Observe that an action is (k, C)–acylindrical if it's k–acylindrical on the subgroups of G with size strictly greater than C.*

Observe that an action being (partially) k–acylindrical prevents the unlimited subdivision of edges which motived the definition of a reduced action. As such we can replace it with the following weaker notion.

Definition 3.1.3. *Suppose $\mathcal{P}$ is a class of subgroups of G which is closed under conjugation. A minimal action is said to be* partially-reduced *over $\mathcal{P}$ if either*

- *T/G is a circle consisting of a single vertex and edge; or*
- *whenever a vertex v of T has stabiliser equal to that of an edge and which is contained in a subgroup of a member of $\mathcal{P}$ then v/G has valence at least 3 in T/G.*

Remark 3.1.4. *Observe that an action is reduced if it's partially-reduced over the class of all subgroups of G. We say an action is C–*partially-reduced *if it's partially-reduced over the class of all subgroups of size at most C.*

The following gives us a way of measuring how "large" a given group is relative to $\mathcal{P}$.

Definition 3.1.5. *Let $\mathcal{P}$ be a conjugation invariant set of subgroups of G. For a subgroup $K \leqslant G$ suppose there is some maximal integer n such that there are $H_1, \cdots, H_n \in \mathcal{P}$ with*

$$K \leqslant H_1 < H_2 < \cdots < H_n$$

*We define the $\mathcal{P}$–*weight *of K to be 2^n and we denote this quantity by $W_{\mathcal{P},K}$. If chains of arbitrary length exist we say the $\mathcal{P}$–weight is equal to ∞. Finally if K is not a subgroup of a member of $\mathcal{P}$ we say that $K \leqslant G$ is* larger than $\mathcal{P}$ *and define $W_{\mathcal{P},K} = 1$.*

If G acts on T then we define the $\mathcal{P}$–weight of each edge to be the $\mathcal{P}$–weight of its stabiliser. We say that $\mathcal{P}$ has height *M if the maximal weight of any $K \leqslant G$ is 2^M; equivalently if $W_{\mathcal{P},1} = 2^M$.*

Remark 3.1.6. *Since we insist that $\mathcal{P}$ is conjugation invariant we see that the $\mathcal{P}$-weight of a subgroup is conjugation invariant. As such we define the $\mathcal{P}$-weight of a conjugacy class of subgroups to be the $\mathcal{P}$-weight of any representative of that class.*

We now state an easier version of our main results. We will prove this before moving on to the full theorems as it will demonstrate the important ideas of the argument without being obscured by as many technical details.

Theorem 3.1.7. *Let G be a finitely presented group and let $\mathcal{P}$ be a set of subgroups for G with height M and which is closed under conjugation and taking subgroups. Suppose G acts on a tree T and that this action is both k–acylindrical on groups larger than $\mathcal{P}$ and partially reduced on groups in $\mathcal{P}$. Then there is some $C(G)$ (which depends only on G) such the that number of edges of T/G is bounded above by $(2k+1)2^M C(G)$.*

Our main results are two different generalisations of the above. In the first we extend the result to certain cases where $\mathcal{P}$ isn't closed under taking subgroups, which is necessary for including infinite subgroups in $\mathcal{P}$. In the second we extend the result to groups which are merely finitely generated instead of just those which are finitely presented. In order to state the former we first need to make the following definitions.

Definition 3.1.8. *Let $\mathcal{P}$ be a set of subgroups for G. Let K be a subgroup of G. We say $H \in \mathcal{P}$ is a* minimal extension of K (to $\mathcal{P}$) *if $K \leqslant H$ and whenever $\tilde{H} \in \mathcal{P}$ with $K \leqslant \tilde{H} \leqslant H$ then $\tilde{H} = H$. We say that K is $\mathcal{P}$*–closed *if any subgroup of K has a minimal extension which is contained in K. We say an action of G on a tree T is $\mathcal{P}$*–closed *if all its edge stabilisers are $\mathcal{P}$–closed.*

Remark 3.1.9. *If $\mathcal{P}$ has finite height then minimal extensions always exist for any group which isn't larger than $\mathcal{P}$.*

Definition 3.1.10. *We say that $\mathcal{P}$ satisfies condition $(\dagger)$ if the following conditions hold.*

- *$\mathcal{P}$ has finite height.*
- *Suppose G acts on a tree T and let e be an edge of T. Then for any subgroup $K \leqslant \operatorname{Stab} e$ which is not larger than $\mathcal{P}$ there is a vertex v of T which is fixed by some minimal extension of K to $\mathcal{P}$. In particular this always holds if each minimal extension is of finite index as a finite index extension of an elliptically acting group also fixes a point [21].*
- *If H_1 and H_2 are in $\mathcal{P}$ then so is $H_1 \cap H_2$. An equivalent condition is that minimal extensions to $\mathcal{P}$ are unique.*

Example 3.1.11. *Suppose that G is a torsion-free hyperbolic group. Take $\mathcal{P}$ to be the set of cyclic subgroups of G which are root-closed. Equivalently $\mathcal{P}$ is the set of maximal cyclic subgroups of G. This $\mathcal{P}$ satisfies $(\dagger)$ since every cyclic subgroup is contained (with finite index) in a unique maximal cyclic subgroup. Moreover a tree which G acts on is $\mathcal{P}$–closed if and only if all its edge stabilisers are root-closed.*

We are now ready to state our main result, which as mentioned before is a pair of extensions of Theorem 3.1.7.

Theorem 3.1.12. *Let G be a finitely generated group and let $\mathcal{P}$ be a set of subgroups for G which is closed under conjugation and has height M. Suppose G acts on a tree T and that this action is both k–acylindrical on groups larger than $\mathcal{P}$ and partially reduced on $\mathcal{P}$. Then the following statements hold.*

(a) Suppose G is finitely presented, $\mathcal{P}$ satisfies (†) *and T is $\mathcal{P}$–closed. Then there is some integer $C(G)$ such that the number of edges of T/G is bounded above by $(2k+1)2^M C(G)$.*

(b) If $\mathcal{P}$ is closed under taking subgroups then the number of edges of T/G is bounded above by $\binom{2k+1}{2} 2^M(\operatorname{rank} G - 1)$. Moreover suppose either of the following conditions hold

- *the action on T is reduced and $k > 1$; or*
- *every edge stabiliser of T is not in $\mathcal{P}$;*

then if G isn't cyclic the number of edges of T/G is bounded above by $2^M k(\operatorname{rank} G - 1)$.

Remark 3.1.13. *In Theorem* 3.1.7 *and Theorem* 3.1.12 *(a) the bound $C(G)$ is given by Dunwoody's resolution lemma (Theorem* 2.1.12*). Dunwoody's resolution lemma holds for so called* almost finitely presented groups *and this extends to both Theorem* 3.1.7 *and Theorem* 3.1.12 *(a). (A group is* almost finitely presented *if it's both finitely generated and acts freely, simplicially and cocompactly on a simplicial complex X with $H^1(X, \mathbb{Z}_2) = 0$.)*

The following is an immediate consequence of Theorem 3.1.12 (b) and is an extension of Weidmann's result on (k, C)–acylindrical actions [25]. In particular this shows that the number of prime factors is the limiting factor, not the absolute size of the group.

Corollary 3.1.14. *Let G be a finitely generated group and $M \in \mathbb{N}$. Suppose G acts on a tree T and that this action is both k–acylindrical on groups which are infinite or have at least M prime factors and partially reduced on subgroups with at most $M - 1$ prime factors. (Where the number of prime factors is counted without multiplicity.) Then the number of edges of T/G is bounded above by $\binom{2k+1}{2} 2^M(\operatorname{rank} G - 1)$. Moreover if either T is reduced and $k > 1$ or every edge stabiliser of T is either infinite or has at least M prime factors then the number of edges of T/G is bounded above by $2^M k(\operatorname{rank} G - 1)$.*

Proof. Let $\mathcal{P}$ be the set of finite subgroups of G whose order has at most $M - 1$ prime factors. Observe that $\mathcal{P}$ has height of at most M and is closed under taking subgroups. The result now immediately follows from Theorem 3.1.12 (b). □

We also apply Theorem 3.1.12 (a) to a couple of specific cases to get some interesting results. The first case is a generalisation of Example 3.1.11; where a torsion-free hyperbolic group acts on a tree with root-closed edge stabilisers. We now allow the group to have finite order elements and the root closed condition is replaced by one which says that the maximal virtually $\mathbb{Z}$ subgroups of edge stabilisers should be "almost" maximal in G. In the second we consider splittings of RAAGs which are k–acylindrical on its non-abelian subgroups and the maximal abelian subgroups of edge stabilisers should be maximal for their rank.

Definition 3.1.15. *Let G be a hyperbolic group. We say a virtually $\mathbb{Z}$ subgroup $H \leqslant G$ is m-* almost maximal *if whenever we have a virtually cyclic $K \leqslant G$ with $H \leqslant K$ then $[H : K] \leq m$. Let $\mathcal{P}_m$ be the collection of subgroups of G which are either finite or m-almost maximal.*

Corollary 3.1.16. *Suppose G acts on a tree T with an action which is partially reduced on $\mathcal{P}_m$ and k–acylindrical on groups larger than $\mathcal{P}_m$. Suppose also that the action on T is $\mathcal{P}_m$-closed. Then the number of edges of T/G is bounded above by $(2k+1)C'(G)$. (Where $C'(G) = 2^n C(G)$ for some $n \in \mathbb{N}$.)*

Proof. Observe that $\mathcal{P}_m$ is not closed under intersections and so doesn't satisfy condition $(\dagger)$. Instead define $\mathcal{P}'_m$ to be the set of subgroups which are a (finite) intersection of groups in $\mathcal{P}_m$. It's clear that if $\mathcal{P}'_m$ has finite height then it satisfies $(\dagger)$. As a hyperbolic group has only finitely many conjugacy classes of finite subgroups [4] we just need to show that chains of infinite subgroups in $\mathcal{P}'_m$ have bounded length. It therefore suffices to show that given any index in $\mathbb{N}$ there is a uniform bound on the number of subgroups of that index for any virtually cyclic subgroup of G.

Recall that every virtually cyclic $H \leqslant G$ is either of the form $H \cong K*_K$ where K is finite or $H \cong A *_C B$ where A, B, C are finite and C is an index 2 subgroup of both A and B [17]. The general description of a subgroup of the fundamental group of a graph of groups [21] now tells us there's a bound on the number of subgroups of H of a given index which depends only on the index and the size of either K or C respectively. Again a hyperbolic group has finitely many conjugacy classes of finite subgroups, which uniformly bounds the order of K and C. This implies the result. □

Before stating the other application we briefly recall the definition of a RAAG.

Definition 3.1.17. *Let Γ be a finite graph. Let $v_1, \cdots, v_n$ be the vertices of Γ. The* right-angled Artin group *(RAAG) associated to Γ is the group $A(\Gamma)$ where*

$$A(\Gamma) := \left\langle v_1, \cdots, v_n \mid [v_i, v_j] \text{ wherever there's an edge between } v_i \text{ and } v_j \text{ in } \Gamma. \right\rangle$$

Corollary 3.1.18. *Let $G = A(\Gamma)$ be a RAAG. An abelian subgroup $H \leqslant G$ is said to be* rank maximal *if the following condition holds. Whenever an abelian subgroup $K \leqslant G$ which contains H with finite index then we have $K = H$. Let $\mathcal{P}$ be the collection of rank maximal abelian subgroups of G. Suppose G acts on a tree T which is partially reduced on abelian* *subgroups and k–acylindrical on non-abelian subgroups. Suppose also that T is $\mathcal{P}$-closed. Then the number of edges of T/G is bounded above by $(2k+1)2^nC(G)$ (where n is the size of the largest complete subgraph of Γ).*

Proof. Recall that a RAAG acts freely and cocompactly on a simply connected CAT(0)-cube complex X_Γ whose dimension is equal to the size of the largest complete subgraph of Γ. The flat torus theorem [7, Theorem II.7.1] says that an abelian subgroup $H \leqslant A(\Gamma)$ must act properly and cocompactly by isometries on a Euclidean hyperplane of X_Γ. In particular H must have rank at most equal to the size of the largest complete subgraph of Γ. Hence $\mathcal{P}$ has finite height.

Now every rank 2 subgroup $\langle u, v\rangle \leqslant A(\Gamma)$ is either free abelian or free [1, Theorem 1.2]. Suppose $H \leqslant A(\Gamma)$ is an abelian subgroup, u is a root of an element of H and pick any $v \in H$. Since a power of u is in H and H is abelian we see that $\langle u, v\rangle$ cannot be non-abelian free and so u and v must commute. Hence every member of $\mathcal{P}$ is root closed.

It remains to check that $\mathcal{P}$ satisfies condition $(\dagger)$, then we can apply Theorem 3.1.12 (a) to get the result. Pick any $H \in \mathcal{P}$ and $K \leqslant H$ and let $M \in \mathcal{P}$ be a minimal extension of K. We must have $M \leqslant H$ as H is root closed and K is a finite index subgroup of M. Hence $\mathcal{P}$ is $\mathcal{P}$-closed and hence satisfies $(\dagger)$. □

Remark 3.1.19. *For a general group G it need not be the case that $\mathcal{P}$ as defined in Corollary* 3.1.18 *satisfies* $(\dagger)$. *For example if $G \cong \mathbb{Z} *_{2\mathbb{Z}} \mathbb{Z}$ then G acts freely and cocompactly on a* CAT(0) *space; but contains two rank maximal copies of $\mathbb{Z}$ whose intersection is another copy of $\mathbb{Z}$ which is not rank maximal.*

3.2 Forests of influence

We'll now give an *extremely* rough outline of the core ideas of the argument. Suppose that G acts on a minimal tree T which is k–acylindrical on groups larger than $\mathcal{P}$. Use Dunwoody's resolution lemma to obtain a tree T' which has a bound on the number of edges and a map $\Psi : T' \to T$. If some edge of T' (before subdividing) has a stabiliser larger than $\mathcal{P}$ then its image in T cannot have more than k edges because of the acylindrical condition. Thus we can collapse this edge in T' and only collapse at most k edges of T.

Now subdivide T' to make Ψ simplicial, but note that the initial vertices are 'more important' in the sense that every vertex stabiliser is contained in one of these. So we can build a collection of disjoint subtrees for T' by starting with this set of initial vertices and then iteratively expanding to include vertices whose stabiliser is contained in the stabiliser of the corresponding initial vertex.

Now we subdivide Ψ into folds using Stallings' folding theorem (Theorem 2.2.2) and apply the first fold. If every vertex stabilizer is still contained in a stabilizer for one of the initial vertices then we have still have a collection of subtrees with the same properties as before. Otherwise some vertex stabiliser isn't contained in one of the initial ones. This only happens if two of our subtrees gets folded together in some way which is unavoidable. We then add this vertex to our set of "initial" ones and then rebuild our collection of subtrees with the same properties as before. However we will see that the intersections of the stabilisers between one of the original initial vertices and this "new initial vertex" is strictly larger than the intersection of the original initial vertices. (See Figure 3.1 for an example or Lemma 3.2.14 for a more precise statement.) If $\mathcal{P}$ has finite height this means that this can only happen boundedly often before one of these intersections is larger than $\mathcal{P}$ and so can collapse down a path of length at most k. So either we can keep doing this until we are left with a single point or we get a set of "initial" vertices for T. In the latter case if T is $\mathcal{P}$–partially reduced we can find a bound for the number of edges using our set "initial" vertices. (See Lemma 3.3.1.)

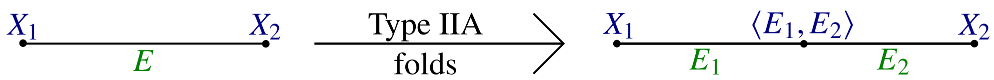

Fig. 3.1 After applying a series of type II folds to an edge (with subdivisions) there may be a vertex whose stabiliser isn't contained in the stabiliser of either of the initial vertices. If this happens we see that the intersection of stabilisers between this vertex and either of the initial vertices must contain the original edge group as a *proper* subgroup.

In order to make the above precise we introduce the following notions.

Definition 3.2.1. *Suppose G acts on a tree T. We call a subset of vertices S a set of* seed vertices *if it's G–invariant and for every vertex v (with non-trivial stabiliser) there is some $u \in S$ with* $\operatorname{Stab} v \leqslant \operatorname{Stab} u$. *In particular if the action on T is free we also allow the empty set to be a set of seed vertices, otherwise S is necessarily non-empty.*

Definition 3.2.2. *Suppose G acts on a tree T. A G–invariant subgraph $\Gamma \subseteq T$ is a* forest of influence *if the following conditions hold.*

- Γ *deformation retracts to a non-empty set of seed vertices S, equivalently every component of Γ contains exactly one member of S. We say that Γ is* grown *from S.*
- *If vertices u and v are in the same connected component of Γ with $u \in S$ then* $\operatorname{Stab} v \leqslant \operatorname{Stab} u$. *We call such a component the* tree of influence *of u, say that v is* influenced *by u and call the reduced edge path from v to u the* branch *of v.*
- *Every vertex of T is contained in Γ.*

Remark 3.2.3. *The branch of any vertex v is stabilised by* $\operatorname{Stab} v$. *As such the first edge on the branch of v must have the same stabiliser as v as any edge cannot be fixed by more than either of its endpoints.*

Definition 3.2.4. *Suppose G acts on a tree T and that $\Gamma \subseteq T$ is a forest of influence. We call the edges of $T \setminus \Gamma$ the* connecting edges *of Γ. Fix a representative set of connecting edges. The* connecting groups *are the conjugacy classes of the stabilisers of this (any) representative set of connecting edges, counted with multiplicity.*

Remark 3.2.5. *Observe that a connecting group is a conjugacy class and not a group. However we will often abuse this notation by saying that a representative of a connecting group (i.e. the stabiliser of a connecting edge)* is *the connecting group. The reason for this strange definition and abuse of notation is that we want a finite collection of objects which are uniquely determined, but also find it easier to work with groups directly.*

The following simple observation allows to better visualise forests of influence.

Proposition 3.2.6. *Suppose G acts on a tree T and that $\Gamma \subseteq T$ is a forest of influence grown from the set of seed vertices S. Let v be a vertex of T and suppose it's influenced by $u \in S$ and let $g \in G$. Then v and gv are in the same tree of influence if and only if $g \in$* $\operatorname{Stab} u$.

Proof. If $g \in \operatorname{Stab} u$ then the reduced edge paths $[u, v]$ and $[u, gv]$ are both contained in Γ; so v and gv are in the same tree of influence. Now suppose $g \notin \operatorname{Stab} u$ and that v and gv are in the same tree of influence. Observe that u and gu are distinct seed vertices of S and that Γ contains the reduced edge paths $[u, v]$, $[gu, gv]$ and $[v, gv]$. In particular u and gu are in the same component of Γ; contradicting the fact that each component of Γ contains exactly one member of S. Hence if $g \notin \operatorname{Stab} u$ then v and gv are in different trees of influence. □

Corollary 3.2.7. *Suppose G acts on a tree T and that $\Gamma \subseteq T$ is a forest of influence. Then Γ/G is a forest contained in T/G.*

Proof. If Γ/G contains a non-trivial loop then this corresponds to a path in Γ between vertices v and hv where h acts hyperbolically on T. But this is impossible because of Proposition 3.2.6. □

In general there is not a distinguished choice for a forest of influence. However the following proposition says there is something canonical lurking underneath. This will allow us to move between different choices with minimal complications.

Proposition 3.2.8. *Suppose that T has finitely many orbits of vertices. Suppose also that Γ_1 and Γ_2 are forests of influence which are both grown from the same set of seed vertices S. Then Γ_1 and Γ_2 have the same connecting groups. In other words the connecting groups are determined by S.*

Before proving this we'll first we'll define an *elementary transformation* of a forest of influence. Take a forest of influence Γ and pick a vertex $v \in \Gamma \setminus S$. Suppose that v is contained in the tree of influence of u and let e_1 be the first edge on the branch of v. Observe that $\operatorname{Stab} e_1 = \operatorname{Stab} v$ and let e_2 be a connecting edge with endpoint v and with $\operatorname{Stab} e_2 = \operatorname{Stab} v$. (If no such e_2 exists then we cannot apply an elementary transformation at the vertex v.) We now define $\Gamma' := (\Gamma \setminus G\{e_1\}) \cup G\{e_2\}$. In other words we replace the orbit of e_1 in Γ with the orbit of e_2 in Γ'. (See Figure 3.2.) Since $\operatorname{Stab} e_1 = \operatorname{Stab} e_2 = \operatorname{Stab} v$ we see that Γ' is also a forest of influence grown from S and that both Γ and Γ' have the same connecting groups.

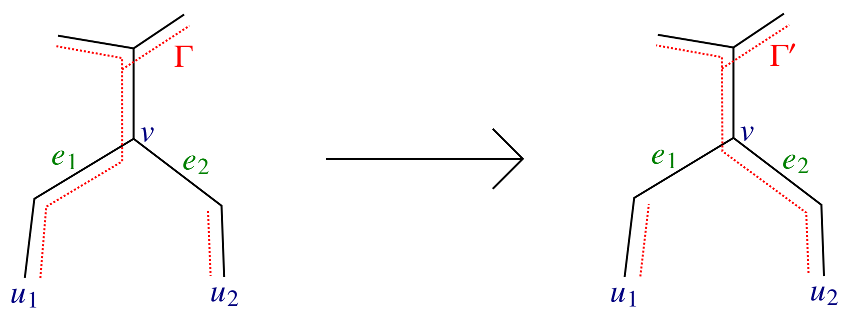

Fig. 3.2 An example of an elementary transformation. The edge e_1 is removed and replaced with e_2. For Γ' to be a forest of influence we must have $\operatorname{Stab} e_1 = \operatorname{Stab} e_2 = \operatorname{Stab} v$.

Proposition 3.2.8 is now an immediate consequence of the following.

Lemma 3.2.9. *Suppose that T has finitely many orbits of vertices and that Γ_1 and Γ_2 are forests of influence which are both grown from the same set of seed vertices S. Then we can apply a finite series of elementary transformations to Γ_1 to obtain Γ_2.*

Proof. Let $d(\Gamma_1,\Gamma_2)$ be the number of (orbits of) vertices which are in trees of influence of different seed vertices in Γ_1 and Γ_2. If $d(\Gamma_1,\Gamma_2) = 0$ then $\Gamma_1 = \Gamma_2$ and there is nothing to show.

If $d(\Gamma_1,\Gamma_2) > 0$ pick a vertex v which is in the tree of influence of u_1 in Γ_1 and of $u_2 \neq u_1$ in Γ_2. Let e_2' be the final edge in the branch of v (in Γ_2) which is not contained in Γ_1 and so is a connecting edge of Γ_1. Let v' be the endpoint of e_2' which is not in the tree of influence of u_2 in Γ_1. Observe that v' is in the tree of influence of u_2 in Γ_2 as v is. Suppose that v' is in the tree of influence of u_1' in Γ_1 and let e_1' be the first edge on the branch of v' (in Γ_1). Since $\operatorname{Stab} v' = \operatorname{Stab} e_1' = \operatorname{Stab} e_2'$ we can apply an elementary transformation to Γ_1 by removing the orbit of e_1' and adding the orbit of e_2' to get Γ_1'. (See Figure 3.3.)

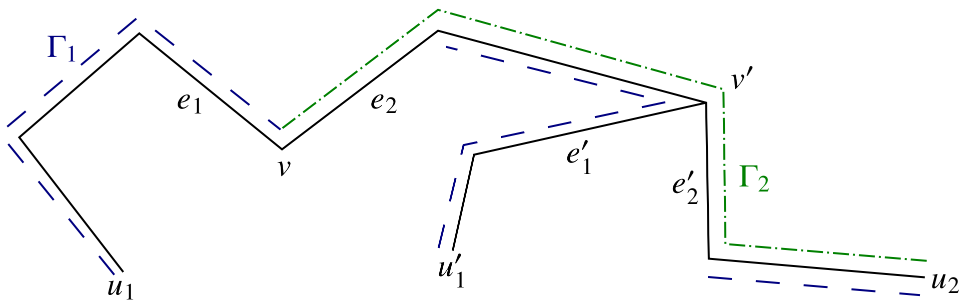

Fig. 3.3 An example of a situation where $d(\Gamma_1,\Gamma_2) > 0$. By replacing e_1' with e_2' in Γ_1 we can make it "more similar" to Γ_2. This idea of how similar two forests of influence are is formalised by the metric d.

If we can show that $d(\Gamma_1',\Gamma_2) < d(\Gamma_1,\Gamma_2)$ then we are done by induction. Observe that v' is in the tree of influence of u_2 in Γ_1' and Γ_2 but not in Γ_1. Thus we just need to show that any vertex which is influenced by the same seed vertex in Γ_1 and Γ_2 is also influenced by the same one in Γ_1'. This holds because the only vertices whose influencing vertex changed under the elementary transformation were those in (the orbit of) the tree of influence of u_1' in Γ_1 at and beyond v'. These can't be influenced by u_1' in Γ_2 as v' is influenced by u_2 in Γ_2 and so the tree of influence of u_1' in Γ_2 cannot contain them. □

Recall the definition of the $\mathcal{P}$–weight of a subgroup $K \leqslant G$ from Definition 3.1.5 as $W_{\mathcal{P},K} \leq 2^N$ where N is the length of the longest chain of groups in $\mathcal{P}$ which contain K. From this definition we note that the following properties are all obvious.

Proposition 3.2.10. *Let $\mathcal{P}$ be a conjugation invariant set of subgroups of G.*

(a) If $\mathcal{P}$ has height M then $W_{\mathcal{P},K} \leq 2^M$ for any $K \leqslant G$.

(b) *$K \leqslant G$ has $\mathcal{P}$–weight* 1 *if and only if it's larger than* $\mathcal{P}$.

(c) *If $H \in \mathcal{P}$ and $K < H \leq G$ then $W_{\mathcal{P},H} \leq \frac{1}{2} W_{\mathcal{P},K}$.*

We will now extend our definition of $\mathcal{P}$–weight to sets of seed vertices. Proposition 3.2.8 ensures this is well defined.

Definition 3.2.11. *If G acts on a tree T and S is a non-empty set of seed vertices for T then we define its $\mathcal{P}$–*weight *$W_{\mathcal{P},S}$ to be the sum of the $\mathcal{P}$–weights of the corresponding connecting groups (and ∞ if any of the connecting groups have infinite $\mathcal{P}$–weight). If S is empty then we instead define $W_{\mathcal{P},S} := (\beta_1(T/G) - 1)W_{\mathcal{P},1}$.*

Remark 3.2.12. *The case of a free action is special because the stabiliser of each vertex is trivial. As there are no "interesting" stabilisers we aren't really missing anything by just forgoing seed vertices entirely. If the action is free and S is non empty then we see that*

$$W_{\mathcal{P},S} := (\beta_1(T/G) + |S/G| - 1)W_{\mathcal{P},1}.$$

This justifies the definition of $W_{\mathcal{P},S}$ for empty S by setting $|S/G| = 0$ in the above equation. On a more practical level we allow the empty set to be a set of seed vertices for a free action to prevent an otherwise guaranteed drop in $\mathcal{P}$–weight if a fold causes a free action to become non-free. (See Lemma 3.2.15.*)*

With this in hand we are ready to state the key lemma. From this Theorem 3.1.7 will follow quickly.

Lemma 3.2.13. *Suppose G is a non-cyclic countable group. Let $\mathcal{P}$ be a conjugation invariant set of subgroups of G which is closed under taking subgroups. Let G act on a tree T where this action is both $\mathcal{P}$–partially-reduced and k–acylindrical on a subgroups larger than $\mathcal{P}$. Let G act on another tree T' and suppose that there is a G-equivariant combinatorial map $\Psi : T' \to T$. Suppose also that T' has a set of seed vertices S' with finite $\mathcal{P}$–weight $W_{\mathcal{P},S'}$. Then T/G has at most $\left(\frac{2k+1}{2}\right) W_{\mathcal{P},S'}$ edges.*

The remainder of this section as well as the entirety of Section 3.3 will be dedicated to providing the necessary tools to prove this.

Recall that our plan involves decomposing Ψ into folds. The following says that we can recursively find a nice set of seed vertices for each intermediate step.

Lemma 3.2.14. *Suppose that $\alpha : R \to \tilde{R}$ is a folding map between trees acted on by some group G. Suppose that S is a non-empty set of seed vertices for R where all of the connecting* *groups are in $\mathcal{P}$. Then there is a set of seed vertices $\tilde{S}$ for $\tilde{R}$ with $\alpha(S) \subseteq \tilde{S}$ and $W_{\mathcal{P},\tilde{S}} \leq W_{\mathcal{P},S}$. Moreover if $W_{\mathcal{P},\tilde{S}} = W_{\mathcal{P},S}$ then $\alpha|_S$ is injective.*

Proof. Suppose that α folds together the edges $e_1 = [x, y_1]$ and $e_2 = [x, y_2]$. Suppose $\alpha(e_1) = \alpha(e_2) = e'$ and $\alpha(y_1) = \alpha(y_2) = y'$. Let y_i be in the tree of influence of u_i and if $y_i \neq u_i$ we also let f_i be the first edge in the branch of y_i. We will split into cases depending on if there is a forest of influence containing e_1 and/or e_2. In the first two cases, (as well as for $i = 2$ in the third and final case), we will assume that $y_i \neq u_i$ and so f_i exists. The cases where $y_i = u_i$ turn out to be essentially the same except the lack of f_i sometimes causes $W_{\mathcal{P},\tilde{S}}$ to be smaller.

Case 1 There is a forest of influence Γ containing both e_1 and e_2.

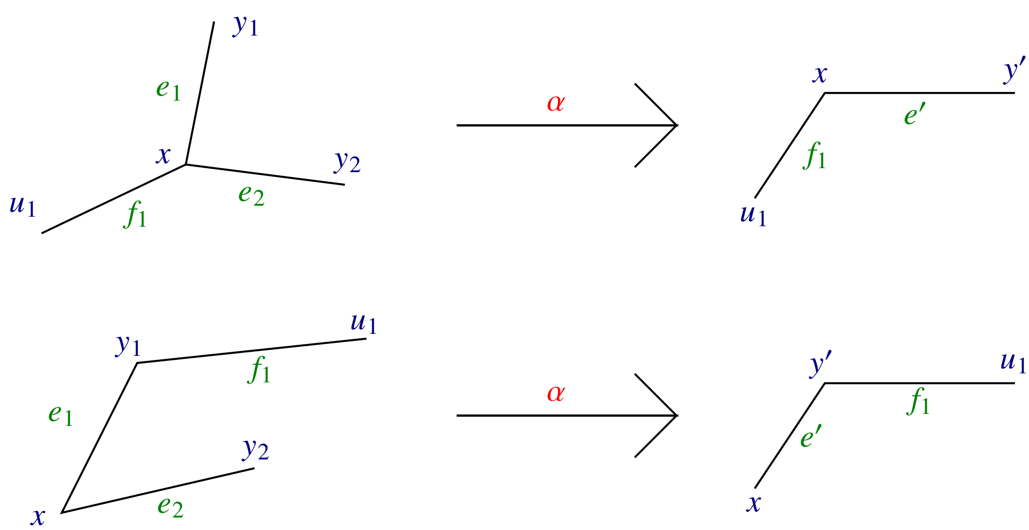

Fig. 3.4 The two pictures that can arise when Γ contains both e_1 and e_2. In either case Stab u_1 contains Stab x, Stab y_1 and Stab y_2 and hence also contains Stab y'.

The fold cannot be of type III as otherwise y_1 and $y_2 = hy_1$ would be in the same tree of influence, contradicting Proposition 3.2.6. So $\alpha(\Gamma)$ is a forest of influence for $\tilde{R}$ which is grown from $\alpha(S)$. The connecting edges of Γ are untouched by α and so S and $\alpha(S)$ have the same connecting groups.

Case 2 There is no forest of influence containing either e_1 or e_2.

Pick any forest of influence Γ. Observe that $\tilde{S} = \alpha(S) \cup G\{y'\}$ is a set of seed vertices which grows into $\alpha(\Gamma) \setminus G\{f_1, f_2\}$. If Stab e_i = Stab f_i then we could apply an

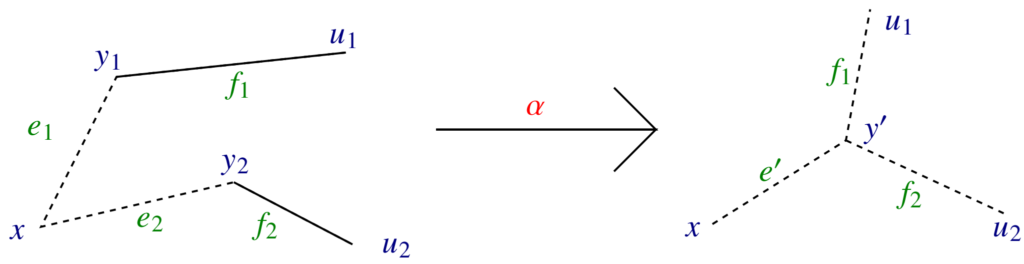

Fig. 3.5 The picture that arises when Γ cannot possibly contain either e_1 or e_2. Observe that $\operatorname{Stab} e_i < \operatorname{Stab} u_i$. After applying the fold we remove the edges f_i (if they exist) from the forest of influence. Separate arguments depending on the type of fold are now required to show that this doesn't increase the $\mathcal{P}$-weight. (In this diagram and the ones which follow the dashed lines indicate a connecting edge.)

elementary transformation to get a new forest of influence which includes e_i instead of f_i. However this contradicts our assumption that e_i is not contained in any forest of influence. Hence $\operatorname{Stab} e_i < \operatorname{Stab} f_i$ and so $W_{f_i,\mathcal{P}} \leq \frac{1}{2} W_{e_i,\mathcal{P}}$ by Proposition 3.2.10 (c).

If α is a fold of type I then e_1, e_2, f_1 and f_2 are pairwise inequivalent. Moreover e' contains the image of both e_1 and e_2, so $W_{e',\mathcal{P}} \leq \min(W_{e_1,\mathcal{P}}, W_{e_2,\mathcal{P}})$ and hence

$$\begin{aligned} W_{S,\mathcal{P}} - W_{\tilde{S},\mathcal{P}} &= W_{e_1,\mathcal{P}} + W_{e_2,\mathcal{P}} - W_{f_1,\mathcal{P}} - W_{f_2,\mathcal{P}} - W_{e',\mathcal{P}} \\ &\geq \frac{1}{2} W_{e_1,\mathcal{P}} + \frac{1}{2} W_{e_2,\mathcal{P}} - \min(W_{e_1,\mathcal{P}}, W_{e_2,\mathcal{P}}) \\ &\geq 0 \end{aligned}$$

If α is a fold of type II then e_1 is equivalent to e_2 and f_1 is equivalent to f_2. Additionally $\operatorname{Stab} e' > \operatorname{Stab} e_1$. So by Proposition 3.2.10 (c) we have

$$\begin{aligned} W_{S,\mathcal{P}} - W_{\tilde{S},\mathcal{P}} &= W_{e_1,\mathcal{P}} - W_{f_1,\mathcal{P}} - W_{e',\mathcal{P}} \\ &\geq W_{e_1,\mathcal{P}} - \frac{1}{2} W_{e_1,\mathcal{P}} - \frac{1}{2} W_{e_1,\mathcal{P}} \\ &= 0 \end{aligned}$$

Now assume α is a fold of type III. We see that f_1 and f_2 are equivalent, while e_1 and e_2 are inequivalent. Thus

$$\begin{aligned}
W_{S,\mathcal{P}} - W_{\tilde{S},\mathcal{P}} &= W_{e_1,\mathcal{P}} + W_{e_2,\mathcal{P}} - W_{f_1,\mathcal{P}} - W_{e',\mathcal{P}} \\
&\geq W_{e_1,\mathcal{P}} - W_{f_1,\mathcal{P}} \\
&\geq W_{e_1,\mathcal{P}} - \frac{1}{2} W_{e_1,\mathcal{P}} \\
&> 0
\end{aligned}$$

Case 3 There is a forest of influence Γ containing e_1 but not e_2; also there isn't one which contains both of them.

Note that α cannot be a fold of type II (as then the e_i are equivalent) or type IIIB (as then both e_i are connecting edges by Corollary 3.2.7). We will split into five subcases; corresponding to combinations of whether or not f_1 is equal to e_1 (if it exists at all) and whether or not $\operatorname{Stab} y_1$ is a subgroup of $\operatorname{Stab} y_2$.

Case 3ai $y_1 \neq u_1$, $e_1 = f_1$ and $\operatorname{Stab} y_1 \leqslant \operatorname{Stab} y_2$.

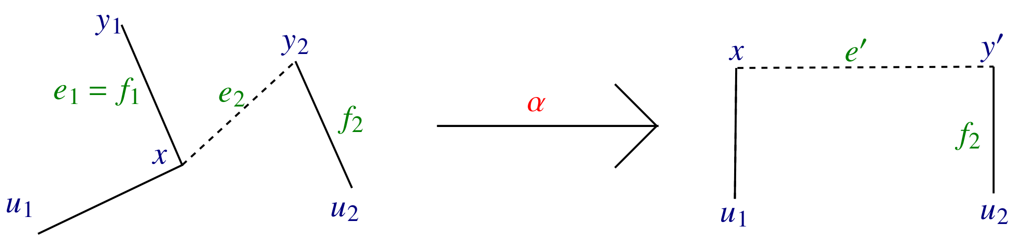

Fig. 3.6 The picture that arises when Γ contains e_1 but not e_2, the branch of y_1 contains x, $\operatorname{Stab} y_2$ contains $\operatorname{Stab} y_1$ and the fold is of type I. Here we take the image of each connecting edge to still be a connecting edge.

If the fold is of type I then observe that $\tilde{S} := \alpha(S)$ is a set of seed vertices for $\tilde{R}$. Observe that the image of the connecting edges of Γ are the connecting edges of a forest of influence grown from $\tilde{S}$; hence the $\mathcal{P}$–weight can't increase. If instead the fold is of type IIIA then $\tilde{S} := \alpha(S) \cup G\{y'\}$ is a set of seed vertices. In this case we have a forest of influence $\tilde{\Gamma} := \alpha(\Gamma) \setminus G\{f_2\}$. Since the fold is of type III the y_i are in the same orbit and hence the f_i are as well. Hence the connecting edges of $\tilde{\Gamma}$ are the image of the connecting edges of Γ with the orbit of e_2 removed and the orbit of $\alpha(f_1)$ added. (This gives the same picture as the one seen in Figure 3.7.) Since $\mathrm{Stab}\, e_2 \leqslant \mathrm{Stab}\, f_2$ we see that the $\mathcal{P}$–weight can't increase in this case either.

Case 3aii $y_1 \neq u_1$, $e_1 = f_1$ and $\mathrm{Stab}\, y_1$ is not contained in $\mathrm{Stab}\, y_2$.

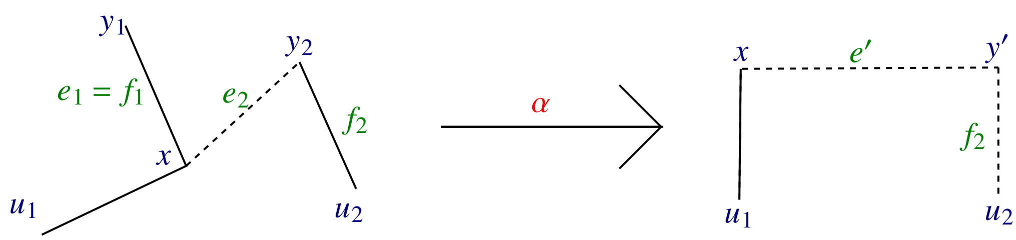

Fig. 3.7 The picture that arises when Γ contains e_1 but not e_2, the branch of y_1 contains x and $\mathrm{Stab}\, y_2$ doesn't contain $\mathrm{Stab}\, y_1$. Here we take the new connecting edges to be the image of the old connecting edges together with f_2 (if it exists).

If the fold is of type IIIA then proceed as in case 3ai. Otherwise observe $\tilde{S} = \alpha(S) \cup G\{y'\}$ is a set of seed vertices for $\tilde{R}$ and that $\tilde{\Gamma} = \alpha(\Gamma) \setminus G\{e', f_2\}$ is a forest of influence grown from $\tilde{S}$. Since $\mathrm{Stab}\, y_1$ is not contained in $\mathrm{Stab}\, y_2$ and $\mathrm{Stab}\, y_1 = \mathrm{Stab}\, e_1$ we have $\mathrm{Stab}\, e_2 < \mathrm{Stab}\, e'$. We also have $\mathrm{Stab}\, e_2 < \mathrm{Stab}\, f_2$ because otherwise we could apply an elementary transformation to Γ to get a new forest of influence which is in case 1. Hence by Proposition 3.2.10 (c)

$$\begin{aligned} W_{S,\mathcal{P}} - W_{\tilde{S},\mathcal{P}} &= W_{e_2,\mathcal{P}} - W_{f_2,\mathcal{P}} - W_{e',\mathcal{P}} \\ &\geq W_{e_2,\mathcal{P}} - \frac{1}{2} W_{e_2,\mathcal{P}} - \frac{1}{2} W_{e_2,\mathcal{P}} \\ &= 0 \end{aligned}$$

Case 3bi $y_1 \neq u_1$, $e_1 \neq f_1$ and $\mathrm{Stab}\, y_1 \leqslant \mathrm{Stab}\, y_2$.

If the fold is of type I then $\tilde{S} := \alpha(S)$ is a set of seed vertices for $\tilde{R}$. If instead the fold is of type IIIA then $\tilde{S} := \alpha(S) \cup G\{y'\}$ is a set of seed vertices instead. In either case observe that $\alpha(\Gamma) \setminus G\{f_1\}$ is a forest of influence grown from $\tilde{S}$. (Note that if the fold is of type IIIA then f_1 and f_2 are in a common orbit, so f_2 also becomes a connecting edge.) Since $\mathrm{Stab}\, e_2 \leqslant \mathrm{Stab}\, x \leqslant \mathrm{Stab}\, f_1$ we have $W_{\tilde{S},C} \leq W_{S,C}$.

Case 3bii $y_1 \neq u_1$, $e_1 \neq f_1$ and $\mathrm{Stab}\, y_1$ is not contained in $\mathrm{Stab}\, y_2$.

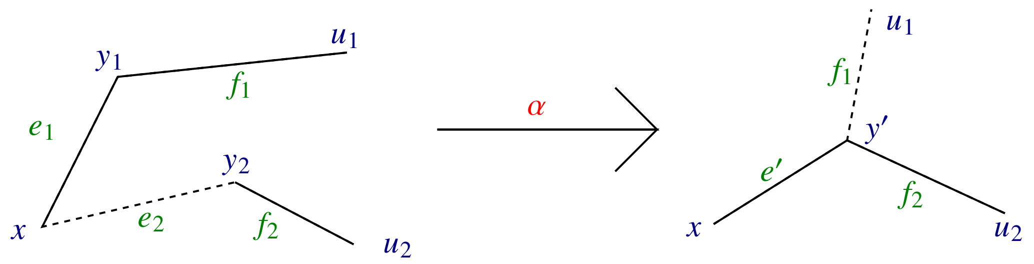

Fig. 3.8 The picture that arises when Γ contains e_1 but not e_2, the branch of y_1 doesn't contain x and $\operatorname{Stab} y_2$ contains $\operatorname{Stab} y_1$. Here we take the new forest of influence to be the image of the old one with f_1 removed.

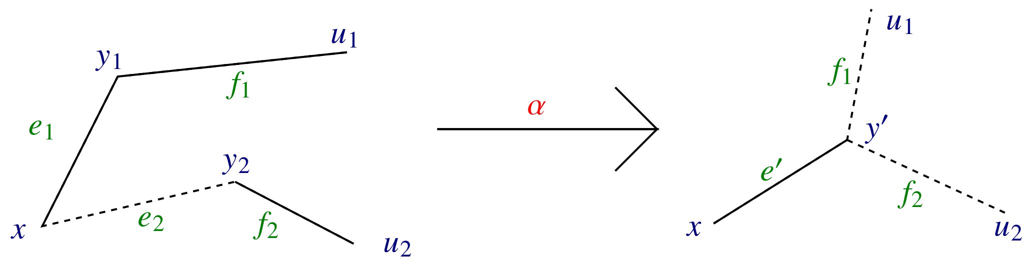

Fig. 3.9 The picture that arises when Γ contains e_1 but not e_2, the branch of y_1 doesn't contain x and $\operatorname{Stab} y_1$ doesn't contain $\operatorname{Stab} y_2$. Here we take the new forest of influence to be the image of the old one with both f_1 and (if it exists) f_2 removed.

If the fold is of type IIIA then proceed as in case 3bi. Otherwise observe that $\tilde{S} = \alpha(S) \cup G\{y'\}$ is a set of seed vertices which grows into a forest of influence $\tilde{\Gamma} = \alpha(\Gamma) \setminus G\{f_1, f_2\}$. If $\operatorname{Stab} e_2 = \operatorname{Stab} f_2$ then we could apply an elementary transformation to get both e_1 and e_2 in the same forest of influence and so we are in case 1; hence $\operatorname{Stab} f_2 < \operatorname{Stab} e_2$. Also since $\operatorname{Stab} y_1$ is not contained in $\operatorname{Stab} y_2$ and $\operatorname{Stab} x \leqslant \operatorname{Stab} y_1$ we have

$$\begin{aligned} \operatorname{Stab} e_2 &= \operatorname{Stab} y_1 \cap \operatorname{Stab} y_2 \\ &< \operatorname{Stab} y_1 \\ &= \operatorname{Stab} f_1. \end{aligned}$$

Hence by Proposition 3.2.10 (c)

$$\begin{aligned} W_{S,\mathcal{P}} - W_{\tilde{S},\mathcal{P}} &= W_{e_2,\mathcal{P}} - W_{f_1,\mathcal{P}} - W_{f_2,\mathcal{P}} \\ &\geq W_{e_2,\mathcal{P}} - \frac{1}{2}W_{e_2,\mathcal{P}} - \frac{1}{2}W_{e_2,\mathcal{P}} \\ &= 0 \end{aligned}$$

Case 3c $y_1 = u_1$ and so f_1 doesn't exist.

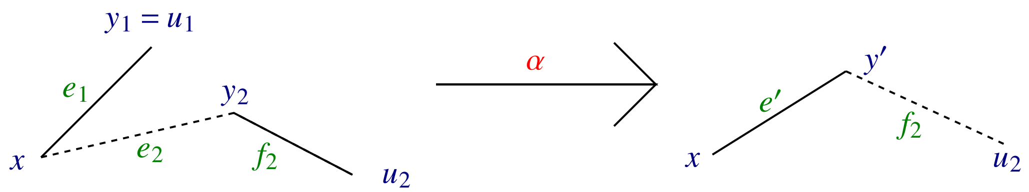

Fig. 3.10 The picture that arises when $y_1 = u_1$. Here we take the new forest of influence to be the image of the old one with f_2 removed (if it exists).

Observe that $\tilde{S} = \alpha(S)$ is a set of seed vertices which grows into a forest of influence $\tilde{\Gamma} = \alpha(\Gamma) \setminus G\{f_2\}$. Since $\operatorname{Stab} e_2 \leqslant \operatorname{Stab} y_2 = \operatorname{Stab} f_2$ we have $W_{e_2,\mathcal{P}} \geq W_{f_2,\mathcal{P}}$ and so the $\mathcal{P}$–weight can't increase.

□

Recall that a free action is a special case as we allow the set of seed vertices to be empty. Thus we must deal with this case separately.

Lemma 3.2.15. *Suppose that $\alpha : R \to \tilde{R}$ is a fold and S is a set of seed vertices for R. Suppose also that $\mathcal{P}$ has finite height. If G acts freely on R but not on $\tilde{R}$ then there is a set of seed vertices $\tilde{S}$ for $\tilde{R}$ such that $W_{\mathcal{P},\tilde{S}} \leq W_{\mathcal{P},S}$.*

Proof. Recall from Definition 3.1.5 that we must have $W_{\mathcal{P},S} \geq (\beta_1(R/G) - 1)W_{\mathcal{P},1}$. As the action on $\tilde{R}$ is non-free and α is a fold it follows that α is a fold of type III. Moreover all the edge stabilisers of $\tilde{R}$ are trivial and there is a single vertex u (up to equivalence) with a non-trivial stabiliser. We define $\tilde{S} := G\{u\}$. Since all the connecting groups of $\tilde{S}$ are trivial we get that $W_{\mathcal{P},\tilde{S}} = \beta_1(\tilde{R}/G)W_{\mathcal{P},1}$. Since α is a fold of type III we have $\beta_1(\tilde{R}/G) = \beta_1(R/G) - 1$ and so the result follows. □

Now suppose that we have a map $\Psi : T \to T'$ where T has a set of seed vertices S and the action on T' is k-acylindrical on groups larger than $\mathcal{P}$. If a connecting group in S is larger than $\mathcal{P}$ then there are a pair of seed vertices, say u_1 and u_2, whose images in T' are separated by distance at most k. Since we have control of the length of the path between these images we wish to collapse it to avoid unnecessary extra counting. However in general the image of the path between u_1 and u_2 need not lie in the path between $\Psi(u_1)$ and $\Psi(u_2)$. For our core argument to work we require this containment and the following says we can do this with some extra folds.

Lemma 3.2.16. *Let $\Psi : T \to T'$ be a simplical map where the action on T' is k-acylindrical on groups larger than $\mathcal{P}$. Let S be a set of seed vertices for T where at least one of the connecting groups are larger than $\mathcal{P}$. Then there are $\overline{T}$, $\overline{T'}$ and a simplical $\overline{\Psi} : \overline{T} \to \overline{T'}$ such that the action on $\overline{T'}$ is k-acylindrical on groups larger than $\mathcal{P}$, there's a set of seed vertices $\overline{S}$ for $\overline{T}$ with $W_{\mathcal{P},\overline{S}} \leq W_{\mathcal{P},S'} - 1$ and T' has at most k more edges than $\overline{T'}$.*

Proof. Let e be a connecting edge of $\mathcal{P}$–weight 1, (in some forest of influence Γ), which connects the trees of influence of u_1 and u_2. Let γ be the reduced edge path between u_1 and u_2. Recall statement ($\star$) from the proof of Stallings folding theorem (Theorem 2.2.2) and apply it to γ to get a composition of folds $\rho : T \to \tilde{T}$ so that the induced map $\tilde{\Psi} : \tilde{T} \to T'$ is locally injective on $\rho(\gamma)$. The intersection of the stabilisers for u_1 and u_2 is larger than $\mathcal{P}$ since e has $\mathcal{P}$–weight 1. So since the action on T' is k–acylindrical on groups larger than $\mathcal{P}$ the distance between the $\rho(u_i)$ is at most k. Hence we can collapse at most k edges of T' to get a new tree $\overline{T'}$ and an induced simplicial map $\overline{\Psi} : \overline{T} \to \overline{T'}$. (The action on $\overline{T'}$ is k–acylindrical on groups larger than $\mathcal{P}$ as it is formed by collapsing edges of T', a process which cannot increase the diameter of a fixed point set.) Let $\overline{S}$ be the image of S in $\overline{T}$. By construction we see that the stabiliser of every edge in $\overline{T}$ is either conjugate into either a stabiliser of a member of S or $\langle \mathrm{Stab}\, u_1, \mathrm{Stab}\, u_2 \rangle$. Hence $\overline{S}$ is a set of seed vertices with $W_{\mathcal{P},\overline{S}} \leq W_{\mathcal{P},S} - 1$ as the image of the connecting edges of Γ except the orbit of e is a set of connecting edges for $\overline{S}$. □

3.3 Building partially reduced trees

It remains to bound the number of edges of a k–acylindrical action on a tree given a set of seed vertices.

Lemma 3.3.1. *Let $\mathcal{P}$ be a class of subgroups for a group G which is closed under conjugation. Suppose G acts on a tree T and that this action is both partially-reduced on $\mathcal{P}$ and k–acylindrical on groups larger than $\mathcal{P}$. Suppose that S is a non-empty set of seed vertices for T with n orbits of connecting edges. Then T/G has at most $(2k+1)n$ edges. Furthermore if T is reduced and $k > 1$ then T/G has at most $2kn$ edges.*

Proof. First observe that we can assume that each connecting group is a subgroup of a group in $\mathcal{P}$. Indeed suppose that there are $r > 0$ connecting groups which are larger than $\mathcal{P}$. For each of the corresponding connecting edges we see that path consisting of it together with the branches of both its endpoints must be fixed by the connecting group, which is larger than $\mathcal{P}$. Since the action is k–acylindrical on groups larger than $\mathcal{P}$ each of these paths have length at most k. Thus we can collapse these paths to get a new tree with at most kr fewer edges and a set of seed vertices with r fewer connecting groups.

Let $F \subset T$ be the forest consisting of S together with every edge and vertex whose stabiliser is larger than $\mathcal{P}$. Since all of the connecting groups are contained in a member of $\mathcal{P}$ we see that F must deformation retract to S. Let $R \subseteq T$ be a maximal subtree where every edge stabiliser is contained in a member of $\mathcal{P}$. Let $A = R \cap F$, the vertices with stabiliser larger than $\mathcal{P}$ and seed vertices which are in R. We define $\tilde{R}$ as the union of R and the branches of each $v \in A$.

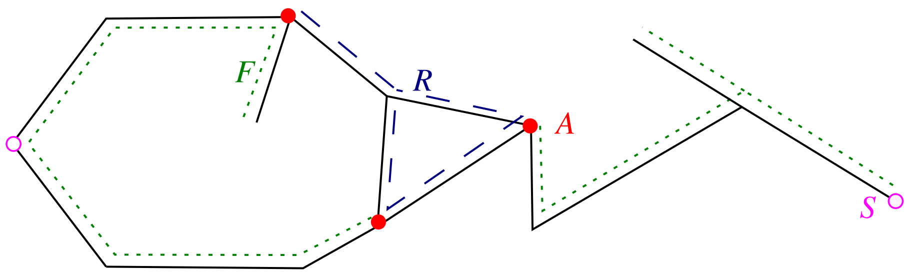

Fig. 3.11 An example of what this construction may look like in the graph of groups. Here we see that $p := |A| = 3$, $\chi(R/G) = 0$ and that R contains three orbits of connecting edges; the bottom right one and any two of the other three in R/G. Note the leftmost member of A has valence 1 in R/G and isn't in S, so if the tree is reduced F must extend beyond it. Because of this the length of the branch of this vertex is actually bounded above by $k-1$ instead of just the usual k.

Let $|A/G| = p$ and suppose R contains q connecting edges (up to equivalence) of some forest of influence Γ grown from S. Now the branch of each vertex of R must contain a vertex in A. Moreover this is unique as two distinct members of A are influenced by different seed vertices and so must lie in different components of Γ. Therefore $R \cap \Gamma$ deformation retracts to A and so $\chi(R/G) = p - q$.

Recall from Remark 3.2.3 that every non-seed vertex has an edge with equal stabiliser to it and hence every vertex of R/G with valence 1 or 2 (in R/G) must be in A as T is partially reduced on $\mathcal{P}$. Suppose that R/G has n_i vertices of valence i and observe that $p \geq n_1 + n_2$. Observe that $\chi(R/G) = \frac{1}{2}(n_1 - n_3 - 2n_4 - \cdots)$ and hence $\sum_{i \geq 3} n_i \leq n_1 - 2\chi(R/G)$. So

$$\begin{aligned} \#(\text{edges of } R/G) &= \#(\text{vertices of } R/G) - \chi(R/G) \\ &\leq 2n_1 + n_2 - 3\chi(R/G) \end{aligned}$$

Now the length of the branch of each $v \in A$ is at most k as the action is k–acylindrical on $\mathcal{P}$. (If $v \in A \setminus S$ then $\operatorname{Stab} v$ is lager than $\mathcal{P}$ and fixes its branch.) Hence we see that

$$\begin{aligned} \#(\text{edges of } \tilde{R}/G) &\leq \#(\text{edges of } R/G) + kp \\ &\leq 2n_1 + n_2 + kp - 3\chi(R/G) \\ &\leq (k+2)p - 3\chi(R/G) \end{aligned}$$

We now split into cases depending on the value of $\chi(R/G)$. First suppose that $\chi(R/G) \leq 0$. Then $q \geq p$ and so

$$\begin{aligned} \#(\text{edges of } \tilde{R}/G) &\leq (k+2)p - 3\chi(R/G) \\ &= (k-1)p + 3(p - \chi(R/G)) \\ &\leq (k-1)q + 3q \\ &= (k+2)q \end{aligned}$$

Otherwise $\chi(R/G) = 1$ and we have $q = p - 1$ and so

$$\begin{aligned} \#(\text{edges of } \tilde{R}/G) &\leq (k+2)p - 3\chi(R/G) \\ &\leq (k+2)(q+1) - 3 \\ &= (k+2)q + (k-1) \\ &\leq (k+2)q + (k-1)q \\ &= (2k+1)q \end{aligned}$$

Let $\tilde{T}$ be the tree obtained by collapsing each edge of $G\tilde{R}$ and let $\pi : T \longrightarrow \tilde{T}$. Observe that $\tilde{S} := \pi(S)$ is a set of seed vertices for $\tilde{T}$ and that the number of connecting edges of $\tilde{S}$ is $n - q$. Hence by induction $\tilde{T}/G$ has at most $(2k+1)(n-q)$ edges. Combining this with the above we see that T/G has at most $(2k+1)n$ edges as required.

It remains to show the improved bound if T is reduced and $k > 1$. In this case any $v \in A \setminus S$ where v/G has valence 1 in R/G must be the endpoint of at least 2 edges not contained in

R/G. This means that the path from v to the corresponding $u \in S$ actually has length at most $k-1$. Hence in this case

$$\begin{aligned}\#(\text{edges of } \tilde{R}/G) &\leq \#(\text{edges of } R/G) + k(p-n_1) + (k-1)n_1 \\ &\leq (k+1)n_1 + n_2 + k(p-n_1) - 3\chi(R/G) \\ &\leq (k+1)p - 3\chi(R/G)\end{aligned}$$

The rest of the calculations are essentially the same as before and so are omitted for the sake of brevity. We will note however that we only actually obtain the improved bound if $k > 1$. (In the case where $\chi(\tilde{R}/G) = 1$ we need $(k-2) \leq (k-2)q$. Since $q \geq 1$ we need $k \geq 2$.) □

We now have all the pieces we need to prove Lemma 3.2.13 and hence Theorem 3.1.7.

Proof of Lemma 3.2.13. We will proceed by induction on $W_{\mathcal{P},S'}$. If $W_{\mathcal{P},S'} = 0$ then since G isn't isomorphic to $\mathbb{Z}$ the action isn't free and so there is a single seed vertex in T', which must be fixed by G. The image of this vertex in T is fixed by G and so as T is minimal it must just consist of a single vertex.

So without loss of generality $W_{\mathcal{P},S'} > 0$. Start by using Stallings folding theorem (Theorem 2.2.2) to decompose α into folds $\alpha_i : T_{i-1} \to T_i$ and let $S_0 := S'$. Recursively for each $i > 0$ if S_{i-1} is defined and has connecting edge groups contained in $\mathcal{P}$ we obtain a set of seed vertices S_i for T_i at each step using either Lemma 3.2.14 or Lemma 3.2.15 (depending on if the action on T_{i-1} is free). If the $\mathcal{P}$–weight at any step decreases then we are done by induction on $W_{\mathcal{P},S'}$.

If instead S_{i-1} has a connecting edge which is larger than $\mathcal{P}$ we apply Lemma 3.2.16. We obtain $\overline{\Psi} : \overline{T}_{i-1} \to \overline{T}$ where the action on $\overline{T}$ is k-acylindrical on groups larger than $\mathcal{P}$. There's a set of seed vertices $\overline{S}_{i-1}$ for $\overline{T}_{i-1}$ with $W_{\mathcal{P},\overline{S}_{i-1}} \leq W_{\mathcal{P},S'} - 1$ and T has at most k more edges than $\overline{T}$. Hence by induction on $W_{\mathcal{P},S'}$ we see that $\overline{T}/G$ has at most $(W_{\mathcal{P},S'} - 1)k$ edges, hence T/G has at most $W_{\mathcal{P},S'}k$ edges as desired.

So without loss of generality we can assume that S_i is always defined and that both the $\mathcal{P}$–weight is constant and that we never have a connecting edge of $\mathcal{P}$–weight 1. Recall that Lemma 3.2.14 says that $\alpha_i(S_{i-1}) \subseteq S_i$ and since the number of connecting edges is bounded above by $\frac{1}{2}W_{\mathcal{P},S'}$ we must have $S_i = \alpha_i(S_{i-1})$ for all sufficiently large i. Thus by taking limits we see that there is a set of seed vertices S for T with $W_{\mathcal{P},S} = W_{\mathcal{P},S'}$. Now Lemma 3.3.1

implies that the number of edges of T/G is bounded above by $\left(\frac{2k+1}{2}\right)W_{\mathcal{P},S'}$. (Since each connecting edge has weight of at least 2.)

□

Proof of Theorem 3.1.7. First we use Dunwoody's resolution lemma (Theorem 2.1.12) to get G acting on a tree T' which has at most $C(G)$ orbits of edges together with a combinatorial map $\Psi : T' \to T$. Let S be the set of vertices of T' before subdividing. Observe that S is a set of seed vertices for T' and that it has $\mathcal{P}$–weight of at most $2^M C(G)$ since $\mathcal{P}$ has height M. Hence by Lemma 3.2.13 we see that T/G has as most $(2k+1)2^{M-1}C(G)$ edges. □

3.4 Extending to the main results

Now that we have finished proving our simplified result it's time to extend it to get our main theorems. The first way we're going to do this is to show that we don't require $\mathcal{P}$ to be closed under taking subgroups; although it still must satisfy condition (†). (See Definition 3.1.10 for the statement of (†).) The following is the analogue to Lemma 3.2.13 in this context.

Lemma 3.4.1. *Let G be a non-cyclic group and let $\mathcal{P}$ be a conjugation invariant set of subgroups of G which satisfies* (†)*. Let G act on a tree T and suppose this action is $\mathcal{P}$–partially-reduced and k–acylindrical on subgroups larger than $\mathcal{P}$. Let G act on another tree T' and let there be a G–equivariant combinatorial map $\Psi : T' \to T$. Suppose that T' has a set of seed vertices S with finite $\mathcal{P}$–weight $W_{\mathcal{P},S}$ and T is $\mathcal{P}$–closed. Then T/G has at most $\left(\frac{2k+1}{2}\right)W_{\mathcal{P},S}$ edges.*

The added difficulty is that Lemma 3.2.14 requires every connecting group to be in $\mathcal{P}$. Previously this was not an issue as every subgroup of G was either in $\mathcal{P}$ or larger than it. We will solve this problem by adding extra folds at each step which forces the connecting groups to be in $\mathcal{P}$.

Lemma 3.4.2. *Suppose $\Psi : T' \to T$ and $\beta : T' \to R$ are G-equivariant combinatorial maps where β factors through Ψ. Let S be a set of seed vertices for R with finite $\mathcal{P}$–weight $W_{\mathcal{P},S}$ and where none of the connecting groups are larger than $\mathcal{P}$. Suppose that $\mathcal{P}$ satisfies* (†) *and T is $\mathcal{P}$–closed. Then there is a simplicial map $\rho : R \to R'$ which factors through Ψ such that R' has a set of seed vertices $S' := \rho(S)$ with $W_{\mathcal{P},S'} \leq W_{\mathcal{P},S}$ and all its connecting groups are in $\mathcal{P}$.*

Proof. Let Γ be any forest of influence which is grown from S. If each connecting group of S is in $\mathcal{P}$ then we are done; so without loss of generality there is some connecting edge group $\operatorname{Stab} e$ which is not in $\mathcal{P}$. Since $\mathcal{P}$ satisfies (†) we have $H \in \mathcal{P}$ which is a minimal extension of $\operatorname{Stab} e$ to $\mathcal{P}$ and acts elliptically on R. Suppose H fixes the vertex v in R. Let p be the reduced edge path which starts at v and has final edge e. Let $\tilde{p}$ be the union of p together with the branch of each vertex on p. (See Figure 3.12.) Since T is $\mathcal{P}$–closed and the stabiliser of each edge f in $\tilde{p}$ contains $\operatorname{Stab} e$ we see that image of f in T must be stabilised by H. Let ρ be the (possibly infinite) composition of type II folds which "pulls" H onto each edge of $\tilde{p}$ and observe that this factors through Ψ since T is $\mathcal{P}$–closed. Hence if e' is a connecting edge of Γ with stabiliser in $\mathcal{P}$ then either $\operatorname{Stab} e' = \operatorname{Stab} \rho(e')$ or $W_{\mathcal{P},e'} < W_{\mathcal{P},e}$. Moreover $\rho(\Gamma)$ is a forest of influence grown from the seed vertices $\rho(S)$ with $W_{\mathcal{P},S'} \leq W_{\mathcal{P},S}$. Hence we can apply this process finitely many times until we get the result.

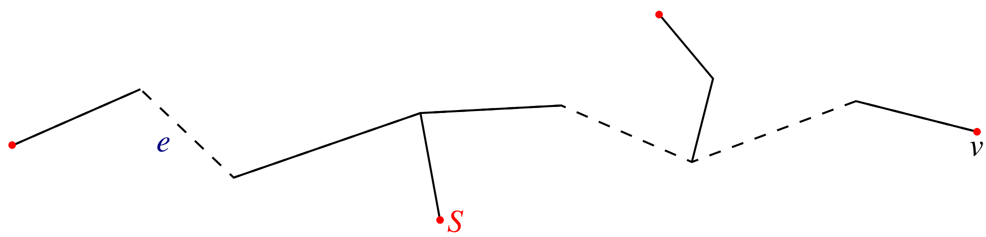

Fig. 3.12 An example of the domain we need to apply extra folds to. The reduced path from v to e is p while the entire diagram is $\tilde{p}$. We need to apply folds to get H fixing this whole region and not just p as otherwise we would need additional seed vertices.

□

Proof of Lemma 3.4.1. As before we will proceed by induction on $W_{\mathcal{P},S}$. If $W_{\mathcal{P},S} = 0$ then there is some vertex of T' which is fixed by G. So the image of this vertex in T is fixed by G and so as T is minimal it must just consist of just this single vertex.

Use Stallings folding theorem (Theorem 2.2.2) to decompose α into folds $\alpha_i^{(0)} : T_{i-1}^{(0)} \to T_i^{(0)}$. We will iteratively define trees $T_i^{(j)}$ (for $i \geq j-1$) and sets of seed vertices $S_i^{(j)}$ for $T_i^{(j)}$ (for $i = j$ and $i = j+1$) together with maps $\alpha_i^{(j)} : T_{i-1}^{(j)} \to T_i^{(j)}$ and $\rho_i^{(j)} : T_i^{(j)} \to T_i^{(j+1)}$ as follows for each $j > 0$.

$$\begin{array}{ccccccc}
T_0^{(0)} & \xrightarrow{\alpha_1^{(0)}} & T_1^{(0)} & \xrightarrow{\alpha_2^{(0)}} & T_2^{(0)} & \xrightarrow{\alpha_3^{(0)}} & \cdots \\
\downarrow{\scriptstyle \rho_0^{(0)}} & & \downarrow{\scriptstyle \rho_1^{(0)}} & & \downarrow{\scriptstyle \rho_2^{(0)}} & & \\
T_0^{(1)} & \xrightarrow{\alpha_1^{(1)}} & T_1^{(1)} & \xrightarrow{\alpha_2^{(1)}} & T_2^{(1)} & \xrightarrow{\alpha_3^{(1)}} & \cdots \\
 & & \downarrow{\scriptstyle \rho_1^{(1)}} & & \downarrow{\scriptstyle \rho_2^{(1)}} & & \\
 & & T_1^{(2)} & \xrightarrow{\alpha_2^{(2)}} & T_2^{(2)} & \xrightarrow{\alpha_3^{(2)}} & \cdots \\
 & & & & \downarrow{\scriptstyle \rho_2^{(2)}} & & \\
 & & & & \vdots & & \ddots
\end{array}$$

First define $\beta_j := \alpha_{j-1}^{(j-1)} \circ \rho_{j-2}^{(j-2)} \circ \cdots \circ \rho_1^{(1)} \circ \alpha_1^{(1)} \circ \rho_0^{(0)}$. (Part of the red path along the bottom of the diagram which ends with a right facing arrow.) Applying Lemma 3.4.2 to β_j we obtain a map $\rho_j^{(j)} : T_{j-1}^{(j-1)} \to T_{j-1}^{(j-1)}$ where $S_{j-1}^{(j)} := \rho^{(j)}(S_{j-1}^{(j-1)})$ is a set of seed vertices where the connecting groups are in $\mathcal{P}$ and $W_{\mathcal{P},S_{j-1}^{(j)}} \leq W_{\mathcal{P},S_{j-1}^{(j-1)}}$. If $\alpha_i^{(j)}$ folds together the edges e_1 and e_2 we define $\alpha_i^{(j+1)}$ (for $i > j$) to be the fold (or identity) obtained by identifying the $\rho_{i+1}^{(j)}$ images of the e_i. Finally define $\rho_i^{(j+1)}$ (for $i > j+1$) to make the above diagram commute. Finally Lemma 3.2.14 says that the fold $\alpha_j^{(j)}$ induces a set of seed vertices $S_j^{(j)}$ on $T_j^{(j)}$ with $W_{\mathcal{P},S_j^{(j)}} \leq W_{\mathcal{P},S_{j-1}^{(j)}}$. (If a separating edge is larger than $\mathcal{P}$ then we can reduce the $\mathcal{P}$–weight by collapsing at most k edges using Lemma 3.2.16 and then proceeding by induction on $W_{\mathcal{P},S}$.)

As $\alpha = \cdots \alpha_2^{(0)} \circ \alpha_1^{(0)}$ we see that $\alpha = \cdots \alpha_2^{(2)} \circ \rho(1) \circ \alpha_1^{(1)} \circ \rho^{(0)}$ by the definitions of $\alpha_i^{(j)}$ and $\rho^{(j)}$. Moreover at each step we have a set of seed vertices with non-increasing $\mathcal{P}$–weight, the number of orbits of connecting edges are non-decreasing and hence all but finitely many of the sets of seed vertices are the image of the seed vertices at the previous level. At this point the proof is exactly the same as the proof of Lemma 3.2.13. □

Proof of Theorem 3.1.12 *(a).* First we use Dunwoody's resolution lemma (Theorem 2.1.12) to get G acting on a tree T' which has at most $\alpha(G)$ orbits of edges together with a combinatorial map $\Psi : T' \to T$. Let S be the set of vertices of T' before subdividing. Observe that S is a set of seed vertices for T' and that it has $\mathcal{P}$–weight of at most $2^M C(G)$ since $\mathcal{P}$ has height M. Hence by Lemma 3.2.13 we see that T has as most $(2k-1)2^{M-1}C(G)$ edges. □

It remains to extend Theorem 3.1.7 to finitely generated groups. An immediate hurdle for this is the lack of Dunwoody's resolution lemma (Theorem 2.1.12), as this only holds for (almost) finitely presented groups. Instead take a finite generating set X for G and consider the free group $F(X)$ acting freely on a tree T'. Whenever G acts on a tree T we see that there is an $F(X)$–equivariant combinatorial map $T' \to T$. It's this map which we intend to decompose into folds and apply our prior methods to.

Before stating the analogue to Lemma 3.2.13 we first need to extend the definition of $\mathcal{P}$–weights.

Definition 3.4.3. *Let $\phi : H \to G$ be a surjective homomorphism of groups. Let $\mathcal{P}$ be a set of subgroups for G which is closed under conjugation. We define the $\mathcal{P}$–*weight *of $K \leq H$, (denoted $W_{\mathcal{P},\phi,K}$ or $W_{\mathcal{P},K}$ if ϕ is understood,) to be equal to $W_{\mathcal{P},\phi(K)}$. If H acts on a tree with a non-empty set of seed vertices S then we define its $\mathcal{P}$–*weight *$W_{\mathcal{P},\phi,S}$ (or $W_{\mathcal{P},S}$ if ϕ is understood) to be equal to the sum of the $\mathcal{P}$–weights of the connecting groups.*

Lemma 3.4.4. *Let $\mathcal{P}$ be a conjugation invariant set of subgroups of G which is closed under taking subgroups. Let G act on a tree T and suppose this action is $\mathcal{P}$–partially-reduced and k–acylindrical on subgroups larger than $\mathcal{P}$. Suppose also that G' is a countable group acting on a tree T' and that the following conditions hold.*

- *There is a surjective homomorphism $\phi : G' \to G$.*
- *The kernel of ϕ has trivial intersection with every edge stabiliser of T'.*
- *There is a G'-equivariant combinatorial map $\Psi : T' \to T$. (Where the action of G' on T is the natural one given by ϕ.)*
- *T' has a set of seed vertices S' with $\mathcal{P}$–weight $W_{\mathcal{P},S'}$.*

Then T/G has at most $\left(\frac{2k+1}{2}\right)W_{\mathcal{P},S'}$ edges. Furthermore if either T is reduced and $k > 1$ or all of the edges of T have stabiliser of size greater than $\mathcal{P}$ then T/G has at most $kW_{\mathcal{P},S'}$ edges.

First observe that the following variation of Lemma 3.2.14 and Lemma 3.2.15 holds with the exact same proof as before.

Lemma 3.4.5. *Let $\phi : H \to G$ be a surjective homomorphism and $\mathcal{P}$ is a conjugation invariant set of subgroups for G. Suppose that there is a H–equivariant map $\alpha : R \to \tilde{R}$ which is a fold. Suppose that the kernel of ϕ has trivial intersection with each vertex stabiliser of R. Suppose that S is a set of seed vertices for R where the image of each connecting group is in $\mathcal{P}$. Then there is a set of seed vertices $\tilde{S}$ for $\tilde{R}$ with $W_{\mathcal{P},\phi,\tilde{S}} \leq W_{\mathcal{P},\phi,S}$. Moreover if $W_{\mathcal{P},\phi,\tilde{S}} = W_{\mathcal{P},\phi,S}$ then $\alpha|_S$ is injective.*

Observe from our discussion of different fold types on page 19 that the new edge stabilisers will always be contained in an old vertex group. The same is not true of the vertex groups however. In particular after applying Lemma 3.4.5 it may be the case that elements in the kernel of ϕ may end up acting elliptically on the intermediate tree, although they will only fix a single vertex. As such we need a way of modifying a group G_i and tree T_i which essentially keeps the action and map $\Psi_i : T_i \to T$ but removes elliptically acting group elements found in the kernel of $\phi : G_i \to G$.

Lemma 3.4.6. *Let $\phi : G' \to G$ be a surjective group homomorphism and suppose that G' acts on a tree T' where no element in the kernel of ϕ fixes an edge. Then there's a group G'' acting on a tree T'' together with surjective homomorphisms $\phi' : G'' \to G$ and $\sigma : G' \to G''$ with $\phi = \phi' \circ \sigma$ and a G'–equivariant simplical map $\rho : T' \to T''$ where the following conditions hold. (The action of G' on T'' is given by σ.)*

- *The map $\rho/G : T'/G \to T''/G$ is a homeomorphism of graphs.*
- *For each edge e of T' we have $\sigma(\mathrm{Stab}\, e) = \sigma(\mathrm{Stab}\, \rho(e))$.*
- *The kernel of ϕ' has trivial intersection with every vertex (and edge) stabiliser of T''.*

In particular if T' has a set of seed vertices S' then $S'' := \rho(S')$ is a set of seed vertices for T'' with $W_{\mathcal{P},\phi',S''} = W_{\mathcal{P},\phi,S'}$.

Proof. We define G'' as the fundamental group of a graph of groups decomposition corresponding to T' but with each vertex label replaced with its image under ϕ. Let T'' be the Bass-Serre tree corresponding to this modified graph of groups. This naturally induces maps $\phi' : G'' \to G$ and $\tilde{\phi} : G' \to G''$ and $\sigma : T'' \to T$ with all the desired properties. □

Proof of Lemma 3.4.4. As before we will proceed by induction on $W_{\mathcal{P},S}$. If $W_{\mathcal{P},S} = 0$ then there is some vertex of T' which is fixed by G. So the image of this vertex in T is fixed by G and so as T is minimal it must just consist of this single vertex.

Use Stallings folding theorem (Theorem 2.2.2) to decompose α into folds $\alpha_i^{(0)} : T_{i-1}^{(0)} \to T_i^{(0)}$ and let $G^{(0)} := G'$. We will iteratively define groups $G^{(j)}$ which act on trees $T_i^{(j)}$ with sets of seed vertices $S_i^{(j)}$ for $T_i^{(j)}$. Also we will define maps $\sigma^{(j)} : G^{(j)} \to G^{(j+1)}$ and $\phi^{(j)} : G^{(j)} \to G$

together with $\alpha_i^{(j)} : T_{i-1}^{(j)} \to T_i^{(j)}$ and $\rho_i^{(j)} : T_i^{(j)} \to T_i^{(j+1)}$ (for $i \geq j-1$).

$$\begin{array}{ccccccc}
T_0^{(0)} & \xrightarrow{\alpha_1^{(0)}} & T_1^{(0)} & \xrightarrow{\alpha_2^{(0)}} & T_2^{(0)} & \xrightarrow{\alpha_3^{(0)}} & \cdots \\
\downarrow{\scriptstyle \rho_0^{(0)}} & & \downarrow{\scriptstyle \rho_1^{(0)}} & & \downarrow{\scriptstyle \rho_2^{(0)}} & & \\
T_0^{(1)} & \xrightarrow{\alpha_1^{(1)}} & T_1^{(1)} & \xrightarrow{\alpha_2^{(1)}} & T_2^{(1)} & \xrightarrow{\alpha_3^{(1)}} & \cdots \\
 & & \downarrow{\scriptstyle \rho_1^{(1)}} & & \downarrow{\scriptstyle \rho_2^{(1)}} & & \\
 & & T_1^{(2)} & \xrightarrow{\alpha_2^{(2)}} & T_2^{(2)} & \xrightarrow{\alpha_3^{(2)}} & \cdots \\
 & & & & \downarrow{\scriptstyle \rho_2^{(2)}} & & \\
 & & & & \vdots & & \ddots
\end{array}$$

First we define $\beta_j := \alpha_j^{(j)} \circ \rho_{j-1}^{(j-1)} \circ \cdots \circ \rho_1^{(1)} \circ \alpha_1^{(1)} \circ \rho_0^{(0)}$. (Part of the red path along the bottom of the diagram which ends with a right facing arrow.) First use Lemma 3.4.6 on β_j and $\phi^{(j)}$ to obtain a map $\rho^{(j)} : T_j^{(j)} \to T_j^{(j+1)}$ together with group homomorphisms $\sigma^{(j)} : G^{(j)} \to G^{(j+1)}$ and $\phi^{(j+1)} : G^{(j+1)} \to G$. Also $S_j^{(j+1)} := \rho^{(j)}(S_j^{(j)})$ is a set of seed vertices where the images of the connecting groups are in $\mathcal{P}$ and $W_{\mathcal{P},S_{j-1}^{(j)}} \leq W_{\mathcal{P},S_{j-1}^{(j-1)}}$. If $\alpha_i^{(j)}$ folds together the edges e_1 and e_2 we define $\alpha_i^{(j+1)}$ (for $i > j$) to be the fold (or identity) obtained by identifying the $\rho_{i+1}^{(j)}$ images of the e_i. Finally define $\rho_i^{(j+1)}$ (for $i > j+1$) to make the above diagram commute.

Finally Lemma 3.4.5 says that the fold $\alpha_j^{(j)}$ induces a set of seed vertices $S_j^{(j)}$ on $T_j^{(j)}$ with $W_{\mathcal{P},S_j^{(j)}} \leq W_{\mathcal{P},S_{j-1}^{(j)}}$. (If a separating edge is larger than $\mathcal{P}$ then we can reduce the $\mathcal{P}$–weight by collapsing at most k edges using Lemma 3.2.16 and then proceeding by induction.)

As $\alpha = \cdots \alpha_2^{(0)} \circ \alpha_1^{(0)}$ we see that $\alpha = \cdots \alpha_2^{(2)} \circ \rho_1^{(1)} \circ \alpha_1^{(1)} \circ \rho_0^{(0)}$ by the definitions of $\alpha_i^{(j)}$ and $\rho_i^{(j)}$. Moreover at each step we have a set of seed vertices with non-increasing $\mathcal{P}$–weight, the number of orbits of connecting edges are non-decreasing and hence all but finitely many of the sets of seed vertices are the image of the seed vertices at the previous level. At this point the proof is exactly the same as the proof of Lemma 3.2.13. □

Proof of Theorem 3.1.12 *(b).* Pick a minimal generating set X for G and let $G' := F(X)$. Let $\phi : G' \to G$ be the natural projection and let T' be the tree corresponding to the rose with $\operatorname{rank} G$ petals labelled by the elements of X. Let $\Psi : T' \to T$ be any G'–equivariant combinatorial map. If G acts freely on T then the action is reduced and so every vertex in T/G has valence at least 3. Now Proposition 2.3.7 says that $\beta_1(T/G) \leq \operatorname{rank} G$ and thus Proposition 2.3.6 implies that T/G has at most $3(\operatorname{rank} G - 1)$ edges. Now suppose the action on T is not free and let S be the set of vertices of T' before subdividing. Observe that S is a set of seed vertices for T' and that it has $\mathcal{P}$–weight of at most $2^M \operatorname{rank}(G)$ since $\mathcal{P}$ has height M. Hence by Lemma 3.4.4 T/G has as most $(2k+1)2^{M-1}(\operatorname{rank}(G)-1)$ edges. If all of the edges of T have stabiliser larger than $\mathcal{P}$ or T is reduced with $k > 1$ then the number of edges is in fact bounded by $k2^M(\operatorname{rank}(G)-1)$. □

3.5 Sharpness of bounds

We will now restrict our attention to the case where $\mathcal{P} = 1$, the collection which only contains the trivial subgroup. In other words we are to consider actions which are k–acylindrical. In [26] Weidmann showed that a finitely generated group acting k–acylindrically on a tree where all the edges have non-trivial stabiliser has at most $2k(\operatorname{rank} G - 1)$ orbits of edges. Theorem 3.1.12 (b) immediately implies this by setting $k = 1$. The purpose of this section is to improve this bound to one which is the best possible.

Theorem 3.5.1. *Let G be a (non-cyclic) finitely generated group acting k–acylindrically on a minimal tree T (where $k \geq 1$.) Suppose that each edge of T has non-trivial stabiliser. Then T/G has at most $\left\lfloor \left(2\operatorname{rank} G - \frac{5}{2}\right) k \right\rfloor$ edges. If G is torsion-free then this bound can be improved to $(2\operatorname{rank} G - 3)\, k$.*

Theorem 3.5.2. *For any $k > 0$ and $r \geq 2$ there is a finitely presented group G with $\operatorname{rank} G = r$ which acts k–acylindrically on a minimal tree T where each edge of T has non-trivial stabiliser and T/G has exactly $\left\lfloor \left(2\operatorname{rank} G - \frac{5}{2}\right) k \right\rfloor$ edges.*

Similarly F_r admits a k–acylindrical action on a minimal tree T where each edge of T has non-trivial stabiliser and T/F_r has exactly $(2r-3)\, k$ edges.

Remark 3.5.3. *Unlike in the previous results there is no requirement that T needs to be reduced. Instead the conditions that T is k–acylindrical and has no edges with trivial stabiliser are enough to completely prevent the unrestricted edge subdivision which motivated the definition of a reduced action.*

We will start by constructing the examples with many edges as this will motivate the argument for the sharp bound.

Proof of Theorem 3.5.2. We need to show that for any integers $k \geq 1$ and $r \geq 2$ that there is a group of rank r acting k-acylindrically on a tree with $\left\lfloor 2k\left(\operatorname{rank} G - \frac{5}{4}\right)\right\rfloor$ orbits of edges, none of which have trivial stabilisers. Pick distinct primes p and q such that $(p-1)(q-1) \geq$

$2(r-2)$. Let $G := \langle a, h_1, \cdots, h_{r-1} \mid a^{pq} = 1 \rangle \cong \left(\frac{\mathbb{Z}}{pq\mathbb{Z}}\right) * F_{r-1}$ and note that $\operatorname{rank} G = r$. We will now construct a tree T for G to act on. Start with the graph of groups decomposition consisting of the rose with $r-1$ petals representing the h_i and with a single vertex on the loop representing h_1 with label $\langle a \rangle$. (In the diagrams we take $k = 4$ and $r = 3$.)

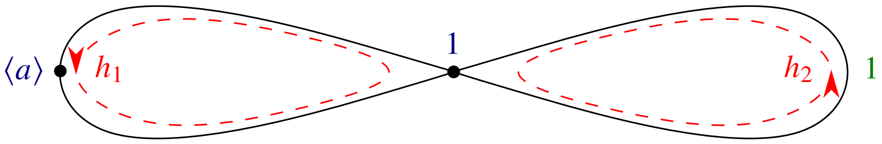

Subdivide the loop representing h_1 so that it consists of $\left\lceil \frac{k}{2} \right\rceil + 2$ edges. Apply folds of type II to "pull" a onto each vertex on the loop except the central one.

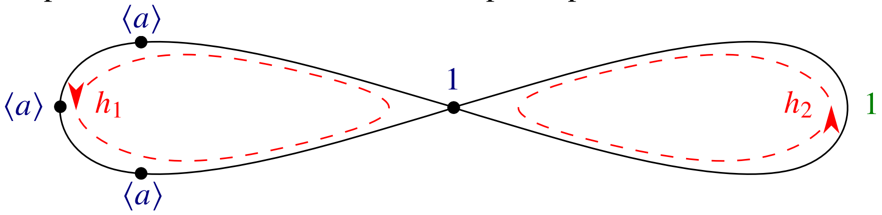

Next subdivide the edges on the loop representing h_1 which are adjacent to the central vertex into $\left\lfloor \frac{k}{2} \right\rfloor$ sub-edges. Apply folds of type II which "pull" a^p along one of these series of edges and "pull" a^q along the other. We see that the central vertex has stabiliser $H := \langle b, c \mid b^q = c^p = 1 \rangle \cong \left(\frac{\mathbb{Z}}{p\mathbb{Z}}\right) * \left(\frac{\mathbb{Z}}{q\mathbb{Z}}\right)$ where $b = a^p$ and $c = (a^q)^{h_1}$.

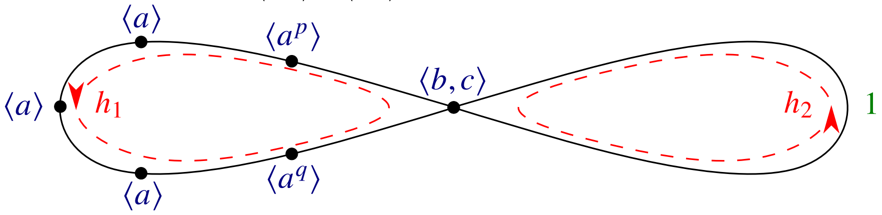

For $i \leq 2(r-2)$ we define $g_i = b^{x+1} c^{y+1}$ where $i = (p-1)x + y$ for $x, y \in \mathbb{Z}$ with $0 \leq y \leq p-2$. Since $(p-1)(q-1) \geq 2(r-2)$ we see that these g_i represent pairwise non-conjugate elements of H. Subdivide the loop representing each h_j (for $2 \leq j \leq r-1$) into $2k$ sub-edges, then apply folds of type II which "pulls" g_{2j-3} along k edges starting at one end and "pulls" g_{2j-2} along k edges starting at the other.

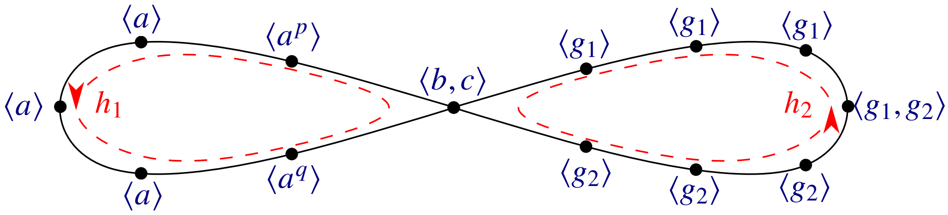

Observe that this decomposition has $\left\lfloor (2r - \frac{5}{2})k \right\rfloor$ edges. (It's not reduced in general, but recall that this isn't a condition for this result.) It remains to check that the corresponding Bass-Serre tree is k–acylindrical. The elements of G which act elliptically (upto conjugacy) are powers of a and elements of H; so these are the ones we need to check that they fix a region of bounded diameter.

First consider elements of H. The elements which fix an edge of our tree are (powers of) the g_i, b and c (upto conjugacy). As b and c are conjugate to powers of a we'll leave these for now. Now observe that each g_i has a different image in the abelianization of H; hence distinct g_i are in different conjugacy classes. Moreover each cyclic root-closed subgroup of H is malnormal in it. Hence each (power of) g_i only fixes k edges.

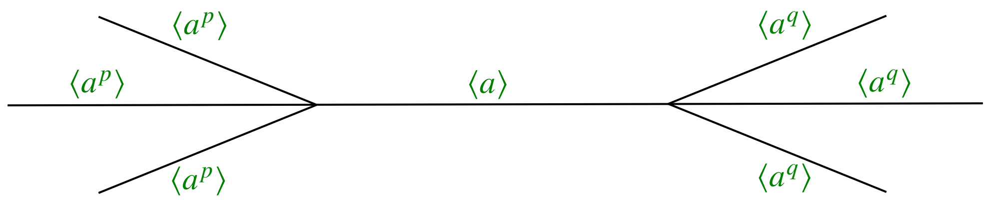

Fig. 3.13 The region of the Bass-Serre tree which is fixed by some power of a, together with their stabilisers. The central arc with stabiliser $\langle a \rangle$ has length $\lceil \frac{k}{2} \rceil$ while each "offshoot" with stabiliser either $\langle a^p \rangle$ or $\langle a^q \rangle$ has length $\lfloor \frac{k}{2} \rfloor$.

We now need to consider powers of a. Let $1 \leq m < pq$ and look at the fixator of a^m. If $p \,|\, m$ then the fixator of a^m consists of a central vertex with $q+1$ "offshoots", one of length $\lceil \frac{k}{2} \rceil$ and the rest of length $\lfloor \frac{k}{2} \rfloor$. In other words the fixator consists of the left and the centre parts of Figure 3.13. This region has diameter k and so we are fine. Likewise for the case where $q \,|\, m$. We cannot have $pq \,|\, m$ as $m < pq$. Finally if m is coprime to pq then p^m just fixes a path of length $\lceil \frac{k}{2} \rceil$; the middle section of Figure 3.13.

Building an example which is maximal for torsion-free groups is similar. First we need a to have infinite order and so $G \cong F_r$. The initial splitting is defined in the same way as before. Next we subdivide the loop representing h_1 into k subedges and apply folds of type II so that each edge in this loop has label $\langle a \rangle$. (If $k = 1$ then we collapse either of the initial edges of the loop instead.) The label of the central vertex is now isomorphic to the free group of rank 2 which is generated by $b := a$ and $c := a^{h_1}$. We now subdivide and fold onto the loops representing the rest of the h_j as before.

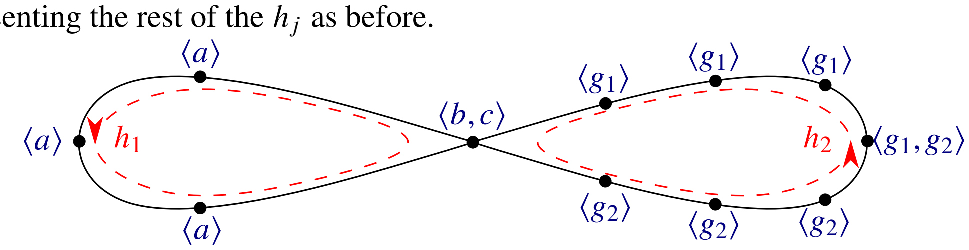

□

Before proving Theorem 3.5.1, which will show the above examples are the best possible, it's useful to compare their constructions to the proof of Lemma 3.4.4. We start with a single orbit of seed vertices with representative stabiliser $\langle a \rangle$. Our initial folds induce a new orbit of seed vertices on the central vertex. Moreover the two connecting edges on the loop representing h_1 are now non-trivial and so we collapse it. In doing this we'll reduce the 1–weight by two but only collapse either $\lfloor \frac{3k}{2} \rfloor$ or k edges, depending on which construction we're talking about. This is less than the $2k$ edges theoretically allowed by the lemma. Continuing we then successively collapse each loop; getting rid of the maximally possible $2k$ edges each time.

With this comparison in mind we will now show that such an inefficiency must occur at a particular point in Lemma 3.4.4. Specifically whenever a vertex first obtains a non-cyclic stabiliser.

Lemma 3.5.4. *Let $\alpha : R \to \tilde{R}$ be a fold which factors through $\Psi : T' \to T$ and let S be a set of seed vertices for R. Suppose that the action on T is k–acylindrical. Suppose also that both every vertex stabiliser of R is cyclic and every connecting group of S is trivial. Then one of the following holds*

- *Every vertex stabiliser of R is cyclic and we can find a set of seed vertices $\tilde{S}$ for $\tilde{R}$ such that every connecting group of $\tilde{S}$ is trivial and $W_{1,\tilde{S}} \leq W_{1,S}$.*
- *There is a simplicial map $\rho : \tilde{R} \to \tilde{R}'$ which factors through Ψ and we can collapse at most $\lfloor \frac{3k}{2} \rfloor$ orbits of edges of T to get a new tree $\overline{T}$ such that the following holds. Let $\overline{R}$ be the tree obtained by collapsing the edges of $\tilde{R}'$ corresponding to $T \to \overline{T}$. There is a set of seed vertices $\overline{S}$ for $\overline{R}$ with $W_{1,\overline{S}} \leq W_{1,S} - 2$. Moreover if G is torsion-free then we can obtain $\overline{T}$ by collapsing at most k edges of T.*

Proof. Suppose that α folds together the edges $e_1 = [x, y_1]$ and $e_2 = [x, y_2]$ to an edge $e' = [x', y']$. If there's a forest of influence containing both e_1 and e_2 then we can just take $\tilde{S} := \alpha(S)$ and we end up in the first outcome listed in the statement. The same applies if the fold is of type I or II and either of the y_i have trivial stabiliser. Similarly we can take $\tilde{S} := \alpha(S) \cup G\{y'\}$ if α is a fold of type III and the y_i have trivial stabilisers.

Consider the case where α is a fold of type III, the stabiliser of the y_i are non-trivial and $hy_1 = y_2$. Suppose that y_1 is influenced by u and $u' := \alpha(u)$. Both u' and y' are fixed by the stabiliser of y_1; hence as in the proof of Theorem 3.1.7 we can apply a series of folds

$\rho : \tilde{R} \to \tilde{R}'$ which factors through Ψ such that the reduced edge path f' from u' to y' consists of at most k edges and is injective under $\delta : \tilde{R}' \to T$. Now we define $\overline{T}$ by collapsing the image of f' in T and $\overline{R}$ as in the statement. Note that the image of S in $\overline{T}$ is a set of seed vertices with $W_{1,\overline{S}} \le W_{1,S} - 2$ as every connecting edge is trivial and $\chi(\overline{R}/G) = \chi(R/G) + 2$.

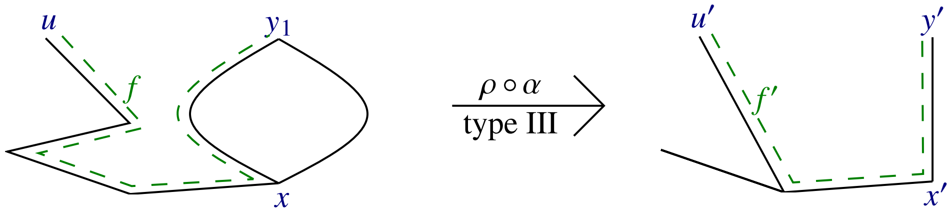

Fig. 3.14 The case where the fold α is of type III. The edge path labelled f is fixed by some non-trivial member of G, hence its image in $\tilde{R}$ has diameter at most k. In particular if we replace f with its image in $\tilde{R}$ we see that the path f' has length at most k.

So now we can assume that the fold is of type I or II where the stabiliser of both y_1 and y_2 are non-trivial. If α is a fold of type I we say that y_i is influenced by u_i and f_i be the branch of y_i. If instead α is a fold of type II we say that x is influenced by u_1 and y_2 is influenced by u_2. We also let f_1 be the union of the branch of x together with e_1 and let f_2 be the branch of y_2.

First consider the case where u_1 is inequivalent to u_2. Let $\rho : \tilde{R} \to \tilde{R}'$ be the composition of folds on f_1 and f_2 (separately) which causes $\gamma : \overline{R'} \to T$ to be locally injective on f_1 and f_2. Let u_i' be the vertex closest to $y'' := \rho(y')$ with stabiliser equal to $\operatorname{Stab} y_i$ and let f_i' be the reduced edge path from u_i' to y''. We now define $\overline{T}$ by collapsing the (orbits of the) images of f_1' and f_2' in T and $\overline{R}$ as in the statement. Observe that we have a set of seed vertices $\overline{S}$ for $\overline{R}$ defined to be the union of the image of $S \setminus G\{u_1, u_2\}$ together with the image of $G\{y''\}$ and observe that $W_{1,\overline{S}} \le W_{1,S} - 2$. It remains to bound the number of edges we've collapsed. If f_i' consists of more than a single vertex let g_i be a group element which fixes u_i' but no edge of f_i'. Then since $\operatorname{Stab} u_i' = \operatorname{Stab} u_i$ is cyclic we see that $\operatorname{Stab} y_i$ fixes the reduced edge path $f_i' \cup g_i f_i'$. So since the action on T is k–acylindrical we see that f_1' and f_2' each consist of at most $\frac{k}{2}$ edges each and so we have collapsed at most k edges in total.

Now consider the case $u_2 = hu_1$ for some $h \in G$. Define $\rho : \tilde{R} \to \tilde{R}'$, u_i', y'' and f_i' as before. Let $\tilde{f}$ be the path from u_1' to $h^{-1}u_2'$. Define $\overline{T}$ by collapsing the images of f_1', f_2' and $\tilde{f}$ in T. We have a set of seed vertices $\overline{S}$ for $\overline{R}$ defined to be the union of the image of $S \setminus G\{u_1\}$ together with the image of $G\{y''\}$ and again we have $W_{1,\overline{S}} \le W_{1,S} - 2$. It now

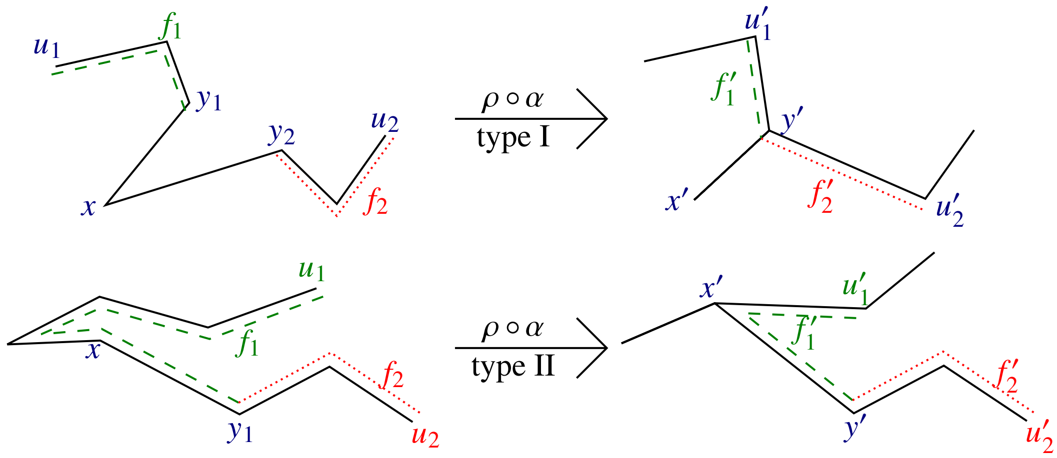

Fig. 3.15 The graph of groups in the cases where the u_i aren't equivalent.

remains to bound the number of edges collapsed. As before the paths f_i' have at most $\frac{k}{2}$ edges. If $\tilde{f}$ consists of just a single vertex u_1' then we are done as before. If not then observe that $f_1'\cup\tilde{f}\cup h^{-1}f_2'$ is a reduced edge path from y'' to $h^{-1}y''$. If G is torsion-free then as $\operatorname{Stab} u_1$ is cyclic then there is some non-trivial subgroup which fixes $f_1'\cup\tilde{f}\cup h^{-1}f_2'$ and so we've collapsed at most k edges. If G isn't torsion-free then we are only guaranteed to have a non-trivial subgroup which fixes $f_1'\cup\tilde{f}$. Thus $f_1'\cup\tilde{f}\cup h^{-1}f_2'$ has at most $k+\frac{k}{2}=\frac{3k}{2}$ edges.

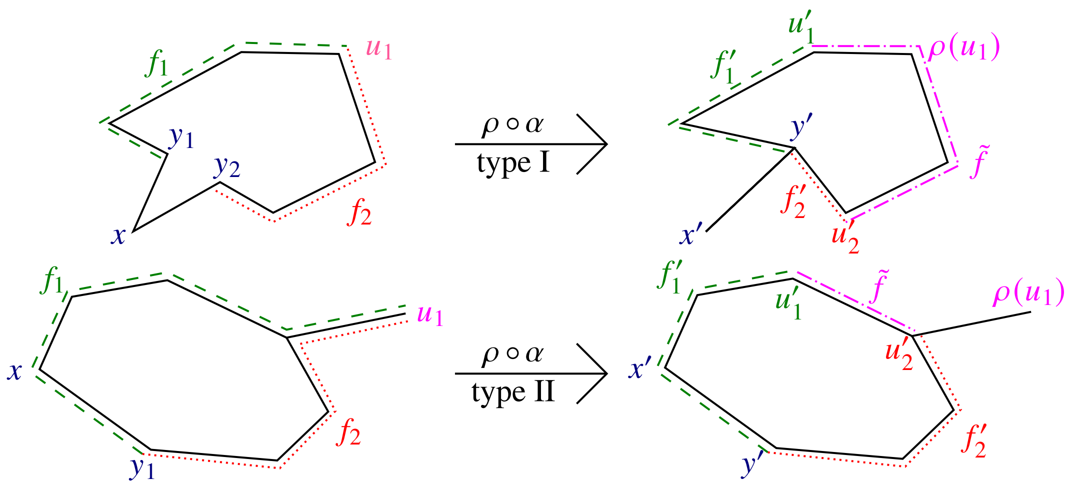

Fig. 3.16 The graph of groups in the cases where u_1 and u_2 are equivalent.

□

Proof of Theorem 3.5.1. Proceed as in the proof of Theorem 3.1.12 (b). We have a homomorphism $\phi: F(X)\to G$ where X is a minimal generating set for G and combinatorial map

$\Psi : T' \to T$. As in the proof of Lemma 3.4.4 we now decompose Ψ into folds $\alpha_i : T_{i-1} \to T_i$. We then apply Lemma 3.5.4 to each α_i in turn until one of them causes us to collapse edges. (This must happen eventually as all the edges of T have non-trivial stabiliser.) Then apply Theorem 3.1.12 (b) to the collapsed tree to get the desired bound. □

Chapter 4

Strong accessibility

So far we have analysed the conditions required to give a bound on the complexity of a single splitting. Now consider the vertex group of a maximal splitting over some class of edge groups. It is possible that such a vertex also admits a splitting over the same class of groups. It is therefore natural to ask if such a process can go on indefinitely. In this chapter we give an account of a theorem of Louder and Touikan [19]

4.1 Preliminaries

Before we give the statement of the main theorem we need several new definitions.

Definition 4.1.1. *A group is said to be* slender *if all its subgroups are finitely generated.*

Remark 4.1.2. *Note a slender group can only act elliptically, linearly or dihedrally on a tree. The analysis on page 19 shows that parabolic and hyperbolic actions always have infinitely generated subgroups.*

We will use the definition of a JSJ-decomposition given by Guirardel and Levitt. [14] First recall the definition of a minimal and reduced tree.

Definition 4.1.3. *The action of a group G on a tree T is* minimal *if there are no G–invariant proper subtrees. Such an action is said to be* reduced *if either T/G is a circle consisting of a single vertex and edge or the label of every vertex of valence 2 in T/G properly contains its edge groups.*

Definition 4.1.4. *Given a group G and sets of subgroups $\mathcal{A}$ and $\mathcal{B}$ we let $\mathrm{S}_{\mathcal{A},\mathcal{B}}$ be the set of reduced trees which G acts on with edge groups in $\mathcal{A}$ and where each group in $\mathcal{B}$ acts elliptically. We will assume that $\mathcal{A}$ is closed under conjugation and taking subgroups. If $\mathcal{B}$*

is empty then we shorten $\mathrm{S}_{\mathcal{A},\mathcal{B}}$ *to* $\mathrm{S}_{\mathcal{A}}$.

A tree $T \in \mathrm{S}_{\mathcal{A},\mathcal{B}}$ *is* universally elliptic *if every edge group of* T *is elliptic in any given tree in* $\mathrm{S}_{\mathcal{A},\mathcal{B}}$.

A G*-tree* T_1 dominates *another* T_2 *if there is a* G*-map* $T_1 \to T_2$. *Equivalently every vertex group of* T_1 *is elliptic in* T_2.

A tree $T \in \mathrm{S}_{\mathcal{A},\mathcal{B}}$ *is a* JSJ-tree *(over the class* $\mathcal{A}$ *relative to* $\mathcal{B}$*) if it's universally elliptic and dominates all other universally elliptic trees in* $\mathrm{S}_{\mathcal{A},\mathcal{B}}$. *The graph of groups corresponding to a JSJ-tree is called a* JSJ-splitting *or a* JSJ-decomposition *(over the class* $\mathcal{A}$ *relative to* $\mathcal{B}$*).*

Roughly speaking, we restrict our attention to splittings which do not "exclude" any other, then choose a maximal one amongst these. A priori it need not be the case that a JSJ-splitting exists and in complete generality they do not. However in many important cases they do in fact exist. In particular the following is true.

Lemma 4.1.5. *[14, Theorems 2.16 & 2.20] Let* G *be finitely presented (relative to some finite set of subgroups* $\mathcal{B}$*). Then for any class* $\mathcal{A}$ *there exists some JSJ-splitting for* G *over* $\mathcal{A}$ *(relative to* $\mathcal{B}$*) with finitely generated edge groups.*

Uniqueness of JSJ-trees does not hold in general, although one can often find a canonical choice for one-ended groups. However all of the JSJ-trees in a class live in a common deformation space. (See [14].) From this it follows that the stabilisers of the vertices (which aren't in $\mathcal{A}$) do not depend on the choice of JSJ-tree. Further the vertices of a JSJ-tree can be split into two classes.

Definition 4.1.6. *A vertex of a JSJ-tree is called* rigid *if it is elliptic in every tree in* $\mathrm{S}_{\mathcal{A},\mathcal{B}}$. *Otherwise it is called* flexible.

The flexible vertices should be thought of as analogous to the Seifert-fibred components of a JSJ-decomposition of a 3–manifold. The following result demonstrates this.

Lemma 4.1.7. *[14, Theorem 6.2] Let* v *be a flexible vertex of a slender JSJ-tree (relative to a finite set of finitely generated subgroups* $\mathcal{B}$*.) Then the stabiliser* G_v *of* v *is either slender or QH (quadratically hanging) with slender fibre. In other words, if* G_v *is not slender there is a short exact sequence*

$$0 \to A \to G_v \to \pi_1(\Sigma) \to 0$$

where A *is slender and* Σ *is a 2–dimensional orbifold with a non-trivial boundary.*

Such groups have very controlled JSJ-decompositions.

Lemma 4.1.8. *Let G be a non-slender group which is QH with slender fibre. The vertices of a slender JSJ-decomposition of G have slender stabiliser.*

Proof. First suppose that the fibre is trivial, so that $G = \pi_1(\Sigma)$ where Σ is a 2–dimensional orbifold with a non-trivial boundary. Observe that we can decompose G as $G_1 * \cdots * G_k * F_r$ where $G_1, \cdots, G_k$ are the finite groups associated to the orbifold points of Σ and F_r is the free group of rank r. Hence G has Grushko decompositions over trivial edge groups with finite vertex stabilisers. Such decompositions are exactly the slender JSJ-splittings of G.

Now suppose the fibre A is non-trivial. Let G act on a (reduced) tree T with slender edge stabilisers. If A acts trivially on such a T then we get an action of $\frac{G}{A} \cong \pi_1(\Sigma)$ on T. So if A acts trivially on T we can lift from $\pi_1(\Sigma)$ to see that the slender JSJ-trees of G have slender vertex groups, containing A with finite index. Now fix T. If A acts elliptically on T then since A is normal in G we have some vertex $v \in T$ such that $A \leqslant G_{gv}$ for any $g \in G$. Since T is reduced for any vertex $u \in T$ there is $g \in G$ such that $u \in [v, gv]$ and so $A \leqslant G_u$. Hence A acts trivially on T. If A doesn't act elliptically on T it must fix a line as it's slender. As A is normal in G we see that T must be exactly this line. However this implies that G is slender which contradicts our assumption on G. □

We will define a hierarchy using Bass-Serre trees instead of the standard method using graphs of groups. While there are advantages to both approaches the present author believes this to be better one for this argument, particularly for Section 4.5.

Definition 4.1.9. *A* hierarchy $\mathcal{H}$ *for a group G (over a class* $\mathcal{A}$*) is a rooted tree where for each vertex v of* $\mathcal{H}$ *we assign a subgroup* $G_v \leqslant G$ *with a minimal action on a tree* T_v *(with edge groups in* $\mathcal{A}$ *and) where the following conditions are satisfied.*

- *The initial vertex of* $\mathcal{H}$ *is assigned G as its group.*
- *If* T_v *is a point for some* $v \in \mathcal{H}$ *then v is a terminal vertex of* $\mathcal{H}$*.*
- *Otherwise there is a natural one-to-one correspondence between the immediate descendants of a vertex* $v \in \mathcal{H}$ *and the vertices of the tree* T_v*. Natural in this context means that the associated groups do not change under the correspondence. Henceforth we will often abuse notation and treat the vertices related under this correspondence as the same object.*

- $\mathcal{H}$ *is 'conjugation invariant'. More precisely if* w *and* gw *are vertices of some* T_v *(with* $g \in G_v$*), then the corresponding sub-hierarchies starting at* w *and* gw *are identical except that all the groups and labels for the vertices are conjugated by* g*. Equivalently* G *acts naturally on* $\mathcal{H}$ *and the stabiliser of any vertex is its associated group.*

It is clear that one reobtains the traditional definition of a hierarchy by quotienting everything by G.

Definition 4.1.10. *The* depth *of a vertex* v *in a hierarchy* $\mathcal{H}$*, denoted by* $\operatorname{depth}(v)$*, is its distance from the initial vertex. The* depth *of a hierarchy* $\mathcal{H}$*, denoted by* $\operatorname{depth}(\mathcal{H})$*, is the supremum of the depths of its vertices.*

The n^{th}–level *of* $\mathcal{H}$*, denoted by* $\mathcal{H}^n$*, is the set of all the vertices of* $\mathcal{H}$ *with depth* n*.*

We say that $\mathcal{H}$ *is* finite *if* $\operatorname{depth}(\mathcal{H})$ *is finite and the graph of groups associated to each action is finite. Equivalently* $\mathcal{H}$ *has finitely many* G*–orbits of vertices.*

Definition 4.1.11. *A* JSJ-hierarchy *(over* $\mathcal{A}$ *relative to* $\mathcal{B}$*) is a hierarchy where the action associated to every vertex is on a JSJ-tree (over* $\mathcal{A}$ *relative to* $\mathcal{B}$*.) If the group associated to a vertex is in* $\mathcal{A}$ *then we insist that it is terminal.*

Unless mentioned otherwise all JSJ-hierarchies will be non-relative and over the class of slender subgroups. Also note that the JSJ-trees *don't* need be canonical (such the Bowditch JSJ-tree [5]) and can instead be any maximal splitting.

In general the vertex groups of a splitting need not be finitely presented even if the original group was. However for hyperbolic groups we have the following.

Proposition 4.1.12. *[5, Proposition 1.2] If* G *is a hyperbolic group, then any graph of groups decomposition of* G *with quasiconvex edge groups also has quasiconvex vertex groups.*

Theorem 4.1.13 (Tits alternative)**.** *[12, Theorem 37] Every subgroup of a hyperbolic group* G *is either virtually cyclic or contains* F_2 *as a subgroup.*

In particular every slender subgroup of a hyperbolic group is virtually cyclic and hence quasi-isometrically embeds into it. Thus since every virtually cyclic subgroup of a hyperbolic group embeds quasiconvexly we see that every vertex group of a slender splitting for a hyperbolic group is also hyperbolic. So JSJ-hierarchies always exist for hyperbolic groups.

Subgroups which are elliptic on each level of the hierarchy will play an important role.

Definition 4.1.14. *A subgroup of $H \leqslant G$ is $\mathcal{H}$–elliptic if it's either contained in a terminal vertex of $\mathcal{H}$ or there is an infinite sequence of vertices $\{v_n\}$ with $H \leqslant G_{v_n}$ for all n and where each v_{n+1} is an immediate descendent of v_n.*

Suppose that $H \leqslant G$ is non-slender. Suppose that $H \leqslant G_v$ for some $v \in \mathcal{H}$. Since T_v is a slender tree H can be contained in at most one vertex group of T_v. Thus by iterating we see that H is contained in at most one vertex at each level of $\mathcal{H}$. So if H is also $\mathcal{H}$-elliptic then it is contained in exactly one vertex at each level of $\mathcal{H}$ and so if $H \leqslant G_v$ then it acts elliptically on T_v.

Lemma 4.1.15. *Suppose that $\mathcal{H}$ is a JSJ-hierarchy over slender edge groups (relative to $\mathcal{B}$) for a group G and that $\mathcal{K}$ is a finite slender hierarchy also for G which has terminal vertex groups which are either slender or $\mathcal{H}$-elliptic. Then $\mathcal{H}$ is also finite and moreover*

$$\operatorname{depth}\mathcal{H} \leq \operatorname{depth}\mathcal{K} + 1$$

We'll save the proof of this for Section 4.4 as it fits in naturally with what we are doing there.

The bulk of this chapter will be spent trying to massage hierarchies until they satisfy the conditions of Theorem 4.2.1. As a way of measuring how far away from doing this we are we introduce the following notions.

Definition 4.1.16. *A $\mathcal{H}$–*complex *X is a 2–dimensional connected simplical complex with $H^1(X,\mathbb{Z}_2) = 0$ and which some $K \leqslant G$ acts on with cell stabilisers which are either slender or $\mathcal{H}$–elliptic.*

The covolume *of a $\mathcal{H}$–complex X is the number of orbits of triangles under the action. Denote this quantity as* $\operatorname{covol}(X)$.

*A $\mathcal{H}$-*structure *$\mathcal{K}$ is a finite slender hierarchy for a group $K \leqslant G$ together with a $\mathcal{H}$–complex X_w acted upon by G_w for each terminal vertex $w \in \mathcal{K}$.*

The covolume *of a $\mathcal{H}$–structure $\mathcal{K}$ is the sum of the covolumes of the complexes associated to its terminal vertices (modulo equivalence.) Denote this quantity as* $\operatorname{covol}(\mathcal{K})$.

In many important cases (such as for hyperbolic groups) we will be able to take all of the cell stabilisers of our $\mathcal{H}$–complexes to be slender. This will streamline a few parts of the argument. Adding $\mathcal{H}$–elliptic cell stabilisers is necessary if we wish to consider certain other applications such as relatively hyperbolic groups.

4.2 Main results

The focus of this chapter shall be proving the machinery necessary for showing the following fact to be true.

Theorem 4.2.1. *[19, Corollary 2.7] Let G be a hyperbolic group which is virtually 2–torsion-free. Then any JSJ-hierarchy for G is finite.*

Note in particular that this implies that residually finite hyperbolic groups are strongly accessible. We will do this by proving the main result of Louder and Touikan [19] in full generality, keeping the above goal in mind so as to keep us from getting too lost in the details. Before we can state this result we need a few technical definitions.

Definition 4.2.2. *Let G be a finitely generated group and let $\mathcal{H}$ be a slender hierarchy of G. We say that G is $\mathcal{H}$–*almost finitely presented *if it has a $\mathcal{H}$–structure with finite covolume.*

Note that G being (almost) finitely presented means that it acts freely and cocompactly on some 2–dimensional simplicial complex X with $H^1(X,\mathbb{Z}_2) = 0$. Thus this notion genuinely extends the notion of being (almost) finitely presented.

The following restriction is rather technical and is essentially defined to be the exact condition which causes a step deep in the argument to work.

Definition 4.2.3. *$\mathcal{H}$ satisfies the* ascending chain condition *(henceforth abbreviated to ACC) if every ascending chain of subgroups*

$$S_1 \leqslant S_2 \leqslant S_3 \leqslant \cdots ,$$

where the following conditions are satisfied, eventually stabilises. (i.e. there exists N such that $S_N = S_i$ for all $i \geq N$.)

- *$S_i \leqslant G_{v_i}$ for some vertex v_i of $\mathcal{H}$. Moreover S_i is a subgroup of a non $\mathcal{H}$-elliptic edge group in the action on T_{v_i}.*
- *Each v_{i+1} is a descendent (although not necessarily an immediate one) of v_i.*
- *S_i is $\mathcal{H}$-elliptic.*

This abstract condition is satisfied if an ACC on the slender subgroups of G holds. i.e. it is enough to say that every ascending chain of slender groups of G must stabilise. This is a more restrictive condition, however is much easier to internalise and is satisfied for hyperbolic groups. Indeed, given an ascending chain of slender subgroups $\{S_i\}$, the Tits alternative (Theorem 4.1.13) implies that $S_\infty = \bigcup_i S_i$ is either finite or virtually $\mathbb{Z}$. Thus every infinite S_i has finite index in S_∞ and so the ACC follows for free in this case.

For the argument to work we can only consider elliptic and linear actions. Hyperbolic and parabolic actions are excluded by the fact we are working with slender groups, however we still need to prohibit dihedral actions. The following definition allows us to classify when we can do this.

Definition 4.2.4. *A slender hierarchy $\mathcal{H}$ is said to be* linear *if whenever E is an edge group of some T_v with $v \in \mathcal{H}$ we have $E \cap G_w$ acting either elliptically or linearly on T_w for any $w \in \mathcal{H}$.*

*A group G is said to admit a D_∞–*quotient *(over the class $\mathcal{A}$) if there exists a subgroup $H \leqslant G$ which surjects onto D_∞ (with a kernel in $\mathcal{A}$.) Observe that if a group admits no D_∞–quotients over slender groups then all its slender hierarchies are linear.*

Suppose we have subgroup $D \cong A *_C B$ of a hyperbolic group G which surjects D_∞ with slender kernel C. Observe that D is slender, so the Tits alternative (Theorem 4.1.13) now implies that D is virtually cyclic. Thus we see that A, B and C are all finite and in particular both A and B must contain group elements of order 2. Thus a 2–torsion-free hyperbolic group doesn't admit any D_∞–quotients over its slender subgroups.

We are finally ready to state the main result in full. Note that $\mathcal{H}$ doesn't need to be a JSJ-hierarchy for the following to hold.

Theorem 4.2.5. *[19, Theorem 2.5] Let G be a group and let $\mathcal{H}$ be a linear slender hierarchy for G. Suppose that G is finitely presented (relative to $\mathcal{B}$) and that $\mathcal{H}$ satisfies the ACC. Then there exist N and C such that for every vertex v in $\mathcal{H}$ with depth at least N there exists a finite hierarchy $\mathcal{K}_v$ for G_v with* $\mathrm{depth}(\mathcal{K}_v) \leq C$.

Once we prove this our goal follows swiftly.

Proof of Theorem 4.2.1. First consider the case where G is 2–torsion-free. Our prior discussions show that in this case the conditions for Theorem 4.2.5 are satisfied and so its conclusions follow. Now Lemma 4.1.15 tells us that $\mathcal{H}$ is finite with $\mathrm{depth}(\mathcal{H}) \leq N + C + 1$.

Now suppose G is virtually 2–torsion-free. Claim that if G_0 is a finite index subgroup of G then non-slender vertex groups of a JSJ-tree of G are finite index supergroups of the non-slender vertex groups of a JSJ-tree of G_0. From this we see that a JSJ-hierarchy for

G has a non-slender group at level n if and only if the same holds for a JSJ-hierarchy of G_0; at which point the 2–torsion-free case implies that we are done. It remains to prove the claim. Guirardel and Levitt [14, Proposition 4.16] tell us that we can obtain a JSJ-splitting for G by taking a maximal splitting over finite edge groups and replacing each vertex with a JSJ-splitting for each of the (one-ended) vertex groups. In particular we can take the JSJ-splittings for the one-ended groups to be the canonical Bowditch JSJ-splitting. [5] Now since $G_0 \leqslant_{\text{f.i.}} G$ they act freely and cocompactly on a common geodesic metric space; meaning that if T is a maximal tree over finite edge groups for G then T_{G_0} is the same for G_0. (Possibly after reducing away some slender vertices.) Likewise the Bowditch JSJ depends only on the topology of boundary of the underlying group, hence is invariant under quasi-isometries. Thus we can build a JSJ-tree for G whose restriction to G_0 (after reducing) is also a JSJ-tree. The claim now follows. □

The same argument as the above proof of Theorem 4.2.1 can also be used to prove the following more general statement.

Theorem 4.2.6. *[19, Corollary 2.6] Let G be a group and let $\mathcal{H}$ be a slender JSJ-hierarchy (relative to a class $\mathcal{B}$) for G. Suppose that G is $\mathcal{H}$–almost finitely presented and doesn't admit any D_∞–quotients. Suppose also that $\mathcal{H}$ satisfies the ACC. Then $\mathcal{H}$ is finite.*

Remark 4.2.7. *There is a separate notion of a subgroup $K \leqslant G$ being slender relative to a set of subgroups $\mathcal{B}$; satisfied if K fixes a point or a line on any tree T which G acts on where each group in $\mathcal{B}$ fixes a vertex of T. We do not consider such subgroups here and bring them up only to clear up any confusion between them and the notion of a slender JSJ relative to $\mathcal{B}$.*

Remark 4.2.8. *While it is possible to extract an explicit value for C in terms of G from the upcoming proof of Theorem* 4.2.5, *(assuming that $\mathcal{H}$ is a JSJ-hierarchy,) there is no immediately obvious way to do the same for N. Thus this argument does not give an explicit bound on the height of $\mathcal{H}$.*

4.3 A note on canonicalness of JSJ-hierarchies

JSJ-trees are not in general unique, however they share a common deformation space. Thus every vertex group of a JSJ-tree is either slender or is a vertex group in any other JSJ-tree for the same group [14, pg.6]. Thus we see that most of the resulting hierarchy is identical regardless of our choices of JSJ-trees. In particular this shows that the depth of a JSJ-hierarchy is constant for a given group.

If one still feels the need to consider a more canonical object then we can define one as follows. For one ended hyperbolic groups there is a canonical JSJ-tree, called the Bowditch JSJ-tree, which is invariant under outer automorphisms. (See [5] or [15] for details.) Moreover we can use Dunwoody accessibility [11] (which are exactly JSJ-splittings over finite groups) to split multi-ended groups into one-ended ones and the resulting vertex groups to not depend on the exact choice of tree. Alternating between these we can thus get a hierarchy which is defined in an essentially canonical way; meaning that the vertices and their associated groups of the hierarchy are always the same. One can trivially modify the previous proof of Theorem 4.2.1 to work for such hierarchies, to see that these hierarchies are also finite given that the underlying group is virtually without 2–torsion.

4.4 Passing hierarchies to subgroups

The key to proving Theorem 4.2.5 will be to successively pass some 'bad' auxiliary hierarchies from one level of $\mathcal{H}$ to the next until eventually they become 'good'. We will thus begin by detailing a method for splitting a hierarchy over an unrelated tree.

Lemma 4.4.1. *Let T be a tree which a group G acts on with slender edge stabilisers and let $\mathcal{K}$ be a finite slender hierarchy for G. Then for each vertex $v \in T$ there is a finite slender hierarchy $\mathcal{K}_v$ for the vertex group G_v with the following properties.*

1. $\mathrm{depth}(\mathcal{K}_v) \leq \mathrm{depth}(\mathcal{K})$ *for any $v \in T$.*
2. *For each vertex $w \in \mathcal{K}_v$ the group G_w is a subgroup of some G_u where $u \in \mathcal{K}$ and* $\mathrm{depth}(w) \leq \mathrm{depth}(u)$.
3. *If $\mathcal{K}$ is non-trivial and T is a JSJ-tree then* $\mathrm{depth}(\mathcal{K}_v) < \mathrm{depth}(\mathcal{K})$ *whenever v is a rigid vertex of T.*
4. *Suppose that the terminal vertices of $\mathcal{K}$ have associated groups which are either slender or act elliptically in T. Then the terminal vertices of $\mathcal{K}_v$ have associated groups which are either slender or equal to a terminal group of $\mathcal{K}$.*

Later we will take T to be a tree in the hierarchy $\mathcal{H}$ and $\mathcal{K}$ to be one of the aforementioned auxiliary hierarchies. Louder and Touikan use the symmetric core of a product of trees to produce the $\mathcal{K}_v$; however this is not necessary as a more direct approach also works.

Proof. For each $v \in T$ we define $\mathcal{K}_v$ with property 2 listed above one level at a time. By definition the initial vertex of $\mathcal{K}_v$ has associated group G_v which trivially satisfies the required property. So we just need a procedure to generate the action on a tree for a given vertex of $\mathcal{K}_v$.

Take a vertex $w \in \mathcal{K}_v$ with associated group G_w. We are given that $G_w \leqslant G_u$ where $u \in \mathcal{K}$ and $\text{depth}(w) \leq \text{depth}(u)$. If u is a terminal vertex then we take T_w to be a point and so w is also terminal. Otherwise without loss of generality we may assume that G_w is non-elliptic in T_u by passing to vertex groups. We now just take T_w to be the unique minimal subtree of T_u which is invariant under G_w. (This exists because the action of G_w on T_u doesn't contain any global fixed points.) By definition the edge groups of T_w are slender and the vertex groups of T_w are subgroups of the vertex groups of T_G. The latter implies property 2 by induction, thus completing the construction.

We now prove the remaining properties. Property 1 is just a weaker version of property 2 and so is immediately satisfied. Property 3 holds because each rigid vertex of T is by definition elliptic in the top level of $\mathcal{K}$ and so in fact for every $w \in \mathcal{K}_v$ we have $u \in \mathcal{K}$ with $G_w \leqslant G_u$ and $\text{depth}(w) < \text{depth}(u)$.

For property 4 consider a terminal vertex $w \in \mathcal{K}_v$ where G_w is not slender. As usual consider a vertex $u \in \mathcal{K}$ with $G_w \leqslant G_u$ and observe that we can take u to be a terminal vertex of $\mathcal{K}$. Suppose that G_u fixes a line of T. Then $G_u \cup G_v$ either is or has an index 2 subgroup which is contained in an edge stabiliser of T. Hence $G_w \leqslant G_u \cap G_v$ is slender which contradicts our assumption. Thus G_u fixes some vertex $v' \in T$ and so we have $G_u \leqslant G_{v'}$. Since the edges of T have slender stabiliser, $G_w \leqslant G_v \cap G_{v'}$ and G_w is not slender we must have $v = v'$. Since G_u is $\mathcal{K}$–elliptic and $G_u \leqslant G_v$ we see that G_u is also $\mathcal{K}_v$–elliptic from the construction of $\mathcal{K}_v$. Recall that since neither G_u nor G_w are slender they are both contained in exactly one vertex for each level of $\mathcal{K}_v$. Moreover since $G_w \leqslant G_u$ they must both be contained in the same vertices. Hence $G_u \leqslant G_w$ by the definition of G_w and so $G_u = G_w$. □

Remark 4.4.2. *Essentially the same argument as the last part shows that if $\{w_i\}$ is a collection of distinct terminal vertices of $\coprod_v \mathcal{K}_v$ with $G_{w_i} \leqslant G_u$ where $u \in \mathcal{K}$ is a terminal vertex; then at most one of the w_i can be non-slender and $G_{w_i} = G_u$ for this i. This will be important shortly when extending the construction to $\mathcal{H}$–structures.*

We now have all the tools needed to prove Lemma 4.1.15.

Proof of Lemma 4.1.15. We prove by induction on $\text{depth}\,\mathcal{K}$. If $\text{depth}\,\mathcal{K} = 0$ then $\mathcal{K}$ is trivial and so G is either slender or $\mathcal{H}$–elliptic; hence it cannot split in $\mathcal{H}$ and so $\mathcal{H}$ is trivial. Otherwise consider each vertex $v \in \mathcal{H}^1$ in turn. Let $\mathcal{H}_v$ be the subhierarchy of $\mathcal{H}$ with initial vertex v. If G_v is a rigid group in the action on the tree corresponding to the initial vertex of

$\mathcal{H}$ then Lemma 4.4.1 implies that we have another hierarchy $\mathcal{K}_v$ for G_v with depth$(\mathcal{K}_v) <$ depth$(\mathcal{K})$ and with terminal vertices which have associated groups that are either slender or $\mathcal{H}$–elliptic. Thus depth$(\mathcal{H}_v) \leq$ depth$\mathcal{K}$ by induction. If G_v is a flexible group in the action of top level of $\mathcal{H}$ then Lemma 4.1.7 implies that G_v is slender by orbifold; which implies that depth$(\mathcal{H}_v) \leq 1 \leq$ depth$(\mathcal{K})$. Thus in any case we have depth$(\mathcal{H}_v) \leq$ depth$(\mathcal{K})$ for all $v \in \mathcal{H}^1$ and hence depth$(\mathcal{H}) \leq$ depth$(\mathcal{K}) + 1$. □

We will measure how 'bad' our auxiliary hierarchies are by introducing some actions on some complexes. The following lemma shows us that these actions pass down nicely if the terminal vertices of our initial auxiliary hierarchy are elliptic in our tree.

Lemma 4.4.3. *Let $\mathcal{K}$ be a $\mathcal{H}$–structure for $K \leqslant G$. Suppose K acts on a tree T with slender edge stabilisers. Suppose that the terminal vertices of $\mathcal{K}$ are either slender or elliptic in T. Then for each vertex $v \in T$ we get that $\mathcal{K}_v$ (as defined in Lemma* 4.4.1*) naturally inherits a $\mathcal{H}$–structure from $\mathcal{K}$ and moreover we have*

$$\sum_i \mathrm{covol}(\mathcal{K}_{v_i}) = \mathrm{covol}(\mathcal{K})$$

where $\{v_i\}$ is a set of representatives for the orbits of vertices in T.

Proof. Recall from Lemma 4.4.1 that for every terminal vertex $w \in \mathcal{K}_v$ that K_w is either slender or equal to K_u for some terminal vertex $u \in \mathcal{K}$. Thus $\mathcal{K}_v$ naturally inherits a $\mathcal{H}$–structure by letting K_w act trivially on a point if it's slender or on the same complex as K_u otherwise.

It remains to prove the equality of covolumes. Let u be a terminal vertex of $\mathcal{K}$ and let $\{w_j\}$ be a set of representatives for the terminal vertices of $\coprod_i \mathcal{K}_{v_i}$ which have K_{w_j} conjugating into K_u. Observe that for each $w_j \in \mathcal{K}_{v_{i(j)}}$ we have a corresponding $w'_j \in \mathcal{K}_{g_j v_{j(i)}}$ with $(K_{w_j})^{g_j} = K_{w'_j} \leqslant K_u$. These w'_j are distinct as the w_j are in distinct orbits. Hence by Remark 4.4.2 at most one of the K_{w_j} can be non slender and $K_{w'_j} = K_u$ for this j. Such a j must exist as K_u is contained in some vertex group $G_{v'} \leq G$ and moreover is $\mathcal{K}_{v'}$–elliptic. Thus there is a natural G–equivariant one to one correspondence between the non-slender terminal vertices of $\mathcal{K}$ and $\coprod_{v \in T} \mathcal{K}_v$ which implies the result. □

Of course in general it won't be the case that the terminal vertices of $\mathcal{K}$ will be elliptic in T. So our next goal shall now be to add additional layers to $\mathcal{K}$ so this becomes true while not increasing the covolume. We thus require some sort of resolution in order to nicely split up our complexes. If every cell stabiliser of X acted elliptically on T then we could get this from an equivariant map $X \to T$. Since the cell stabilisers can also act linearly on T such a map needn't exist and so we need to be more careful. We will modify X in order to isolate these bad points.

After making modifications to the complexes we may find that they fail to be simplical. For example if we collapse one edge of a triangle we are left with a bigon. We will get around this by *reducing* complexes as follows. Let X be a 2–dimensional cell complex where all the 2–cells are either triangles or bigons. Start defining a simplical complex X' by letting the vertex set be the same as X. Let $[u,v]$ be an edge of X' if there is an edge between u and v in X and similarly let $[u,v,w]$ be a triangle in X' if there is a triangle in X with vertices u, v and w. This X' is the reduction of X. Note that if X is connected with $H^1(X,\mathbb{Z}_2) = 0$ then the same holds for X'.

In order to split the complex we will make use of the following Dunwoody-Delzant-Potyagailo resolution.

Lemma 4.4.4. *[19, Lemma 3.5] Let G be a group acting on a triangle complex X and a tree T. Suppose that*

- *the cell stabilisers of X fix a point of $\hat{T} := T \cup \partial T$. (Where ∂T is the* Gromov boundary *of T.) i.e. they all act elliptically, linearly or parabolically on T.*
- *if $W \subset X^1$ is a connected subcomplex where the stabiliser of each edge acts linearly or parabolically on T then the stabiliser of W fixes a point on ∂T.*

Then there is a resolution $\rho : X \to \hat{T}$.

Before constructing ρ we will show that the second condition is always satisfied for our purposes.

Lemma 4.4.5. *Let G be a group acting on a tree T and let $K \leq G$ act on a $\mathcal{H}$–complex X. Let $W \subset X^1$ be a connected subcomplex such that the stabiliser of every cell in W acts linearly or dihedrally on T. Then* $\mathrm{Stab}(W)$ *also acts linearly or dihedrally on T. Moreover if G doesn't admit any D_∞-quotients then the action of* $\mathrm{Stab}(W)$ *on T is linear.*

Proof. Suppose that $E \leq V$ are subgroups of G which both act linearly or dihedrally on T. Then V must fix the same line as E as otherwise V would contain hyperbolic group elements with different axes, contradicting the fact that V fixes a line of T. It follows that every cell in W fixes a common line in T which implies the result. □

Proof of Lemma 4.4.4. First for each maximal subcomplex W as in the second condition of the statement we (equivariantly) choose a point in ∂T which it fixes.

Next we need to (equivariantily) map each vertex v of X to either a vertex of T or a point on ∂T. If a vertex is contained in a subcomplex W as above we define $\rho(v)$ to be equal to the point on ∂T corresponding to W. Otherwise we just define $\rho(v)$ to be any vertex of T which $\mathrm{Stab}(v)$ fixes.

Now let $e = [u, v]$ be an edge of X. If $\rho(u) = \rho(v)$ then we can just take ρ to be constant on e. Otherwise we want $\rho(e)$ to be the reduced edge path from $\rho(u)$ to $\rho(v)$; however we need to be careful with the parametrisation if $\mathrm{Stab}(e) \neq \mathrm{Stab}^+(e)$. If $\rho(u)$ and $\rho(v)$ are in T we can just take the parametrisation to be linear as the midpoint of $[\rho(u), \rho(v)]$ is fixed by $\mathrm{Stab}(e)$. If $\rho(u)$ is in T but $\rho(v)$ isn't then u and v are in different orbits and so $\mathrm{Stab}(e) = \mathrm{Stab}^+(e)$. Finally suppose $\rho(u)$ and $\rho(v)$ are in ∂T. By assumption $\mathrm{Stab}(e)$ fixes a vertex $v \in T$. Let $y \in [\rho(u), \rho(v)]$ be the closest vertex to x. Since $\mathrm{Stab}(e)$ fixes x and preserves $[\rho(u), \rho(v)]$ it must also fix y. Now map the midpoint of e to y and map the rest in any way which is symmetric through y.

Extending the map affinely over triangles works for similar reasons. □

There are two different cases we will consider, depending on if an edge gets mapped to ∂T or not. If $\rho^{-1}(\partial T)$ doesn't contain any edges then we say that this is a resolution of type I or a *splitting resolution*. Otherwise $\rho^{-1}(\partial T)$ contains an edge and we call this a type II or a *contracting resolution*. The splitting resolutions will allow us to modify the hierarchy so that its terminal vertices are elliptic in T. The contracting resolutions are an issue but we will modify them so that they become splitting resolutions.

First we will describe what to do in the case of a splitting resolution. Let $\rho : X \to \hat{T}$ be as above. Let $\Lambda \subset X$ be the inverse images of the midpoints of the edges in T and observe that this is a collection of tracks (in the sense of Dunwoody [11]). Let $\Lambda^* \subset \Lambda$ be the tracks which partition X into two infinite parts and let $X^* := X \backslash \rho^{-1}(\partial T)$. Observe that X^*/Λ^* (obtained by collapsing each track in Λ^* to a point) is a 2-dimensional cell complex where all the 2–cells are either bigons or triangles. Finally X_T is defined to be the reduction of X^*/Λ^*. Observe that the image of each triangle of X in X_T contains at most one triangle and so $\mathrm{covol}(X_T) \leq \mathrm{covol}(X)$. (See Figure 4.1)

Remark 4.4.6. *Unlike with triangles, the image of an edge of X in X_T need not be a single edge and is in general a (potentially infinite) sequence of edges. However suppose we are*

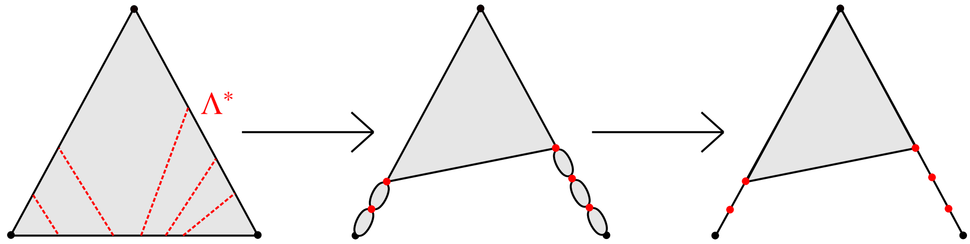

Fig. 4.1 An example of the effects a splitting resolution has on a triangle.

given a triangle in X whose image in X_T contains a triangle. Then the edges of this new triangle in X_T will be a single edge in the image of a corresponding edge of the original triangle.

Lemma 4.4.7. *Suppose that X is connected, has $H^1(X,\mathbb{Z}_2) = 0$ and $\rho : X \to T$ is a splitting resolution. Suppose also that every cutpoint of X has a stabiliser which fixes a point of T. Then X_T is connected with $H^1(X_T,\mathbb{Z}_2) = 0$.*

Proof. Since every cutpoint of X acts elliptically on T we see that X^* is connected and therefore X_T is as well. Since each track is connected it now suffices to show that each cycle in $H^1(X^*,\mathbb{Z}_2)$ is a boundary when mapped into X_T.

Let B be the second barycentric subdivision of X. Let $C \subset B$ be the union of simplices which are disjoint from $\rho^{-1}(\partial T)$ and $A \subset B$ be the union of the simplices (and their subsimplices) which intersect non-trivially with $\rho^{-1}(\partial T)$. Also let $L = A \cap C$.

Consider the Mayer-Vietoris sequence for A and C;

$$\cdots \to H^1(X,\mathbb{Z}_2) \to H^1(A,\mathbb{Z}_2) \oplus H^1(C,\mathbb{Z}_2) \to H^1(L,\mathbb{Z}_2) \to \cdots$$

Since B is the second barycentric subdivision of X we see that the stars of two distinct vertices which are in $\rho^{-1}(\partial T)$ are disjoint. Hence each component of A is the star of a point and so $H^1(A,\mathbb{Z}_2) \cong 0$. Similarly we see that $C \subset X^*$ is a deformation retract. Hence the sequence becomes

$$0 \to H^1(X^*,\mathbb{Z}_2) \to H^1(L,\mathbb{Z}_2) \to \cdots$$

In particular it suffices to show that each loop of L dies when we pass to X_T.

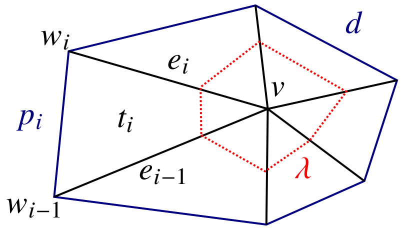

Let d be a reduced edge path of L. There is $v \in \rho^{-1}(\partial T)$ with d homotopic to an edge path $p_1 \cdots p_n$ in the link of v. Let w_i be the vertex common to both p_i and p_{i+1} (where the indices are taken modulo n). Also let e_i be the edge connecting w_i to v and t_i be the triangle with edges w_{i-1}, w_i and v. Since $\rho(w_i) \neq \rho(v)$ for any i and $\rho(v) \in \partial T$ we can choose an edge $f \in T$ such that $f \in [\rho(w_i), \rho(v)]$ and $f \cap \rho(p_i) = \emptyset$ for all i. Thus there is a track λ (which maps to the midpoint of f) whose intersection with $t_1 \cup \cdots \cup t_n$ is homotopic to d. This implies that d is null-homotopic in X_T which implies the result. □

We will now detail what to do in the case of a contracting resolution. Recall that this is the case where $\rho^{-1}(\partial T)$ contains an edge of X. Define X_C to be the complex obtained by collapsing each component of $\rho^{-1}(\partial T) \subset X$ to a point. We summarise the properties of X_C in the following lemma.

Lemma 4.4.8. *Notation as before. Suppose that the cell stabilisers of X are all either slender or elliptic in T. Then every vertex stabiliser of X_C is either equal to a vertex stabiliser of X or fixes a line of T and hence is slender. Moreover ρ descends to a map $X_C \to \hat{T}$ where the midpoint of each edge of X_C gets sent to a point in T and* $\mathrm{covol}(X_C) \leq \mathrm{covol}(X)$.

Proof. These properties are all obvious from the definition of X_C together with Lemma 4.4.5. □

This concludes our discussion on contracting resolutions. We also need a general method for splitting a complex over its cutpoints. Suppose X is a $\mathcal{H}$–complex for a subgroup $K \leqslant G$. Construct a bipartite tree B_X, called the *cutpoint tree*, with vertices which correspond to the cutpoint-free components of X and the cut vertices of X and with the edges of B_X defined in the obvious way by inclusion. Note that K acts on B_X with vertex stabilisers which are equal to the corresponding stabilisers in X and edge stabilisers which are equal to the stabiliser of a connected component of a link.

The above is fine if every cell of X has slender stabiliser, such as in the case of hyperbolic groups; however in general it may be the case that B_X has edge groups which are H–elliptic but not slender. To counteract this we introduce a *reduced cutpoint tree* B'_X by collapsing the edges of B_X which have non-slender stabiliser. Observe that B'_X can naturally be thought of as a $\mathcal{H}$–structure (of depth 1) with the properties summarised in the following lemma.

Lemma 4.4.9. *Let B'_X be the cutpoint tree defined above.*

(a) $\mathrm{covol}(B'_X) \leq \mathrm{covol}(X)$

(b) If X_T is the splitting resolution of some complex X where the cutpoints of X have non-slender stabilisers then the vertex groups of B'_{X_T} are either slender or elliptic in T.

Proof. For part (a) we first observe that each triangle of X sits in at most one subcomplex corresponding to a non-slender vertex of B'_X. Thus we just need to know if two triangles t and $t' = gt$ (with $g \in K$) live in the same subcomplex $X_v \subseteq X$ corresponding to a vertex $v \in B'_X$ then they still lie in the same orbit when restricted to the vertex group K_v. This follows because if $g \in K$ sends a triangle of X_v to another in X_v then it must preserve X_v, so $g \in K_v$.

For part (b) we first observe that the stabiliser of each connected component $Y \subseteq X^* \backslash \Lambda^*$ is elliptic in T by construction. The image of $\bar{Y}$ in X_T is a maximal subcomplex Y' which doesn't contain any cutpoints with slender stabiliser. Moreover the stabiliser of Y is the same as the stabiliser of Y'. A subcomplex of X_T corresponding to a non-slender vertex of B'_{X_T} is contained in such a Y'. Hence the non-slender vertices of B'_{X_T} are elliptic in T. □

Putting all of this together we obtain the following.

Lemma 4.4.10. *[19, Lemma 7.2] Suppose G does not admit any D_∞–quotients. Let T be a tree which $K \leqslant G$ acts on with slender edge stabilisers and let $\mathcal{K}$ be a $\mathcal{H}$–structure for K. Then for each vertex $v \in T$ there is a $\mathcal{H}$–structure $\mathcal{K}_v$ with*

$$\sum_i \mathrm{covol}(\mathcal{K}_{v_i}) \leq \mathrm{covol}(\mathcal{K})$$

where $\{v_i\}$ is a set of representatives for the orbits of vertices of T.

Proof. In light of Lemma 4.4.3 it suffices to show that there is another $\mathcal{H}$–structure $\tilde{\mathcal{K}}$ with terminal vertices which are elliptic in T and with $\mathrm{covol}(\tilde{\mathcal{K}}) \leq \mathrm{covol}(\mathcal{K})$.

Begin by considering the resolution of each complex associated to the terminal vertices of $\mathcal{K}$. We define $\mathcal{K}'$ to be the same as $\mathcal{K}$ except whenever a complex X associated to a terminal vertex has an edge mapped into ∂T we replace it with X_C defined above. Recall from Lemma 4.4.8 that X_C is a $\mathcal{H}$–complex and $\text{covol}(X_C) \leq \text{covol}(X)$. Thus $\mathcal{K}'$ is a $\mathcal{H}$–structure with $\text{covol}(\mathcal{K}') \leq \text{covol}(\mathcal{K})$. Moreover the resolutions from each complex are now all splitting resolutions.

Before collapsing tracks it is first necessary to split the complexes over cutpoints. (Otherwise our X_T may not be connected.) Following the procedure from Lemma 4.4.9 for each complex of $\mathcal{K}'$ containing cutpoints we construct a new $\mathcal{H}$–structure $\mathcal{K}''$ where all the cutpoints in the complexes have $\mathcal{H}$–elliptic stabilisers. Moreover the resolutions still map each edge of the complexes to a point of T and $\text{covol}(\mathcal{K}'') \leq \text{covol}(\mathcal{K}')$.

Finally for each complex X of $\mathcal{K}''$ we consider X_T as defined above. Recall from Lemma 4.4.9 that each non-slender vertex of B'_{X_T} acts elliptically on T. Thus we get a new $\mathcal{H}$–structure $\tilde{\mathcal{K}}$ by replacing each terminal vertex of $\mathcal{K}''$ with the corresponding B'_{X_T}. This $\tilde{\mathcal{K}}$ has all the properties we require. □

Remark 4.4.11. *Note that there is a natural partial map from the set of triangles in the complexes of $\mathcal{K}$ and those in $\coprod_{v\in T}\mathcal{K}_v$. Moreover this map is G–equivariant and is both total and bijective if $\sum_i \text{covol}(\mathcal{K}_{v_i}) = \text{covol}(\mathcal{K})$ where $\{v_i\}$ is a set of representatives for the orbits of vertices of T. An understanding of this map will be crucial for Section 4.5.*

4.5 Extracting trees from complexes

Now with Lemma 4.4.10 in hand we are ready to start the proof of Theorem 4.2.5. Let v_0 be the initial vertex of $\mathcal{H}$. Start by letting $\mathcal{K}_{v_0}$ be any $\mathcal{H}$–structure for G with finite covolume. (Recall that for a finitely presented group we can take $\mathcal{K}_{v_0}$ to have trivial tree structure and have G act freely on a cocompct (2–dimensional) simply connected simplical complex.) Now we recursively define $\mathcal{K}_w$ for each vertex $w \in \mathcal{H}$. Suppose w' is the immediate ancestor of w and $\mathcal{K}_{w'}$ is already defined. We now define $\mathcal{K}_w$ from Lemma 4.4.10 by setting $\mathcal{K}$ to be $\mathcal{K}_{w'}$ and T to be $T_{w'}$.

Let $\mathcal{T}^n$ be the set of all the triangles in all the complexes acted on by the terminal vertices of $\mathcal{K}_w$ where $w \in \mathcal{H}^n$. Note that G naturally acts on $\mathcal{T}^n$ with finitely many orbits of triangles; call this number $\text{covol}(\mathcal{T}^n)$. Moreover the inequality of covolumes in Lemma 4.4.10 extends to an inequality $\text{covol}(\mathcal{T}^{n+1}) \leq \text{covol}(\mathcal{T}^n)$ for all n. Thus $\text{covol}(\mathcal{T}^n)$ must eventually reach some minimum. Pick N_Δ so that $\text{covol}(\mathcal{T}^{N_\Delta}) = \text{covol}(\mathcal{T}^n)$ for any $n \geq N_\Delta$.

Recall from Remark 4.4.11 that for $n \geq N_\Delta$ we can always pass a triangle to the next level of $\mathcal{H}$. More precisely our construction above actually induces a G–equivariant bijective map $\tau_{n,n+1} : \mathcal{T}^n \to \mathcal{T}^{n+1}$. Moreover let $\tau_{n,m} : \mathcal{T}^n \to \mathcal{T}^m$ be the composition $\tau_{m-1,m} \circ \cdots \circ \tau_{n,n+1}$.

A *pair* in $\mathcal{T}^n$ is an element $(t,t') \in \mathcal{T}^n \times \mathcal{T}^n$ where t and t' are distinct triangles in the same complex and which share a common edge e. A pair is called *stable* if it descends to a pair under any $\tau_{n,m}$ where $m > n$. Let $P(\mathcal{T}^n)$ be the set of stable pairs in $\mathcal{T}^n$.

We now define an equivalence relation $\sim_n$ on $\mathcal{T}^n$ to be the one generated by its stable pairs. Note that for each equivalence class of $\sim_n$ we naturally get a connected subcomplex (of some $\mathcal{H}$–complex which is associated to a terminal vertex of $\mathcal{K}_v$ for some $v \in \mathcal{H}^n$) consisting of all the triangles in the class together with all their subsimplices.

We now restrict our attention to a single complex X_w associated to a terminal vertex $w \in \mathcal{K}_v$ for some $v \in \mathcal{H}^n$ with $n \geq N_\Delta$. We define a bipartite graph B_w for each X_w as follows. One set of vertices will be the set of subcomplexes associated to the equivalence classes of $\sim_n$ which are contained in X_w; the other will be the edges of X_w which are contained in more than one of said subcomplexes. The edges of B_w are defined by inclusion in the obvious way.

Observe that G_w acts naturally on B_w. By definition the stabilisers for the subcomplexes associated to the equivalence classes of $\sim_n$ are $\mathcal{H}$–elliptic. If every edge of X_w has slender stabiliser (such as in the case for hyperbolic groups) then the stabilisers of each edge of B_w are slender. So if B_w is a tree for all large enough n then this proves Theorem 4.2.5 by adding the B_w to the bottom layers of the $\mathcal{K}_v$. (Where N in the statement of Theorem 4.2.5 is the first level where this occurs and the corresponding C is the maximal depth of one of the $\mathcal{K}_v$ where v has depth N in $\mathcal{H}$.) If some edge of X_w has a non-slender ($\mathcal{H}$–elliptic) stabiliser then we instead first have to collapse each edge of B_w with non-slender stabilser to get a new graph B'_w. Theorem 4.2.5 will then follow as before.

We shall now work backwards finding a series of sufficient conditions for B_w to be a tree until we arrive at one which we can show is true for large n. First observe that this is true if we can show that, for far enough down the hierarchy, whenever (t,t') is an unstable pair with common edge e that t and t' lie in different connected components of $X_w \backslash e$. We now need a definition.

Definition 4.5.1. *Let D be a triangulated disk with exactly one interior vertex. A* cone $C \subseteq X$ *is the image of some simplicial map $\alpha : D \to X$ which sends triangles to triangles. A cone is said to be* simple *if the image of $\partial\alpha : \partial S^1 \to X$ is a simple loop; equivalently α is injective.*

Lemma 4.5.2. *[19, Lemma 8.5] If every simple cone of X_w is contained in an equivalence class, then B_w is a tree.*

Before proving this we require a simple proposition.

Proposition 4.5.3. *Let γ be the boundary of some cone C. Suppose $e = [u,v]$ and $e' = [v,w]$ are consecutive edges of γ. If $u \neq w$ (so γ is locally injective at vertex v) then there a simple subcone $C' \subseteq C$ containing both e and e'.*

Proof. Suppose $\gamma : S^1 \to C$ is not simple. Then there are distinct $x_1, x_2 \in S^1$ which map to some common vertex $x \in C$. Let A be an arc of S^1 which starts at x_1, finishes at x_2 and which contains e and e' as consecutive edges. Let A' be the circle formed by taking A and gluing its endpoints together. Then $\gamma|_{A'}$ is the boundary for a proper subcone of C which contains e and e' as consecutive edges. Repeat this process until the resulting cone is simple, which must happen eventually as the area of the cone decreases at each step. □

Proof of Lemma 4.5.2. Let (t,t') be pair in X_w with common edge e. In order to prove the result it suffices to show that if t and t' are in the same connected component of $X_w \backslash e$ then $t \sim t'$. In this case $e = [u,v]$ is not a cut edge of X_w. Let a and b be the vertices of t and t' respectively that are not a part of e. Since e is not a cut edge there is an edge path γ (which we'll not assume is injective) from a to b which doesn't intersect e. Let l be the loop consisting of γ composed with $p = [a,u] \cup [u,b]$. Since X is simply connected there is a simply connected simplical complex $D \subset \mathbb{R}^2$ together with a simplical map $\rho : D \to X$ with boundary $\partial\rho : \partial D \to l$. Note that ρ is not required to be an embedding, even locally so. We will now assume that γ, D and ρ as above are chosen to lexicographically optimise the following quantities for which D is homeomorphic to a disc. (D is always homeomorphic to a disc if γ is injective; but this needn't be the case in general.)

- Minimises the number of triangles in D.
- Maximises the length of ∂D.

Note the length of ∂D is bounded above by three times the number of triangles of D. Thus we have a well ordering and so an optimal choice must exist.

For such optimal choices we get the following properties.

- Since ρ is a homeomorphism on each simplex we see that two disjoint components of $\rho^{-1}(e)$ must be separated by an edge path λ in D. Hence $\rho^{-1}(e)$ is connected in an optimal choice as otherwise we can 'cut across' λ to get a new loop which bounds strictly less area.

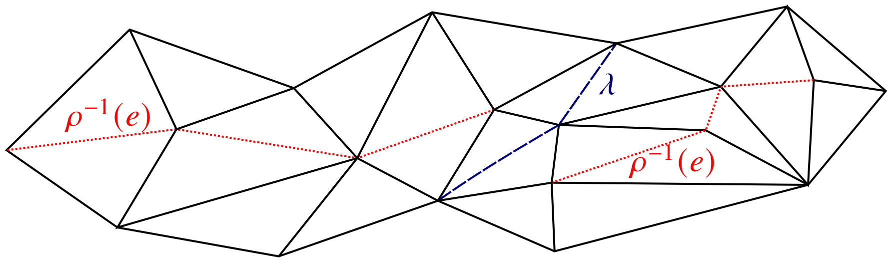

- Every edge $f \in \partial D$ must be in the link of a preimage of either u or v. Otherwise we could remove f and the unique triangle which contains f to obtain a new loop which bounds strictly less area.

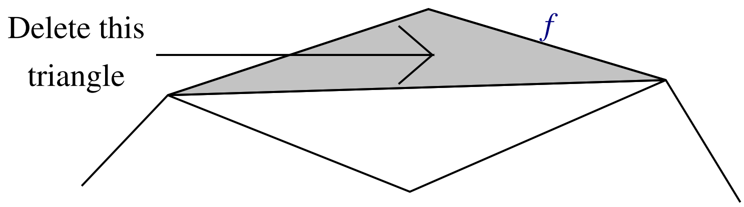

- The only non-boundary vertices in the link of a vertex w' of ∂D are preimages of u and v. Otherwise we could make ∂D longer by adding two copies of an interior edge of D to ∂D.

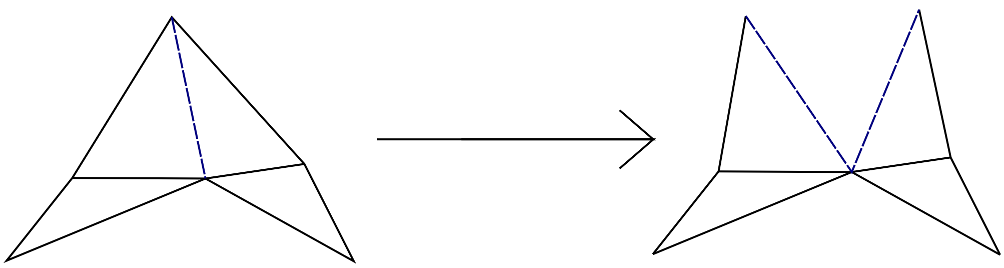

Combining all of the above we see that D must look like the following picture.

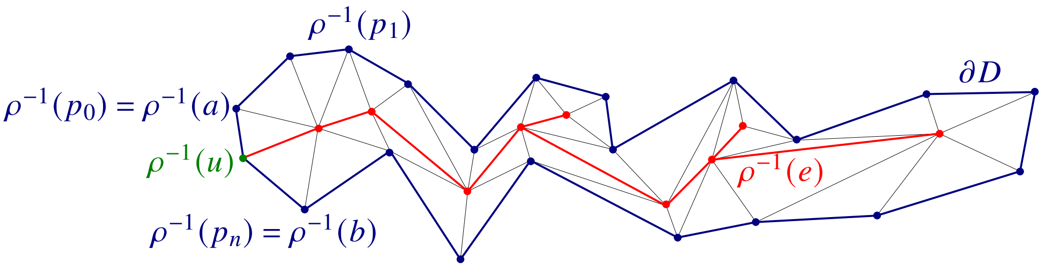

It follows that γ can be decomposed into locally injective subpaths γ_i between p_{i-1} and p_i with the following properties. $(1 \leq i \leq n)$

- Each γ_i is contained in the link of either u or v.
- For each i there is a triangle t_i with vertices p_i, u and v.

Thus for each i there is a cone with central point either u or v containing both t_{i-1} and t_i. Thus by Proposition 4.5.3 either $t_{i-1} = t_i$ or there is a simple cone containing both t_{i-1} and t_i. Thus by the assumption in the statement case we get $t_{i-1} \sim t_i$ for all i and so $t = t_0 \sim t_n = t'$. □

Let $\sigma_{n,m}$ be the surjective map induced by $\tau_{n,m}$ on the equivalence classes of $\sim_n$ and $\sim_m$. Our goal shall be to show that, far enough down the hierarchy, this $\sigma_{n,m}$ is always a bijection. Then we will show that the corresponding subcomplexes are themselves rigid which will allow us to prove Theorem 4.2.5.

Proposition 4.5.4. *For $n > N_\Delta$ every pair which is contained in the subcomplex associated to $\sim_n$ is a stable pair.*

Proof. First note that a subcomplex Y corresponding to an equivalence class must be cutpoint free as any two triangles it contains must be joined by a sequence of stable pairs. This means that the intersection of Y and any track from the resolution ρ is either trivial or parallel to a vertex of Y.

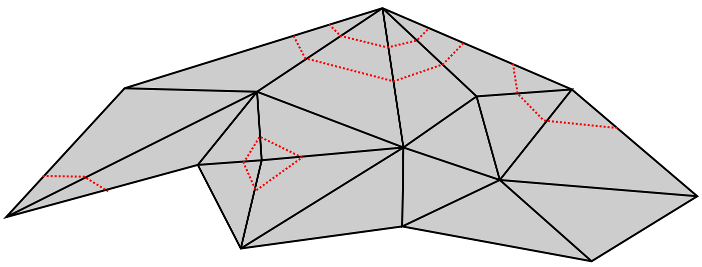

□

Lemma 4.5.5. *There exists $N' \geq N_\Delta$ such that $\sigma_{n,m}$ is a bijection whenever $n, m \geq N'$.*

Proof. We shall proceed by proving the following three claims about the structure of the classes of $\sim_n$.

Claim 1 There is some $N_1 \geq N_\Delta$ such that number of orbits of equivalence classes of $\sim_n$ and $\sim_m$ are equal whenever $n, m \geq N_1$.

Proof of Claim 1. Since $\tau_{n,m}$ induces a surjective equivariant map on the equivalence classes we see that the number of G-orbits of G classes is non-increasing, hence must be eventually constant. □

Let $Y_n^1, \cdots, Y_n^J$ be the associated subcomplexes to a set of representatives for the orbits of $\sim_n$ and without loss of generality we can assume that Y_n^j maps into Y_{n+1}^j.

Claim 2 There is some $N_2 \geq N_1$ such that the number of orbits of edges in each Y_n^j is constant for $n \geq N_2$.

Proof of Claim 2. Recall from Remark 4.4.6 that for any given triangle there is a natural correspondence between its edges at any given level. Thus the only way to increase the number of edges is if two triangles are adjacent on one level but then not on a later one. This contradicts Proposition 4.5.4. □

The proof of Claim 2 also means we can meaningfully talk about the image of an edge under $\tau_{n,m}$ as long as we restrict our attention to a single equivalence class.

Claim 3 [19, Lemma 8.2] For $n > N_2$ if $\sigma_{n,n+1}$ isn't a bijection then there is some j and an edge $e \subset Y_n^j$ such that

$$\mathrm{Stab}^+_{Y_n^j}(e) < \mathrm{Stab}^+_{Y_{n+1}^j}(\tau_{n,n+1}(e))$$

Proof of Claim 3. Claim 1 implies that Y_n^j must join onto a conjugate of itself under $\tau_{n,n+1}$. Claim 2 implies that we have an edge $e \subset Y_n^j$ and a $g \in G \setminus \mathrm{Stab}(Y_{n+1}^j)$ such that $\tau(e) = \tau(ge)$. We thus have $g \in \mathrm{Stab}^+_{Y_{n+1}^j}(\tau_{n,n+1}(e)) \setminus \mathrm{Stab}^+_{Y_n^j}(e)$. □

We are now ready to show that $\sigma_{n,n+1}$ is a bijection for all sufficiently large n. Suppose this isn't the case; then since there are only finitely many orbits of edges in each Y_n^j Claim 3 now implies that there is some j and some subsequence $\{n_{i_k}\}$ of $\{n_i\}$ and some edge $e \in Y_{N_2}^j$ such that

$$\mathrm{Stab}^+_{Y_{n_{i_1}}^j}(e_{n_{i_1}}) < \mathrm{Stab}^+_{Y_{n_{i_2}}^j}(e_{n_{i_2}}) < \mathrm{Stab}^+_{Y_{n_{i_3}}^j}(e_{n_{i_3}}) < \cdots$$

where $e_n = \tau_{N_2,n}(e) \in Y_n^j$. However this is exactly the situation the ACC says cannot happen. □

Before proving Theorem 4.2.5 we need one final statement about the rigidity of the Y_n^j. This essentially says that eventually these complexes look identical at every level.

Lemma 4.5.6. *There is some $N'' \geq N'$ with the following property. Suppose $n, m \geq N''$ and t, t' are triangles at depth n where $(\tau_{n,m}(t), \tau_{n,m}(t'))$ is a stable pair at depth m. Then (t, t') is also a (stable) pair at depth n. In other words $\tau_{n,m}$ induces a bijection on the set of stable pairs for $n, m \geq N''$.*

Proof. Let e, e' be the respective edges of t, t' (at level $n > N'$) which get mapped to the common edge of the pair $(\tau_{n,m}(t), \tau_{n,m}(t'))$. Since $n \geq N'$ we must have t and t' in the same equivalence class of $\sim_n$; call the corresponding subcomplex Y. We also must have $e' = ge$ for some g which stabilises Y as $N' > N_2$. The same argument as the proof of Claim 3 in the proof of Lemma 4.5.5 implies that either $e' = e$ or $\mathrm{Stab}_Y^+(e) < \mathrm{Stab}_{\tau_{n,m}(Y)}^+(\tau_{n,m}(e))$. As in the proof of Lemma 4.5.5 the ACC says this latter case can only occur finitely many times. □

Proof of Theorem 4.2.5. In light of Lemma 4.5.2 and Remark 4.4.11 it suffices to show that every simple cone in a complex of depth at least N'' is contained in an equivalence class. We define the *push-forward* of a cone as follows. Let c be the central vertex of a cone C. If the resolution ρ induces track(s) on X whose intersection with C is homeomorphic to a circle enclosing c then we let s be the outermost such track. Otherwise set $s = c$. We now define the push-forward of C to be the union of the image of the triangles in C which are in the same component as the image of s.

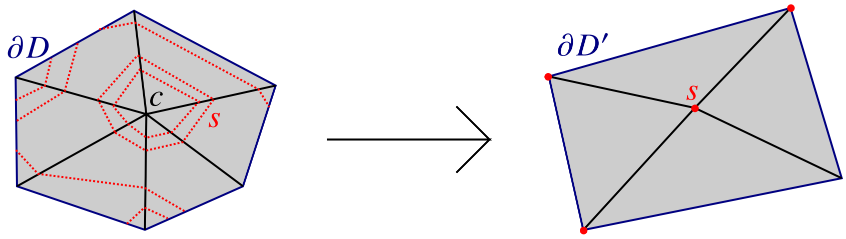

Let C be a simple cone at depth $n \geq N''$. Apply push-forwards to C until we reach a cone with minimal circumference; call this new cone C'. Observe that C' is made of consecutive stable pairs and so is contained in an equivalence class. If C' has the same circumference as C then we are done, so assume that the circumference of C' is strictly smaller than that of C. In this case we see C' must contain a (stable) pair of adjacent triangles that weren't adjacent in C; however this contradicts Lemma 4.5.6. □

References

[1] A.Baudisch: *Subgroups of semifree groups.* Acta Math. Acad. Sci. Hungar. 38 (1981), no. 1–4, 19–28.

[2] M.Bestvina, M.Feighn: *Bounding the complexity of simplicial group actions on trees.* Invent. Math. 103 (1991), no. 3, 449–469.

[3] C.Bleak, M.Kassabov, F.Matucci: *Structure theorems for groups of homeomorphisms of the circle. (English summary)* Internat. J. Algebra Comput. 21 (2011), no. 6, 1007–1036.

[4] O.V.Bogopol'skiĭ, V.N.Gerasimov: *Finite subgroups of hyperbolic groups. (Russian. Russian summary)* Algebra i Logika 34 (1995), no. 6, 619–622, 728; translation in Algebra and Logic 34 (1995), no. 6, 343–345 (1996)

[5] B.H.Bowditch: *Cut points and canonical splittings of hyperbolic groups.* Acta Math. 180 (1998), no. 2, 145–186.

[6] N.Brady: *Finite subgroups of hyperbolic groups. (English summary)* Internat. J. Algebra Comput. 10 (2000), no. 4, 399–405.

[7] M.Bridson, A.Haefliger: *Metric spaces of non-positive curvature.* Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 319. Springer–Verlag, Berlin, 1999. xxii+643 pp. ISBN: 3–540–64324–9

[8] T.Delzant: *Sur l'accessibilité acylindrique des groupes de présentation finie. (French. English, French summary) [On the acylindrical accessibility of finitely presented groups]* Ann. Inst. Fourier (Grenoble) 49 (1999), no. 4, 1215–1224.

[9] T.Delzant, L.Potyagailo: *Accessibilité hiérarchique des groupes de présentation finie. (French. English summary) [Hierarchical accessibility of finitely presented groups]* Topology 40 (2001), no. 3, 617–629.

[10] M.J.Dunwoody: *An inaccessible group.* Geometric group theory, Vol. 1 (Sussex, 1991), 75–78, London Math. Soc. Lecture Note Ser., 181, Cambridge Univ. Press, Cambridge, 1993.

[11] M.J.Dunwoody: *The accessibility of finitely presented groups.* Invent. Math. 81 (1985), no. 3, 449–457.

[12] É.Ghys and P.de la Harpe. *Sur les groupes hyperboliques d'après Mikhael Gromov. (French)* Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988. Edited by É. Ghys and P. de la Harpe. Progress in Mathematics, 83. Birkhäuser Boston, Inc., Boston, MA, 1990. xii+285 pp. ISBN: 0–8176–3508–4

[13] I.A.Grushko: *On the bases of a free product of groups.* Matematicheskii Sbornik, vol 8 (1940), pp. 169–182

[14] V.Guirardel, G.Levitt: *JSJ decompositions of groups. (English, French summary)* Astérisque No. 395 (2017), vii+165 pp. ISBN: 978–2–85629–870–1

[15] V.Guirardel, G.Levitt: *Trees of cylinders and canonical splittings. (English summary)* Geom. Topol. 15 (2011), no. 2, 977–1012.

[16] W.Haken: *Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I. (German)* Math. Z. 80 (1962), 89–120.

[17] D.Juan-Pineda, I.J.Leary: *On classifying spaces for the family of virtually cyclic subgroups. (English summary)* Recent developments in algebraic topology, 135–145.

[18] P.A.Linnell: *On accessibility of groups.* J. Pure Appl. Algebra 30 (1983), no. 1, 39–46.

[19] L.Louder, N.Touikan: *Strong accessibility for finitely presented groups. (English summary)* Geom. Topol. 21 (2017), no. 3, 1805–1835.

[20] Z.Sela: *Acylindrical accessibility for groups. (English summary)* Invent. Math. 129 (1997), no. 3, 527–565.

[21] J-P.Serre: *Trees.* Translated from the French original by John Stillwell. Corrected 2nd printing of the 1980 English translation. Springer Monographs in Mathematics. Springer–Verlag, Berlin, 2003. x+142 pp. ISBN: 3–540–44237–5

[22] J.R.Stallings: *Foldings of G-trees.* Arboreal group theory (Berkeley, CA, 1988), 355–368, Math. Sci. Res. Inst. Publ., 19, Springer, New York, 1991.

[23] J.R.Stallings: *Group theory and three-dimensional manifolds.* A James K. Whittemore Lecture in Mathematics given at Yale University, 1969. Yale Mathematical Monographs, 4. Yale University Press, New Haven, Conn.-London, 1971. v+65 pp.

[24] R.Wade: *Folding free-group automorphisms. (English summary)* Q. J. Math. 65 (2014), no. 1, 291–304.

[25] R.Weidmann: *On accessibility of finitely generated groups. (English summary)* Q. J. Math. 63 (2012), no. 1, 211–225.

[26] R.Weidmann: *The Nielsen method for groups acting on trees.* Proc. London Math. Soc. (3) 85 (2002), no. 1, 93–118.