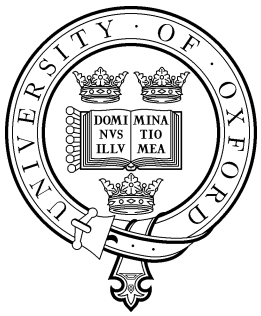
Quasi-Trees and Product Set Growth



Alice Kerr

St John's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2022

To Granny and Grandma.
I wish you could have seen me graduate, I like to think that you would have been proud.

Acknowledgements

Above all I would like to thank my supervisor Cornelia Druţu, for introducing me to the topics in this thesis, and for the many questions, suggestions, and discussions which helped shape this work. I am also grateful for the kindness and understanding you showed me during the pandemic, for all of the advice about the wider maths world, and for the views of Exeter College in the snow.

I would also like to express my gratitude to everyone in the mathematical community who were so generous in sharing their time and knowledge with me, who asked questions that influenced the direction of my work, and who gave valuable comments and feedback on earlier versions of this thesis. In particular, I would like to thank Ric Wade, Jack Button, Panos Papasoglu, David Hume, Johanna Mangahas, Sahana Balasubramanya, Nicolaus Heuer, Elia Fioravanti, Ashot Minasyan, Alessandro Sisto, and Markus Steenbock. Special thanks go to my examiners, Thomas Delzant and Emmanuel Breuillard, along with the referees of my first paper, for their many helpful suggestions and corrections.

Thank you to all of the people who made my experience of the Mathematical Institute as amazing as it was: Christoph Weis, Joe Scull, Luciana Bonatto, Elia Fioravanti, Jan Steinebrunner, Jane Tan, Adele Jackson, Diego Berdeja Suarez, Tomáš Zeman, Ido Grayevsky, David Hume, Ric Wade, Carla Groenland, Nicolaus Heuer, Filippos Sytilidis, Patrick Nairne, Monika Kudlinska, Davide Spriano, Sam Hughes, and many more. Thank you for all the tea breaks, Thursday doughnuts, table football matches, cake seminars, climbing sessions, and pub trips (and even for the occasional help with maths questions).

Special thanks go to Christoph, Joe, and Jordan Wheeler, for introducing me to Rocket League at the start of the pandemic. You may not have meant to change my life, but you did. Also thank you to Chris for sharing an office with me for three years, your work ethic always inspired me, and I'm sorry for trying to get you to watch so many TV shows.

These acknowledgements would not be complete without expressing my gratitude to all of the societies and sports teams that made my Oxford experience about far more than just a degree. Firstly, to the St John's College Middle Common Room, for helping to make college my home, and for all of the wonderful events that made my time here so special (particularly the exchange dinners, my Oxford college spreadsheet would be far more empty without them). I would especially like to thank Jana Muschinski and Patric Bonnier for being such amazing friends over the past few years,

and also Christopher D'Urso, for stepping up to the presidency when we needed him the most, and for doing such an incredible job of it. Thank you also to Eiko Soga for keeping me company during the pandemic, and to Annabel Hancock and Filippos, for the focaccias, Pride and Prejudice, and slightly too much Taylor Swift.

To the St John's College Boat Club, for teaching me how to cox, and how to row with only one blade (although I still think two is better), and most impressively for inspiring me to get on an erg again, ten years later. To the Saints, for that amazing Cuppers win, for allowing me to be immortalised in chalk on the walls of Tommy White quad, and for putting up with all of my shouting. To everyone in the Oxford University Women's Futsal Team, for following me on all those long trips up north, and for allowing me to play the best version of football for a few years longer. Finally, to my Rocket League teammates in the Oxford University Esports Society, for allowing me to occasionally pretend that I'm good at the game. I am especially grateful to Ollie and Cal, for encouraging me to play with them when I could still barely hit the ball.

Unfortunately the pandemic meant that I didn't get as much time with those clubs and societies as I would have wished for, however it is hard to be too sad when I consider what I gained instead. For that I would like to thank Leth and his community, especially Stale, Geo, Death, Kikira, Zulu, Static, Bryan, Jort, Tuba, Josh, and Hest, for giving me a place to belong during the pandemic, for giving my life some semblance of structure when it didn't have any, and for introducing me to so many other amazing people along the way. Also thank you to Matt Ellsworth for stopping me from walking away from it all, I would certainly have regretted it.

I of course would like to thank all of my family, but some deserve a special mention. To Grandpa, whose wisdom, generosity, and joke collection I aspire to emulate. To Grandma, for the debates, the walks on Box Hill, and the birdsong. To Granny, for the phone conversations, the family history, and for teaching me why I should support Brazil over England. To my parents, for putting up with the endless Skype calls when I was deciding where to do my DPhil, and for giving me a love of sports that has shaped so much of my time at university. To my dad, for showing me that there was far more to maths than what we were taught in school. To my mum, for putting up with being constantly surrounded by mathematicians (perhaps I'll go to medical school in my thirties, just to make things even). And finally to my brother, for being the one person I know I could always talk to about anything. I hope you meet as many incredible people in your PhD journey as I did in mine.

Abstract

In this thesis we investigate the geometric properties of quasi-trees, and study product set growth in groups with acylindrical actions on quasi-trees and hyperbolic spaces.

The main tool we use when considering quasi-trees is the end-approximating tree, which is a tree that we can construct from a geodesic space while preserving its space of ends. We use this construction to prove that every quasi-tree is $(1, C)$-quasi-isometric to a simplicial tree, one consequence of which is that having a cobounded quasi-action on a simplicial tree is equivalent to having a quasi-conjugate $(1, C)$-quasi-action on a simplicial tree. The end-approximating tree can also be used to show that the boundary of a quasi-tree is isometric to the boundary of its approximating tree under a certain choice of visual metric, and that this gives us a natural extension of the standard construction of a visual metric on the boundary of a tree.

As another consequence, we are able to show that Gromov's Tree Approximation Lemma for hyperbolic spaces can be improved in the case of quasi-trees to give a uniform approximation for any set of points, independent of cardinality. From this we show that having uniform tree approximation for finite subsets is equivalent to being able to uniformly approximate the entire space by a tree.

This uniform version of Gromov's Tree Approximation Lemma has an application to Delzant and Steenbock's work on product set growth in acylindrically hyperbolic groups. In particular, we are able to show that for any group G that acts acylindrically on a quasi-tree, there exists $\alpha > 0$ such that for any finite $U \subset G$ that has large enough displacement in the quasi-tree, and is not contained in a virtually cyclic subgroup, we have that $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$. By a result of Balasubramanya, this theorem applies to every acylindrically hyperbolic group.

We also study growth in acylindrically hyperbolic groups directly from their actions on hyperbolic spaces, which allows us to get a dichotomy on the finitely generated subgroups of mapping class groups, at the cost of working only with symmetric subsets. Specifically, we show that for every surface S of finite type, there exist $\alpha, \beta > 0$ such that for any finite symmetric subset U of the mapping class group $MCG(S)$ we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$, so long as no finite index pure subgroup of $\langle U \rangle$ has an infinite order central element.

As right-angled Artin groups embed as subgroups of mapping class groups, a similar dichotomy also applies in this case. We separately prove that we can quickly generate loxodromic elements in right-angled Artin groups, which by a result of Fujiwara shows that the set of growth rates for many of their subgroups are well-ordered.

Statement of Originality

I declare that, to the best of my knowledge, the contents of this thesis are original and my own work, unless indicated otherwise. Chapter 2 is based on [54], and Chapters 3-5 are based on [55], both of which have been submitted for publication.

None of the material in this thesis has been submitted wholly or in part towards any other degree.

Contents

1 Introduction **1**

1.1 Geometry of quasi-trees 1

1.2 Product set growth in acylindrically hyperbolic groups 6

1.3 Structure of the Thesis 12

2 Quasi-trees **14**

2.1 Trees, quasi-trees, and hyperbolic spaces 15

2.2 End-approximating tree for a geodesic metric space 22

2.3 Quasi-isometries with trees 27

2.3.1 $\mathbb{R}$-trees 27

2.3.2 Simplicial trees 30

2.3.3 Quasi-actions 34

2.3.4 Countable and locally finite trees 39

2.4 Uniform and non-uniform tree approximation 44

2.5 Geometry of the end-approximating tree 51

2.5.1 Path components 52

2.5.2 Ends of T_X for proper metric spaces 55

2.5.3 Visual metrics on the boundaries of quasi-trees 59

3 Product set growth **64**

3.1 Exponential and uniform exponential growth 65

3.1.1 Growth of groups 65

3.1.2 Growth of subgroups 70

3.2 Growth of arbitrary subsets 73

3.3 Uniform product set growth 77

3.4 Known results for product set growth 80

3.5 Applications of uniform product set growth 85

3.5.1 Helfgott type growth 86

3.5.2 Approximate groups 87
3.6 Obstructions and sufficient conditions 88
3.6.1 Infinite order central subgroups 89
3.6.2 Direct products with non-exponential factors 91
3.6.3 Direct products of groups with uniform product set growth . . 96
3.6.4 Passing to finite index supergroups 100

4 Growth in acylindrically hyperbolic groups 102
4.1 Acylindrical actions and loxodromic elements 102
4.2 Growth from actions on quasi-trees 107
4.2.1 The reduction lemma for quasi-trees 108
4.2.2 Growth of sets with large displacement 114
4.2.3 Loxodromic elements and displacement 115
4.3 Growth from actions on hyperbolic spaces 117

5 Applications to subgroups of mapping class groups 119
5.1 Right-angled Artin groups and mapping class groups 120
5.2 Short loxodromics in right-angled Artin groups 129
5.3 Product set growth in right-angled Artin groups 136
5.4 Product set growth in mapping class groups 141

6 Further questions 146

Bibliography 150

List of Figures

1.1 Tree approximation in a hyperbolic space 3

2.1 The comb space 19
2.2 Manning's bottleneck criterion 22
2.3 End-approximating tree of the hyperbolic plane 27
2.4 Approximating tree of a quasi-tree 27
2.5 Constructing a simplicial tree from an $\mathbb{R}$-tree 30
2.6 A counterexample to uniform subtree approximation 48

3.1 Balls of radius two in the free group and $\mathbb{Z}^2$ 67

4.1 A loxodromic element acting on the hyperbolic plane 106
4.2 The reduction lemma 109

5.1 An example of an active subsurface 128

1

Introduction

Real mathematics must be justified as art if it can be justified at all.
– G.H. Hardy, A Mathematician's Apology

Most mathematicians have the visual sense of a dead codfish. (However, many mathematicians smell rather less and a few have better conversation.)
– T.W. Körner, How to Write a Part III Essay

The inspiration for this thesis came from a paper by Delzant and Steenbock, in which they obtained growth theorems for groups with certain actions on hyperbolic spaces [31]. At one point in their proof they used a result called Gromov's Tree Approximation Lemma, which introduced logarithm terms into the statement of their final theorem. In a remark they stated that if the space being acted on was a quasi-tree, instead of a hyperbolic space, then the statement of Gromov's Tree Approximation Lemma would be improved, and therefore the statement of their theorem would also be improved.

There were two things in this remark that needed to be checked. The first, that this improvement to the lemma would indeed hold for quasi-trees, and the second, that this improvement would lead to a stronger statement of Delzant and Steenbock's growth result in this case. As it turns out, both of these things are true, but more importantly they both lead to some far deeper and more interesting results about quasi-trees and the growth of groups. This thesis is divided into two parts, following these two themes.

1.1 Geometry of quasi-trees

The first part of this thesis consists of Chapter 2, which is largely geometric in nature, and focuses on the properties of quasi-trees, which are the geodesic metric spaces

that are quasi-isometric to some simplicial tree. Quasi-trees are hyperbolic spaces, however the existence of such a quasi-isometry means that they retain many of the stronger geometric properties that trees enjoy. For this reason they often act as the most natural generalisation of trees, especially in cases where the space simply being hyperbolic would be too weak of an assumption.

Quasi-trees are an object of interest not just because of their tree-like properties, but because of the group actions that they admit. Importantly, there are large classes of groups which have interesting actions on quasi-trees, but do not have any similar actions on trees. Many such examples, including mapping class groups, were provided by Bestvina, Bromberg, and Fujiwara. In particular, they provided a set of axioms that, if satisfied, allow the construction of an action on a quasi-tree [9]. It was additionally shown by Kim and Koberda that right-angled Artin groups admit natural acylindrical actions on quasi-trees, where the quasi-trees in question are constructed from the defining graph of the group [56, 57]. Balasubramanya later proved that in fact every acylindrically hyperbolic group admits an acylindrical action on some quasi-tree [6]. This gives a strong motivation for proving new properties for quasi-trees, as they may allow us to obtain new results for these large classes of groups.

When attempting to find new properties of quasi-trees, there are two obvious options: generalise a result that is true for trees, or improve a result that has already been shown for hyperbolic spaces. As mentioned previously, the original aim of Chapter 2 was to do the latter for a result known as Gromov's Tree Approximation Lemma, which essentially says that any finite set of points in a hyperbolic space can be approximated by a finite subset of a tree, with the error dependent on the number of points we are trying to approximate.

Proposition 1.1.1 (Gromov's Tree Approximation Lemma)**.** [43, p. 155–157] *Let (X, d) be a δ-hyperbolic geodesic metric space. Let $x_0, z_1, \dots, z_n \in X$, and let Y be a union of geodesic segments $\bigcup_{i=1}^{n}[x_0, z_i]$. Then there is an $\mathbb{R}$-tree T and a map $f : (Y, d) \to (T, d')$ such that:*

(1) For all $1 \leqslant i \leqslant n$, the restriction of f to the geodesic segment $[x_0, z_i]$ is an isometry.

(2) For all $x, y \in Y$, we have that $d(x, y) - 2\delta(\log_2(n)+1) \leqslant d'(f(x), f(y)) \leqslant d(x, y)$.

This is the result that was used by Delzant and Steenbock [31] to generalise a theorem they proved about groups acting acylindrically on trees, from which they were able to obtain a theorem about groups acting acylindrically on hyperbolic spaces.

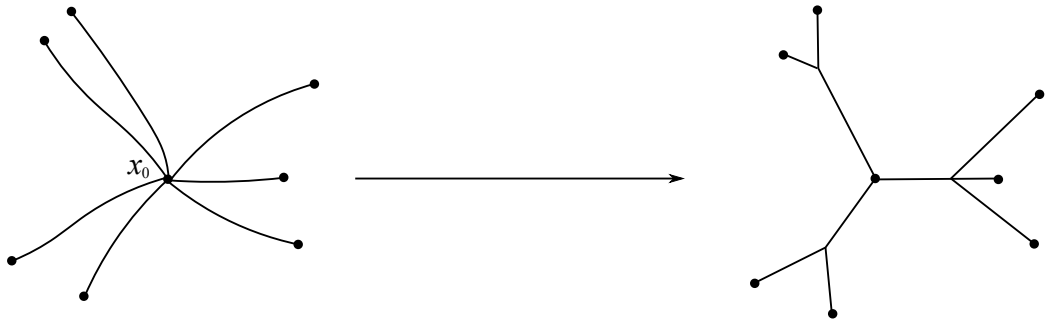


Figure 1.1: Tree approximation in a hyperbolic space

Their result for hyperbolic spaces was not as strong as their result for trees, which was due to the logarithmic error that Gromov's Tree Approximation Lemma introduced.

For general hyperbolic spaces we would not expect to be able to improve the order of this error, however it is natural to ask whether we may be able to do better when we restrict ourselves to trying to approximate subsets of quasi-trees. In particular, we may ask if the approximation could be made to be uniform, in the sense that the error would no longer depend on the number of points being approximated.

It turns out that we can in fact prove a much stronger result than this, as we can show that the entire quasi-tree can be approximated by a tree with uniform error.

Proposition 1.1.2 (Corollary 2.3.3)**.** *For every quasi-tree X there exists an $\mathbb{R}$-tree T and a constant $C \geqslant 0$ such that X is $(1, C)$-quasi-isometric to T.*

The most important implication of this, when combined with the fact that $\mathbb{R}$-trees are $(1, C)$-quasi-isometric to simplicial trees (see Proposition 2.3.4), is that quasi-trees are exactly the geodesic spaces that are $(1, C)$-quasi-isometric to simplicial trees. This is the central result of Chapter 2.

Theorem 1.1.3 (Theorem 2.3.6)**.** *A geodesic metric space X is a quasi-tree if and only if it is $(1, C)$-quasi-isometric to a simplicial tree for some $C \geqslant 0$.*

One application of this is to quasi-actions on trees. It is known that a group has a cobounded quasi-action on a simplicial tree if and only if it has a quasi-conjugate isometric action on a quasi-tree [66]. Theorem 1.1.3 allows us to show that this implies the existence of a quasi-conjugate $(1, C)$-quasi-action on some simplicial tree.

Proposition 1.1.4 (Proposition 2.3.18)**.** *If a finitely generated group admits a cobounded (L, C)-quasi-action on a simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1, C')$-quasi-action on a simplicial tree for some $C' \geqslant 0$.*

We are also able to obtain results similar to Theorem 1.1.3 for countable and locally finite simplicial trees.

Proposition 1.1.5 (Corollary 2.3.30)**.** *We have the following:*

(1) A geodesic metric space is (L, C)-quasi-isometric to a countable simplicial tree for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to a countable simplicial tree for some $C' \geqslant 0$.

(2) A geodesic metric space is (L, C)-quasi-isometric to a locally finite simplicial tree for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to a locally finite simplicial tree for some $C' \geqslant 0$.

This allows us to get analogues to Proposition 1.1.4 in the case of quasi-actions on countable or locally finite trees.

Corollary 1.1.6 (Corollary 2.3.31)**.** *We have the following:*

(1) If a finitely generated group admits a cobounded (L, C)-quasi-action on a countable simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1, C')$-quasi-action on a countable simplicial tree for some $C' \geqslant 0$.

(2) If a finitely generated group admits a cobounded (L, C)-quasi-action on a locally finite simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1, C')$-quasi-action on a locally finite simplicial tree for some $C' \geqslant 0$.

Returning to our original motivation, the tree and quasi-isometry in Proposition 1.1.2 are constructed by adapting the method used to prove Gromov's Tree Approximation Lemma, and so they satisfy the same conditions. Restricting this quasi-isometry to subsets gives us our desired improvement of Gromov's result in the case of quasi-trees.

Proposition 1.1.7 (Proposition 2.4.1)**.** *Let (X, d) be a quasi-tree. Let $x_0 \in X$, and let $Z \subset X$. Let Y be a union of geodesic segments $\bigcup_{z \in Z} [x_0, z]$. Then there is an $\mathbb{R}$-tree T, a map $f : (Y, d) \to (T, d^*)$, and a constant $C \geqslant 0$ such that:*

(1) For all $z \in Z$, the restriction of f to the geodesic segment $[x_0, z]$ is an isometry.

(2) For all $x, y \in Y$, we have that $d(x, y) - C \leqslant d^*(f(x), f(y)) \leqslant d(x, y)$.

We note here that Proposition 1.1.7 applies to subsets of any cardinality, rather than just finite subsets. On the other hand, it turns out that simply having uniform tree approximation for finite subsets is enough to ensure that the space in question is a quasi-tree, so quasi-trees are exactly the geodesic spaces for which uniform tree approximation is possible. In fact, even having uniform tree approximation for finite subsets is a stronger requirement than is necessary, as we still get a quasi-tree even if we relax the requirements on our tree approximation to allow for a multiplicative error.

Proposition 1.1.8 (Proposition 2.4.5)**.** *Let (X, d) be a geodesic metric space. The following are equivalent:*

(1) X is a quasi-tree.

(2) There exists a constant $C \geqslant 0$ such that for every finite subset Z of X, there exists an $\mathbb{R}$-tree (T, d^*) and a $(1, C)$-quasi-isometric embedding $f : (Z, d) \to (T, d^*)$.

(3) There exist constants $L \geqslant 1$ and $C \geqslant 0$ such that for every finite subset Z of X, there exists an $\mathbb{R}$-tree (T, d^*) and a (L, C)-quasi-isometric embedding $f : (Z, d) \to (T, d^*)$.

We are also able to characterise quasi-trees in terms of the Gromov product, which is defined as $(x, y)_{x_0} = \frac{1}{2}(d(x_0, x) + d(x_0, y) - d(x, y))$. This gives us our third equivalent characterisation, along with Theorem 1.1.3 and Proposition 1.1.8.

Proposition 1.1.9 (Proposition 2.3.1)**.** *Let X be a geodesic metric space. Then X is a quasi-tree if and only if there exists $A \geqslant 0$ such that any $x_0, x_1, \ldots, x_n \in X$ satisfy*

$$(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0} - A.$$

The remainder of Chapter 2 will mainly be concerned with the construction of the tree used in Proposition 1.1.2, and the ways in which its geometry relates to that of the original space. Such a tree can be constructed for any geodesic space, and can be thought of as what you get if you collapse the space along spheres around a basepoint, where two points on a sphere collapse together if they can be joined by a path that stays outside the central ball.

This construction can alternatively be visualised as collapsing the space along its ends, and we can show that the space of ends of any proper geodesic space will indeed be homeomorphic to the space of ends of the constructed tree (see Proposition 2.5.22).

For this reason we will often refer to such a tree as the *end-approximating tree*. As a consequence, we can show that the space of ends of any proper geodesic space will in fact be homeomorphic to the space of ends of a locally finite simplicial tree.

Proposition 1.1.10 (Proposition 2.5.23)**.** *Let X be a proper geodesic metric space. There exists a locally finite simplicial tree T such that* $\mathrm{Ends}(X)$ *is homeomorphic to* $\mathrm{Ends}(T)$.

The boundary of the end-approximating tree can be compared to the boundary of the original space. For a hyperbolic space X there exists a family of metrics on the boundary ∂X known as visual metrics, and given a tree T, and a choice of parameter, the visual metrics with that parameter on ∂T are all bi-Lipschitz equivalent to a visual metric that can be written down explicitly in a standard way. It is known that a $(1, C)$-quasi-isometry between two hyperbolic geodesic spaces induces a bi-Lipschitz map between their boundaries [13], and more specifically it is possible to choose a visual metric such that this map is an isometry. Proposition 1.1.2 can therefore be used to show the following.

Corollary 1.1.11 (Corollary 2.5.37)**.** *Let X be a quasi-tree. There exists an $\mathbb{R}$-tree T and a visual metric on ∂X such that ∂X endowed with this metric is isometric to ∂T endowed with its standard visual metric, given a choice of basepoint and visual parameter.*

In particular, the visual metric that we choose on ∂X can be defined directly from the geometry of X, and in the case that X is an $\mathbb{R}$-tree it will be the standard visual metric on the boundary. This means that this metric can be viewed as an extension of the standard visual metric for the boundary of an $\mathbb{R}$-tree, with respect to a choice of basepoint and parameter.

1.2 Product set growth in acylindrically hyperbolic groups

The second part of this thesis will consist of Chapters 3, 4, and 5, which focus on product set growth in groups. For a finite subset U of a group G, we define its nth product set to be

$$U^n = \{u_1 \cdots u_n \ : \ u_1, \ldots, u_n \in U\}.$$

The study of product set growth in groups is the study of how $|U^n|$ behaves as n varies. For infinite groups, the classical notion of group growth is when G is taken

to be finitely generated, and U is taken to be a ball in G with respect to the word metric induced by some finite generating set S. The usual question that is asked is whether $|U^n|$ grows like a polynomial function, an exponential function, or if it lies somewhere between the two.

This is known to have links to the algebraic properties of the group, the most famous example being Gromov's proof that the groups of polynomial growth are exactly the virtually nilpotent groups [42]. One of the nice aspects of this asymptotic growth question is that if U and U' are balls given by different finite generating sets of G, then the answer will be the same for both. This will be discussed briefly in Section 3.1.

In finite groups, on the other hand, there is little interest in estimating $|U^n|$ for large powers of n, as this is bounded above by the size of the group, and so will eventually become constant. Instead, questions about product set growth in this setting tend to focus on getting precise bounds on low powers, such as $|U^2|$ or $|U^3|$.

The question we are interested in here combines aspects of these two cases. Our focus will be on infinite groups, and for U a general finite subset we would like to estimate $|U^n|$ for every $n \in \mathbb{N}$. On the other hand, instead of only being concerned with the asymptotic profile of $|U^n|$, we will be interested in getting an explicit lower bound on $|U^n|$, with this lower bound dependent on the size of U. More specifically, we would like to answer the following question.

Question. *For a group G, do there exist constants $\alpha, \beta > 0$ and a class of subgroups $\mathcal{H}$ such that for every finite (symmetric) $U \subset G$ where $\langle U \rangle \notin \mathcal{H}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$?*

If all of the generating sets of a group satisfy this inequality for the same constants α and β, then we say that the group has *uniform product set growth*. The aim here however is not just to prove this for the group itself, but to find a dichotomy of finitely generated subgroups, so that every subgroup either belongs to $\mathcal{H}$ and does not have uniform product set growth, or does not belong to $\mathcal{H}$ and has uniform product set growth for some common α and β.

This question has strong links to the more commonly studied notion of uniform exponential growth. Specifically, any finitely generated group such that every symmetric generating set satisfies the given inequality will necessarily have uniform exponential growth. Moreover, the finitely generated subgroups that do not lie in $\mathcal{H}$ will have uniform uniform exponential growth, in the sense that they will have a common lower bound on their exponential growth rates. Such properties have already been shown

for a wide range of groups, including mapping class groups and right-angled Artin groups [62], although the question of uniform exponential growth is still open for acylindrically hyperbolic groups in general.

On the other hand, we can see that uniform product set growth is a strictly stronger property than uniform exponential growth. The key difference is that we are asking for the lower bound on the growth to be given in terms of the size of the set, rather than just asking for some common lower bound. Linking growth rates with the size of the generating set in question is a key step in proofs about growth rates being well-ordered, see [38] and [37].

Our main example that uniform product set growth is strictly stronger than uniform exponential growth comes from the following result.

Proposition 1.2.1 (Proposition 3.6.1)**.** *Let G be a finitely generated group, and suppose that it has an infinite normal subgroup H, where H does not have exponential growth, and for every $h \in H$ the conjugacy class of h in G is finite. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

In particular, this applies to any group with an infinite finitely generated central subgroup.

Corollary 1.2.2 (Proposition 3.6.2)**.** *Let G be a finitely generated group, and suppose that the centre of G is infinite. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

The intuitive reason as to why such groups cannot have uniform product set growth is that we can add arbitrarily many central elements to any generating set U. As these elements are contained in an abelian subgroup, which does not have exponential growth, this has minimal impact on the long term growth of U^n. This means that we cannot link the growth of U^n with the size of U in a meaningful way. As a consequence, if our group contains subgroups with exponential growth but infinite centres, we cannot expect the dichotomy of subgroups we get in answer to our question to be the same as the dichotomy given by the Tits alternative.

Results concerning uniform product set growth have consequences for other types of product set growth as well, namely Helfgott type growth, as defined by Jack Button [19].

Proposition 1.2.3 (Proposition 3.5.3)**.** *Let G be a group, and let $U \subset G$ be finite. If there exist $\alpha, \beta > 0$ such that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, then $|U^3| \geqslant (\frac{\alpha^\beta}{8})^{\frac{1}{3}}|U|^{1+\frac{\beta}{6}}$.*

Uniform product set growth also has links with approximate groups, as any approximate group that satisfies the relevant inequality must have bounded size.

Proposition 1.2.4 (Proposition 3.5.6)**.** *Let G be a group, and let $k \geqslant 1$, $\alpha, \beta > 0$. Suppose U is a k-approximate group in G satisfying $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. Then $|U| \leqslant \frac{k^{\lceil 2/\beta \rceil - 1}}{\alpha^2}$.*

There are several known results in the area of product set growth in infinite groups, including complete answers to our dichotomy question for free groups [76], and more recently hyperbolic groups [31]. A more thorough history is given in Section 3.4. Here we will be interested in this second paper, by Delzant and Steenbock, which also contains a result concerning product set growth in acylindrically hyperbolic groups. To state it, we first need to say what they mean by displacement. If (X, d) is a hyperbolic space that a groups G acts on by isometries, and $U \subset G$ is finite, we pick $x_0 \in X$ such that $\sum_{u \in U} d(x_0, ux_0)$ is (almost) minimised. The *displacement* of U is then defined to be $\lambda_0(U) = \max_{u \in U} d(x_0, ux_0)$. For a more detailed definition, see Definition 4.2.1.

Theorem 1.2.5. [31] *Let G be a group acting acylindrically on a hyperbolic space. Then there exist constants $K > 0$ and $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K \log_2(2|U|)$.

(3) $|U^n| \geqslant \left(\frac{\alpha}{\log_2^6(2|U|)}|U|\right)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

The class of groups with interesting acylindrical actions on hyperbolic spaces is a wide one, including the hyperbolic groups and relatively hyperbolic groups that are not virtually cyclic, as well as many cubical groups, right-angled Artin groups, and mapping class groups of surfaces [73]. The above statement is interesting not only because it applies to all of these groups, but also because it is not restricted to symmetric sets. On the other hand it includes the previously mentioned logarithm terms, which we recall were introduced by the use of Gromov's Tree Approximation

Lemma. Once we have shown that this lemma is indeed uniform in the case of quasi-trees, we will use this result to show that we can also obtain a uniform version of Theorem 1.2.5 in the case that the space being acted on is a quasi-tree.

Theorem 1.2.6 (Theorem 4.2.13)**.** *Let G be a group acting acylindrically on a quasi-tree. Then there exist constants $K > 0$ and $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K$.

(3) $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

As mentioned previously, a result of Balasubramanya tells us that in fact every acylindrically hyperbolic group admits some acylindrical action on a quasi-tree [6]. The above theorem therefore applies to all such groups.

We will want to apply Theorem 1.2.6 to investigate uniform product set growth for subgroups of acylindrically hyperbolic groups, however we cannot do this immediately, as the displacement condition is a condition on sets, rather than on the subgroups they generate. In Section 4.2 we will show that one possible way of overcoming this is to find loxodromic elements in the action.

Proposition 1.2.7 (Proposition 4.2.18)**.** *Let G be a group acting acylindrically on a quasi-tree. There exist $\alpha, \beta > 0$ such that for every finite $U \subset G$ such that U^k contains a loxodromic element for $k \in \mathbb{N}$, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\frac{\beta n}{k}}$ for every $n \in \mathbb{N}$.*

In other words, if we can find loxodromic elements uniformly quickly in the generating sets of a subgroup, then we will have the uniform product set growth that we are looking for. Fujiwara has also recently shown a very similar result for actions on hyperbolic spaces, using an alternative method of generating free subgroups of large enough rank. The only downside to this is that it comes at the cost of restricting to symmetric sets, however in most cases this restriction is likely to be necessary anyway, as the common methods of finding loxodromic elements tend to assume symmetry.

Proposition 1.2.8. [37] *Let G be a group acting acylindrically on a hyperbolic space. There exist $\alpha, \beta > 0$ such that for every finite symmetric $U \subset G$ such that U^k contains a loxodromic element for $k \in \mathbb{N}$, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\frac{\beta n}{k}}$ for every $n \in \mathbb{N}$.*

One example where we already know that we can generate loxodromic elements uniformly quickly is in mapping class groups. In the following result we can think of finding an element with the same active subsurface as our subgroup as being roughly the same as finding a loxodromic element in the given subgroup.

Theorem 1.2.9. [63] *Consider a mapping class group $MCG(S)$, where S is a non-sporadic connected surface without boundary. There exists a constant $N = N(S) \in \mathbb{N}$ such that for any finite symmetric $U \subset MCG(S)$ there exists $n \leqslant N$ and $f \in U^n$ such that f has the same active subsurface as $\langle U \rangle$.*

These results combined allow us to answer our main question in the case of mapping class groups. The fact that the following theorem gives us a dichotomy for finitely generated subgroups is a consequence of Proposition 1.2.1, and will be shown in Section 5.4.

Theorem 1.2.10 (Corollary 5.4.3)**.** *Let $MCG(S)$ be a mapping class group. There exist $\alpha, \beta > 0$ such that for every finite symmetric $U \subset MCG(S)$ at least one of the following must hold:*

(1) $\langle U \rangle$ has a finite index pure subgroup with non-trivial centre.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Remark 1.2.11. Here we use the definition of mapping class groups where punctures in the surface are allowed to permute, but boundary components are fixed pointwise. A brief discussion of the variation in definitions of the mapping class group is given in Remark 5.1.29.

Answering our question completely for a group also naturally answers it for all of its subgroups. One class of groups that is known to embed in mapping class groups are the right-angled Artin groups (see [25], for example). In particular, they embed as pure subgroups, so the dichotomy is somewhat simplified in this case.

Theorem 1.2.12 (Theorem 5.3.6)**.** *Let Γ be a finite simple graph, and $A(\Gamma)$ the right-angled Artin group associated to Γ. There exist constants $\alpha, \beta > 0$ such that for every finite symmetric $U \subset A(\Gamma)$ at least one of the following must hold:*

(1) The centre of $\langle U \rangle$ is non-trivial.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

It is however also interesting to consider the right-angled Artin group case separately, as doing so allows us to prove an analogous result to Theorem 1.2.9 for right-angled Artin groups. In other words, we can show that we can quickly generate loxodromic elements in the action of the group on the associated extension graph, which is a quasi-tree.

In the theorem below, for a finite graph $\Gamma = (V, E)$ we have that $A(\Gamma)$ is the associated right-angled Artin group, Γ^e is the associated extension graph, and for a subset $V' \subset V$ we have that $\Gamma(V')$ is the induced subgraph on those vertices. For more detailed definitions see Section 5.1.

Proposition 1.2.13 (Proposition 5.2.14)**.** *Suppose Γ is a finite graph. For $U \subset A(\Gamma)$, let $V_U \subset V$ be minimal under inclusion such that U is conjugate into $A(\Gamma(V_U))$. There exists $N = N(\Gamma) \in \mathbb{N}$ such that for every finite symmetric $U \subset A(\Gamma)$, where $\Gamma(V_U)$ is connected and $\langle U \rangle$ is neither cyclic nor contained non-trivially in a direct product, there exists $n \leqslant N$ such that U^n contains a loxodromic element on $\Gamma(V_U)^e$.*

A recent result of Fujiwara states that if a group G has an acylindrical and non-elementary action on a hyperbolic graph, and the product of any finite generating set of G contains a loxodromic element within a bounded power, then the set of exponential growth rates of G (with respect to its finite generating sets) will be well-ordered so long as G is equationally Noetherian [37]. Linear groups are equationally Noetherian [7], and right-angled Artin groups are linear [48], so we get the following result.

Theorem 1.2.14 (Theorem 5.2.16)**.** *Let Γ be a finite graph. Suppose that $G \leqslant A(\Gamma)$ is finitely generated, and is not contained non-trivially in a direct product where the projection of G to more than one factor has exponential growth. Then the set of exponential growth rates of G (with respect to its finite generating sets) is well-ordered.*

1.3 Structure of the Thesis

As mentioned previously, Chapter 2 will be largely concerned with discussing the geometry of quasi-trees. In Section 2.1 we will collect the definitions and facts about quasi-trees and hyperbolic geometry that are necessary for the rest of the chapter. Once this background knowledge has been established, we will use Section 2.2 to give the construction of our end-approximating tree for a geodesic metric space, which will be the main tool we use in the following sections.

Having outlined this construction, it will then be used in Section 2.3 to prove Theorem 1.1.3, as well as a few analogous results. This section will also contain some consequences of Theorem 1.1.3 for quasi-actions on simplicial trees.

In Section 2.4 we will prove the uniform version of Gromov's Tree Approximation Lemma for quasi-trees. We will also consider an alternative version of tree approximation, and show that this turns out to not be uniform for quasi-trees. In Section 2.5 we will give an alternative description of the end-approximating tree, and use it to discuss the relationship between the ends of the tree and the ends of the space it is constructed from.

We will then move on to considering questions related to growth in groups. In Chapter 3 we will give some of the history of product set growth, as well as a few of its applications. We will also prove some results about which groups can and cannot have uniform product set growth, which will be vital in showing that the results in Chapter 5 represent dichotomies of finitely generated subgroups.

Chapter 4 will begin with some necessary background on acylindrically hyperbolic groups, then in Section 4.2 we will combine our uniform tree approximation result with Delzant and Steenbock's methods in [31] to get Theorem 1.2.6. The remainder of Chapter 4 will detail how being able to quickly find loxodromic isometries in a group can allow us to overcome the displacement condition in this result, and that a similar statement holds for actions on hyperbolic spaces.

The results of the previous chapters will be applied in Chapter 5 to prove Theorem 1.2.10, which will give us a dichotomy for the growth of subgroups of mapping class groups. This result will be built up to by first proving Proposition 1.2.13 regarding loxodromic elements in right-angled Artin groups, then using this to prove Theorem 1.2.12. This will provide a model for the method used to proved Theorem 1.2.10.

We will finish this thesis with Chapter 6, where we briefly go over a number of questions arising from the rest of the text.

2

Quasi-trees

A tree is a tree, how many more do you need to look at?

– Ronald Reagan

This chapter will be focused on the geometry of quasi-trees, with the main aim being to show how close the properties of quasi-trees are to those of trees. With this aim in mind, in Section 2.2 we will give the construction of what we will call the *end-approximating tree* for a geodesic metric space, which can be thought of as what we get if we collapse a space along its ends. The idea here is that when the space in question is a quasi-tree, the end-approximating tree will be a good approximation of this space.

In Section 2.3 we will prove that this is the case, and in fact the end-approximating tree for a quasi-tree is $(1, C)$-quasi-isometric to the original space. We can then use this to show that quasi-trees are precisely those spaces that are $(1, C)$-quasi-isometric to some simplicial tree. We will additionally consider the particular cases where the simplicial tree in question is countable or locally finite, and show that similar results hold there.

In Section 2.4 we will apply the results of Section 2.3 to prove that Gromov's Tree Approximation Lemma is uniform in the case of quasi-trees, meaning that any finite set of points can be approximated by a tree with only a uniform error between the distances. This will be used later in Chapter 4 to get a result about growth in acylindrically hyperbolic groups. In this section we will also consider an alternative version of tree approximation, and show that this turns out to not be uniform for quasi-trees. In Section 2.5 we will give an alternative description of the end-approximating tree, and use it to discuss the relationship between the ends of the tree and the ends of the space it is constructed from. In the case that the original space is a quasi-tree,

we can compare its boundary with the boundary of the end-approximating tree, and show that they can be thought of as being isometric.

2.1 Trees, quasi-trees, and hyperbolic spaces

We first give some basic definitions and lemmas.

Notation 2.1.1. Let (X, d) be a metric space. Let $x_0 \in X$ and $r \geqslant 0$. We will use $B(x_0, r)$ to denote the closed ball of radius r around x_0, and $S(x_0, r)$ to denote the sphere of radius r around x_0.

Definition 2.1.2. Let (X, d) be a metric space. Let $x_0, x, y \in X$. The *Gromov product* of x and y at x_0 is

$$(x, y)_{x_0} = \frac{1}{2}(d(x_0, x) + d(x_0, y) - d(x, y)).$$

Lemma 2.1.3. *Let (X, d) be a metric space. Let $x_0, x, y \in X$. Then the following hold:*

(1) $(x, y)_{x_0} = (y, x)_{x_0}$.

(2) $(x_0, x)_{x_0} = 0$.

(3) $(x, x)_{x_0} = d(x_0, x)$.

(4) $(x, y)_{x_0} \leqslant d(x_0, x)$.

(5) $(x, y)_{x_0} \geqslant 0$.

Proof. The first three statements are immediate from the definition of the Gromov product. The final two statements follow from the triangle inequality. □

The Gromov product can be used to define a certain type of hyperbolicity in metric spaces.

Definition 2.1.4. Let (X, d) be a metric space. Suppose there exists $\delta \geqslant 0$ such that for every $x_0, x, y, z \in X$ we have that

$$(x, z)_{x_0} \geqslant \min\{(x, y)_{x_0}, (y, z)_{x_0}\} - \delta.$$

Then we say that X is *δ-hyperbolic*. We say that X is *hyperbolic* if it is δ-hyperbolic for some $\delta \geqslant 0$.

Example 2.1.5. Any bounded metric space is hyperbolic, where δ can be taken to be the diameter of the space. The hyperbolic plane is also hyperbolic under this definition, as are all other hyperbolic n-spaces. On the other hand, $\mathbb{R}^n$ where $n \geqslant 2$ is not hyperbolic for any δ.

The logarithm in the lower bound of Gromov's Tree Approximation Lemma (Proposition 1.1.1) comes from the following statement, which can be proved by induction directly from the above definition.

Lemma 2.1.6. [43, p. 155] *A metric space X is δ-hyperbolic in the sense of Definition 2.1.4 if and only if any $x_0, x_1, \ldots, x_n \in X$ with $n \leqslant 2^k + 1$ satisfy*

$$(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0} - k\delta.$$

In geodesic metric spaces, there is a commonly used equivalent definition of hyperbolicity that is more intuitive.

Notation 2.1.7. For x, y in a geodesic metric space, we will use $[x, y]$ to represent any geodesic from x to y.

Definition 2.1.8. Let (X, d) be a geodesic metric space. We say that X is *δ-hyperbolic* for some $\delta \geqslant 0$ if for every $x, y, z \in X$ any choice of geodesic triangle $[x, y] \cup [y, z] \cup [z, x]$ is δ-slim, meaning that if $p \in [x, y]$ then there exists $q \in [y, z] \cup [z, x]$ such that $d(p, q) \leqslant \delta$.

Remark 2.1.9. The δ in each definition is not generally the same. The exception is when a geodesic metric space is 0-hyperbolic, as in this case it will be 0-hyperbolic under both definitions.

Convention. We will be mainly working in geodesic metric spaces, so we will take the second definition as standard when we say δ-hyperbolic, unless otherwise specified.

In a geodesic hyperbolic space, the Gromov product of x and y at x_0 is approximately the distance between x_0 and $[x, y]$, as shown in the following standard lemmas.

Lemma 2.1.10. *Let (X, d) be a δ-hyperbolic geodesic metric space, and let $x_0, x, y \in X$. Then there exists $z \in [x, y]$ such that $d(x_0, z) - 2\delta \leqslant (x, y)_{x_0}$.*

Proof. Consider the closed δ-neighbourhoods of $[x_0, x]$ and $[x_0, y]$, which we denote by $N([x_0, x], \delta)$ and $N([x_0, x], \delta)$. The geodesic $[x, y]$ is also closed, and by δ-hyperbolicity

$$[x, y] = ([x, y] \cap N([x_0, x], \delta)) \cup ([x, y] \cap N([x_0, y], \delta)).$$

As $[x, y]$ is connected, it cannot be the union of two disjoint closed sets. Therefore there exists $z \in [x, y] \cap N([x_0, x], \delta) \cap N([x_0, y], \delta)$. We pick $x' \in [x_0, x]$ and $y' \in [x_0, y]$ such that $d(x', z) \leqslant \delta$ and $d(y', z) \leqslant \delta$.

We can then see that

$$\begin{aligned}(x, y)_{x_0} &= \frac{1}{2}(d(x_0, x) + d(x_0, y) - d(x, y)) \\ &= \frac{1}{2}(d(x_0, x') + d(x', x) + d(x_0, y') + d(y', y) - d(x, z) - d(z, y)) \\ &\geqslant \frac{1}{2}(d(x_0, z) - \delta + d(z, x) - \delta + d(x_0, z) - \delta + d(z, y) - \delta \\ &\quad - d(x, z) - d(z, y)) \\ &= d(x_0, z) - 2\delta.\end{aligned}$$

□

Lemma 2.1.11. *Let (X, d) be a geodesic metric space, and let $x_0, x, y \in X$. Then $(x, y)_{x_0} \leqslant d(x_0, [x, y])$.*

Proof. Let $z \in [x, y]$. We have that

$$\begin{aligned}(x, y)_{x_0} &= \frac{1}{2}(d(x_0, x) + d(x_0, y) - d(x, y)) \\ &\leqslant \frac{1}{2}(d(x_0, z) + d(z, x) + d(x_0, z) + d(z, y) - d(x, z) - d(z, y)) \\ &= d(x_0, z).\end{aligned}$$

□

An alternative interpretation is that $(x, y)_{x_0}$ is roughly the distance that the geodesics $[x_0, x]$ and $[x_0, y]$ remain δ-close to each other.

Corollary 2.1.12. *Let (X, d) be a geodesic metric space, and let $x_0, x, y \in X$. For any $x' \in [x_0, x]$ such that $d(x_0, x') < (x, y)_{x_0} - \delta$, there exists $y' \in [x_0, y]$ such that $d(x', y') \leqslant \delta$.*

Proof. Lemma 2.1.11 tells us that there is no $z \in [x, y]$ such that $d(x', z) \leqslant \delta$, so we can apply the δ-slim triangle condition. □

Lemma 2.1.13. *Let (X, d) be a geodesic metric space, and let $x_0, x, y \in X$. For any $x' \in [x_0, x]$ such that $d(x_0, x') > (x, y)_{x_0} + \delta$, there does not exist $y' \in [x_0, y]$ such that $d(x', y') \leqslant \delta$.*

Proof. If such a y' did exist, then we would have that $d(x_0, y') > (x, y)_{x_0}$, and so

$$\begin{aligned} d(x, y) &\leqslant d(x, x') + d(y, y') + d(x', y') \\ &\leqslant d(x_0, x) - d(x_0, x') + d(x_0, y) - d(x_0, y') + \delta \\ &< d(x_0, x) + d(x_0, y) - 2(x, y)_{x_0} - 2\delta. \end{aligned}$$

This contradicts the fact that $d(x, y) = d(x_0, x) + d(x_0, y) - 2(x, y)_{x_0}$, by definition of the Gromov product. □

One special class of hyperbolic spaces are the $\mathbb{R}$-trees, which will play a central role in this chapter.

Definition 2.1.14. A metric space (X, d) is an $\mathbb{R}$-*tree* if for every $x, y \in X$ there exists a unique topological embedding $\alpha : [0, r] \to X$, and this embedding is a geodesic, so $d(x, y) = r$.

The following equivalent characterisation is standard, see for example [8].

Lemma 2.1.15. *A metric space is an $\mathbb{R}$-tree if and only if it is 0-hyperbolic and geodesic.*

Remark 2.1.16. One nice consequence of these definitions is that in an $\mathbb{R}$-tree the Gromov product $(x, y)_{x_0}$ gives us exactly the distance between x_0 and the unique geodesic $[x, y]$. This is also the exact distance for which $[x_0, x]$ and $[x_0, y]$ coincide.

$\mathbb{R}$-trees are not the only way that trees that are defined in the literature, although it is the most general definition. Another commonly studied version of trees are the simplicial trees, which will also feature in this chapter.

Definition 2.1.17. A *simplicial tree* is a 1-dimensional simplicial complex that is an $\mathbb{R}$-tree under the path metric induced by considering each edge to be isometric to $[0, 1]$.

Remark 2.1.18. This matches with the combinatorial definition of a tree as a graph without any loops. Some definitions allow a simplicial tree to be any space homeomorphic to such a simplicial complex, however in this thesis we will use the more restricted definition above.

Example 2.1.19 below shows that simplicial trees are a strict subclass of $\mathbb{R}$-trees, even if we used the broader definition of simplicial trees given in the above remark. On the other hand, in Section 2.3.2 we will see that for many of our purposes simplicial trees and $\mathbb{R}$-trees are pretty much interchangeable.

Example 2.1.19. The comb space $(\{0\} \times [0,1]) \cup \bigcup_{n=1}^{\infty}(\{\frac{1}{n}\} \times [0,1]) \cup ([0,1] \times \{0\})$ (see Figure 2.1), is a subset of $\mathbb{R}^2$ which is an $\mathbb{R}$-tree under the shortest path metric, but can never be realised as a simplicial complex.

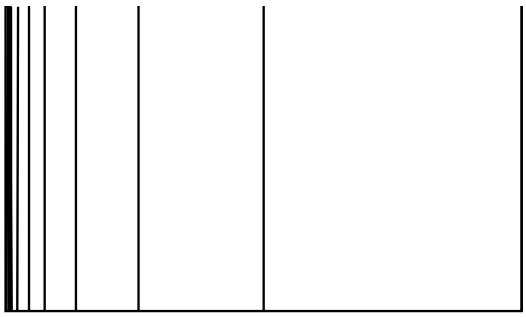


Figure 2.1: The comb space

In this chapter we will not be focusing on studying trees, but instead on using trees to study spaces that are in some way roughly like trees. The following definition will allow us to make this notion precise.

Definition 2.1.20. Let (X, d_X) and (Y, d_Y) be metric spaces. A map $f : X \to Y$ is an (L, C)-*quasi-isometry* if there exist constants $L \geqslant 1$, $C \geqslant 0$ such that:

(1) For every $x_1, x_2 \in X$ we have that

$$\frac{1}{L} d_X(x_1, x_2) - C \leqslant d_Y(f(x_1), f(x_2)) \leqslant L d_X(x_1, x_2) + C.$$

(2) For every $y \in Y$ there exists $x \in X$ such that $d_Y(f(x), y) \leqslant C$.

These two properties can be described as being coarsely bi-Lipschitz and coarsely surjective, respectively. If only the first property is satisfied then the map is an (L, C)-*quasi-isometric embedding*.

Remark 2.1.21. A $(1, C)$-quasi-isometry is sometimes known as a *rough isometry*, amongst other terms.

One well known property of quasi-isometries is that given a quasi-isometry from X to Y, we can also construct a quasi-isometry from Y to X, with related constants.

Lemma 2.1.22. *Let (X, d_X) and (Y, d_Y) be metric spaces. Suppose there exists an (L, C)-quasi-isometry $f : X \to Y$ for some $L \geqslant 1, C \geqslant 0$. Then there exists an $(L, 3LC)$-quasi-isometry $g : Y \to X$.*

Proof. By the coarse surjectivity of f, for any $y \in Y$ there exists $x \in X$ such that $d_Y(f(x), y) \leqslant C$. We therefore pick some x in the preimage of the closed ball $B(y, C)$ under f, and let this choice be $g(y)$. This defines a function $g : Y \to X$.

We want to check that g is coarsely bi-Lipschitz and coarsely surjective. We begin with coarse surjectivity. Let $x \in X$, and consider $g(f(x))$. By our definition of g, we have that $d_Y(f(x), f(g(f(x)))) \leqslant C$. Hence

$$d_X(x, g(f(x))) \leqslant L d_Y(f(x), f(g(f(x)))) + LC \leqslant 2LC,$$

so g is coarsely surjective.

We now prove that g is coarsely bi-Lipschitz. Let $y_1, y_2 \in Y$. By our definition of g and the triangle inequality we can see that

$$d_Y(y_1, y_2) - 2C \leqslant d_Y(f(g(y_1)), f(g(y_2))) \leqslant d_Y(y_1, y_2) + 2C.$$

Therefore

$$d_X(g(y_1), g(y_2)) \leqslant L d_Y(f(g(y_1)), f(g(y_2))) + LC \leqslant L d_Y(y_1, y_2) + 3LC,$$

and in the other direction

$$d_X(g(y_1), g(y_2)) \geqslant \frac{1}{L} d_Y(f(g(y_1)), f(g(y_2))) - \frac{C}{L} \geqslant \frac{1}{L} d_Y(y_1, y_2) - \frac{3C}{L}.$$

As $L \geqslant 1$ and $C \geqslant 0$, we have that $\frac{3C}{L} \leqslant 3LC$, so g is an $(L, 3LC)$-quasi-isometry. □

Convention. When L and C are not given explicitly, Lemma 2.1.22 allows us to assume that they are large enough to hold in both directions.

It is also well known that the composition of two quasi-isometries is a quasi-isometry.

Lemma 2.1.23. *Let (X, d_X), (Y, d_Y), and (Z, d_Z) be metric spaces. Suppose there exists an (L, C)-quasi-isometry $f : X \to Y$, and an (L', C')-quasi-isometry $g : Y \to Z$, for some $L, L' \geqslant 1$, $C, C' \geqslant 0$. Then $g \circ f$ is an $(LL', L'C + 2C')$-quasi-isometry.*

Proof. We first prove that $g \circ f$ is coarsely bi-Lipschitz. Let $x, x' \in X$. Then we have that

$$d_Z(g(f(x_1)), g(f(x_2))) \leqslant L' d_Y(f(x_1), f(x_2)) + C' \leqslant LL' d_X(x_1, x_2) + L'C + C',$$

and similarly

$$d_Z(g(f(x_1)), g(f(x_2))) \geqslant \frac{1}{L'} d_Y(f(x_1), f(x_2)) - C' \leqslant \frac{1}{LL'} d_X(x_1, x_2) - \frac{C}{L'} - C'.$$

We now want to show that $g \circ f$ is coarsely surjective. Let $z \in Z$, then there exists $y \in Y$ such that $d_Z(g(y), z) \leqslant C'$, and there exists $x \in X$ such that $d_Y(f(x), y) \leqslant C$. We can see that

$$d_Z(g(f(x)), g(y)) \leqslant L' d_Y(f(x), y) + C' \leqslant L'C + C',$$

and so by the triangle inequality we get that

$$d_Z(g(f(x)), z) \leqslant L'C + 2C'.$$

Therefore $g \circ f$ is an $(LL', L'C + 2C')$-quasi-isometry from X to Z. □

Definition 2.1.24. Let X and Y be metric spaces. We say that X is (L, C)-*quasi-isometric* to Y for some constants $L \geqslant 1, C \geqslant 0$ if there exists an (L, C)-quasi-isometry $f : X \to Y$. When the constants L and C are not specified we will refer to f simply as a quasi-isometry, and say that X and Y are quasi-isometric.

Remark 2.1.25. Lemma 2.1.22 and Lemma 2.1.23 show that metric spaces being quasi-isometric is an equivalence relation.

Remark 2.1.26. Using the above terminology, Lemma 2.1.22 tells us that if X is (L, C)-quasi-isometric to Y, then Y is $(L, 3LC)$-quasi-isometric to X.

We are now able to formalise our notion of spaces that are roughly like trees.

Definition 2.1.27. A *quasi-tree* is a geodesic metric space that is quasi-isometric to a simplicial tree.

Remark 2.1.28. Any quasi-tree will automatically be δ-hyperbolic for some $\delta \geqslant 0$, as hyperbolicity is a quasi-isometry invariant [43].

A highly useful alternative characterisation of quasi-trees was given by Manning.

Theorem 2.1.29 (Manning's bottleneck criterion)**.** [65] *A geodesic metric space (X, d) is a quasi-tree if and only if there exists $\Delta \geqslant 0$ such that for every geodesic $[x, y]$ in X with midpoint m, every path between x and y intersects the closed ball $B(m, \Delta)$.*

Definition 2.1.30. The constant Δ in the above theorem is known as the *bottleneck constant* for the quasi-tree X.

It is well known that this characterisation can be easily extended to the equivalent statement that every point on a geodesic in a quasi-tree has this property, not just the midpoint (see, for example, [9]).

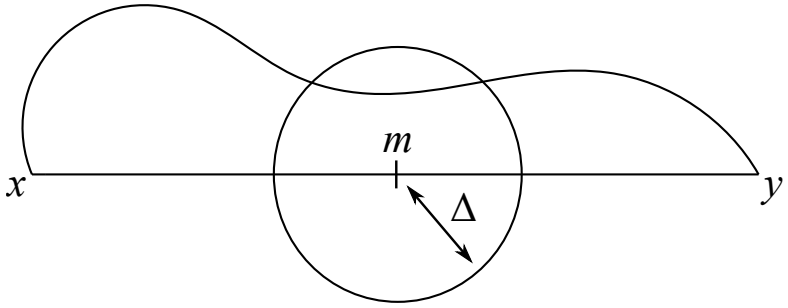


Figure 2.2: Manning's bottleneck criterion

Corollary 2.1.31. *A geodesic metric space (X, d) is a quasi-tree if and only if there exists $\Delta \geqslant 0$ such that for every geodesic $[x, y]$ in X, and every $z \in [x, y]$, every path between x and y intersects the closed ball $B(z, \Delta)$.*

Proof. Suppose X is a quasi-tree, and let $\gamma : [0, 1] \to X$ be a path from x to y. If z is the midpoint of of $[x, y]$, then we are done by Manning's bottleneck criterion. If z is not the midpoint of $[x, y]$, then assume without loss of generality that $d(x, z) < d(z, y)$, and let $y' \in [z, y]$ be the unique point such that $d(x, z) = d(z, y')$. We then will have that z is the midpoint of $[x, y'] \subset [x, y]$.

As $\gamma \cup [y, y']$ is a path from x to y', by Manning's bottleneck criterion it must intersect $B(z, \Delta)$. If γ intersects $B(z, \Delta)$ then we are done. If instead $[y, y']$ intersects $B(z, \Delta)$, then we get that $d(z, y') \leqslant \Delta$. This means that $d(x, z) \leqslant \Delta$, and therefore we again have that γ intersects $B(z, \Delta)$.

The reverse implication is true directly from Manning's bottleneck criterion. □

Remark 2.1.32. It follows from Corollary 2.1.31 that a quasi-tree with bottleneck constant $\Delta \geqslant 0$ is Δ-hyperbolic.

2.2 End-approximating tree for a geodesic metric space

In this section we will give a construction of an $\mathbb{R}$-tree from a general geodesic metric space, in such a way that some of the original structure is retained. The idea is based on Gromov's proof of the Tree Approximation Lemma in [43]. A similar explanation of such a construction can be found in [26], however we state the results here with greater generality.

Definition 2.2.1. Let (X, d) be a metric space. For $x, y \in X$, we will use the notation $S_{x,y}$ to mean the set of all finite sequences in X between x and y, so

$$S_{x,y} = \{x_1, \ldots, x_n \in X : x_1 = x, x_n = y, n \in \mathbb{N}\}.$$

Fix $x_0 \in X$. We define $(x, y)'_{x_0}$ to be

$$(x, y)'_{x_0} = \sup_{S_{x,y}} \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0}.$$

We finally define d' on X as

$$d'(x, y) = d(x_0, x) + d(x_0, y) - 2(x, y)'_{x_0}.$$

We begin with a couple of preliminary lemmas.

Lemma 2.2.2. *Let (X, d) be a metric space, and let $x_0, x, y \in X$. Then $(x, y)'_{x_0} \leqslant d(x_0, x)$.*

Proof. For $x, y \in X$ we can see that the triangle inequality for d implies that $(x, y)_{x_0} \leqslant d(x_0, x)$. Hence for any sequence $x = x_1, \ldots, x_n = y$ in $S_{x,y}$ we have that

$$\min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0} \leqslant (x_1, x_2)_{x_0} = (x, x_2)_{x_0} \leqslant d(x_0, x),$$

so $(x, y)'_{x_0} \leqslant d(x_0, x)$. □

Lemma 2.2.3. *Let (X, d) be a metric space, and let $x_0, x, y \in X$. Then we have that $(x, z)'_{x_0} \geqslant \min\{(x, y)'_{x_0}, (y, z)'_{x_0}\}$.*

Proof. This follows by concatenating sequences. □

We can use these lemmas to prove the following.

Lemma 2.2.4. *The function $d' : X \times X \to \mathbb{R}_{\geqslant 0}$ is a pseudometric on X, with the property that $d' \leqslant d$.*

Proof. We first need to check that d' actually maps to $\mathbb{R}_{\geqslant 0}$. Lemma 2.2.2 tells us that $(x, y)'_{x_0} \leqslant d(x_0, x)$, and similarly that $(x, y)'_{x_0} \leqslant d(x_0, y)$. Hence we have that $d'(x, y) \geqslant 0$. The fact that $d'(x, y) < \infty$ is clear.

We now want to show that $d'(x, x) = 0$. Note that x, y is a valid sequence in $S_{x,y}$, so $(x, y)'_{x_0} \geqslant (x, y)_{x_0}$, and therefore $d'(x, y) \leqslant d(x, y)$. Hence $0 \leqslant d'(x, x) \leqslant d(x, x) = 0$.

It is clear from the symmetry of the Gromov product for d that $(x, y)'_{x_0} = (y, x)'_{x_0}$, and hence clear that $d'(x, y) = d'(y, x)$.

Finally, let $x, y, z \in X$. We can see that

$$\begin{aligned} d'(x,z) \leqslant d'(x,y) + d'(y,z) &\iff -2(x,z)'_{x_0} \leqslant 2d(x_0,y) - 2((x,y)'_{x_0} + (y,z)'_{x_0}) \\ &\iff 0 \leqslant d(x_0,y) + (x,z)'_{x_0} - (x,y)'_{x_0} - (y,z)'_{x_0}. \end{aligned}$$

Lemma 2.2.3 tells us that

$$(x,z)'_{x_0} \geqslant \min\{(x,y)'_{x_0}, (y,z)'_{x_0}\},$$

so suppose without loss of generality that $(x,z)'_{x_0} \geqslant (x,y)'_{x_0}$. Then as $d(x_0,y) \geqslant (y,z)'_{x_0}$, we have shown that d' satisfies the triangle inequality. Therefore d' is a pseudometric on X. □

It is not hard to see that this is not a metric in general, as it is possible to have $x \neq y$ but $d'(x,y) = 0$. One example would be when $X = \mathbb{R}^n$ for some $n \geqslant 2$, and x and y lie on the same sphere around x_0 (see Lemma 2.5.1).

Definition 2.2.5. Let (X,d) be a metric space, and fix $x_0 \in X$. We define (T_X, d^*) to be the metric space induced as a quotient of the pseudometric space (X,d').

Specifically, we consider the equivalence relation $\sim$, where $x \sim y$ if and only if $d'(x,y) = 0$. We let $T_X = X/\sim$, and $d^*([x],[y]) = d'(x,y)$. We let $([x],[y])^*_{[x_0]}$ be the Gromov product under the metric d^*.

Remark 2.2.6. We will consider the choice of basepoint to be fixed. Although the induced space may vary depending on this choice, it will not affect any of the properties that we consider here.

Lemma 2.2.7. *Let (X,d) be a metric space. The metric space (T_X, d^*) is 0-hyperbolic.*

Proof. We first note that for any $x \in X$ we have that $0 \leqslant (x_0,x)'_{x_0} \leqslant d(x_0,x_0) = 0$ by Lemma 2.2.2, so in particular $d'(x_0,x) = d(x_0,x)$. We can now see that for any $x, y \in X$ we have that

$$\begin{aligned} (x,y)'_{x_0} &= \frac{1}{2}(d(x_0,x) + d(x_0,y) - d'(x,y)) \\ &= \frac{1}{2}(d'(x_0,x) + d'(x_0,y) - d'(x,y)) \\ &= \frac{1}{2}(d^*([x_0],[x]) + d^*([x_0],[y]) - d^*([x],[y])) \\ &= ([x],[y])^*_{[x_0]}, \end{aligned}$$

the Gromov product for the metric d^*. For any $x, y, z \in X$ we know from Lemma 2.2.3 that

$$(x, z)'_{x_0} \geqslant \min\{(x, y)'_{x_0}, (y, z)'_{x_0}\},$$

and therefore

$$([x], [z])^*_{[x_0]} \geqslant \min\{([x], [y])^*_{[x_0]}, ([y], [z])^*_{[x_0]}\},$$

so T_X is 0-hyperbolic. □

To get our desired $\mathbb{R}$-tree we will need the additional assumption that the space X is geodesic.

Proposition 2.2.8. *Let (X, d) be a geodesic metric space. The metric space (T_X, d^*) is an $\mathbb{R}$-tree.*

Proof. We know that T_X is 0-hyperbolic by Lemma 2.2.7. It therefore only remains to show that T_X is geodesic.

Let $f : X \to T_X$ be the quotient map $f(x) = [x]$. As $d^*([x], [y]) \leqslant d(x, y)$ it is obvious that f is continuous, and moreover we can show that it is an isometry on any geodesic ray starting from x_0. Let $y \in X$, and let $x \in [x_0, y]$. We can note that

$$(x, y)_{x_0} \leqslant (x, y)'_{x_0} \leqslant d(x_0, x) = (x, y)_{x_0},$$

so $(x, y)_{x_0} = (x, y)'_{x_0}$. Therefore

$$d(x, y) = d'(x, y) = d^*([x], [y]).$$

In particular, the image of $[x, y]$ will be a geodesic between $[x]$ and $[y]$ in T_X.

Now let $[x], [y] \in T_X$ be arbitrary. We want to construct a path in T_X between $[x]$ and $[y]$ that has length $d^*([x], [y])$. Consider representatives x and y of these equivalence classes in X. We consider geodesics $[x_0, x]$ and $[x_0, y]$ in X, and pick the unique points $x' \in [x_0, x]$ and $y' \in [x_0, y]$ such that

$$d(x_0, x') = d(x_0, y') = (x, y)'_{x_0}.$$

This means that

$$d(x, x') + d(y', y) = d(x_0, x) + d(x_0, y) - 2(x, y)'_{x_0} = d'(x, y),$$

so $d^*([x], [x']) + d^*([y'], [y]) = d^*([x], [y])$. Hence if we can show that $[x'] = [y']$ then we can simply take the union of the geodesics $[[x], [x']]$ and $[[y'], [y]]$ to find a geodesic between $[x]$ and $[y]$.

We have that

$$d(x_0, x') = d(x_0, y') = ([x], [y])^*_{[x_0]} = (x, y)'_{x_0} = \sup_{S_{x,y}} \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0}.$$

Recall that this supremum is over all finite sequences in X beginning with x and ending with y. Given that

$$(x, x')_{x_0} = d(x_0, x') = d(x_0, y') = (y, y')_{x_0},$$

in this case we would get the same result if we took the supremum over all finite sequences beginning with x, x', x and ending with y, y', y. This in turn is equivalent to taking the supremum over all finite sequences beginning with x', x and ending with y, y', which is no larger than taking the supremum over all sequences beginning with x' and ending with y', which is $S_{x',y'}$. Therefore

$$d(x_0, x') = d(x_0, y') \leqslant (x', y')'_{x_0} \leqslant d(x_0, x') = d(x_0, y'),$$

so we have equality, and in particular $d'(x', y') = 0$, so $[x'] = [y']$. Hence T_X is a geodesic 0-hyperbolic space, therefore it is an $\mathbb{R}$-tree. □

Informally, we construct T_X from a geodesic metric space X by collapsing X along the spheres $S(x_0, r)$, where two points in $S(x_0, r)$ collapse to the same point in T_X if and only if for every $\varepsilon > 0$ there exists a path γ between them in X such that $d(x_0, \gamma) \geqslant r - \varepsilon$. Alternatively, we can view T_X as the space that we get if we collapse X along its ends.

This idea will be formalised in Section 2.5. It is not necessary for the proofs in Sections 2.3 or 2.4, however it may be a useful image to keep in mind, and it also motivates the language we use to describe T_X.

Definition 2.2.9. Let (X, d) be a geodesic metric space. We call (T_X, d^*) the *end-approximating tree* of (X, d), and $f : X \to T_X$ defined by $f(x) = [x]$ the *end-approximating map*.

Remark 2.2.10. Lemma 2.2.4 tells us that the end-approximating map is non-expanding.

We will sometimes use slightly different terminology for the end-approximating tree of a quasi-tree.

Convention. When (X, d) is a quasi-tree, we will occasionally simply refer to (T_X, d^*) as an *approximating tree*. This terminology is justified in Subsection 2.3.1.

2.3 Quasi-isometries with trees

For a general geodesic metric space X, the end-approximating tree T_X may bear very little resemblance to X. For example, if $X = \mathbb{H}^n$ or $\mathbb{R}^n$ for some $n \geqslant 2$ then T_X is isometric to $[0, \infty)$ (see Figure 2.3 and Corollary 2.5.2). On the other hand, when X is a quasi-tree we might expect T_X to look much more like the original space (see Figure 2.4). This makes sense intuitively, as in a quasi-tree we would expect there to be a limit on how far points can collapse together along spheres. We will show that in this case the end-approximating map $f : X \to T_X$ is in fact a $(1, C)$-quasi-isometry, and then use this to find a $(1, C')$-quasi-isometry from X to a simplicial tree.

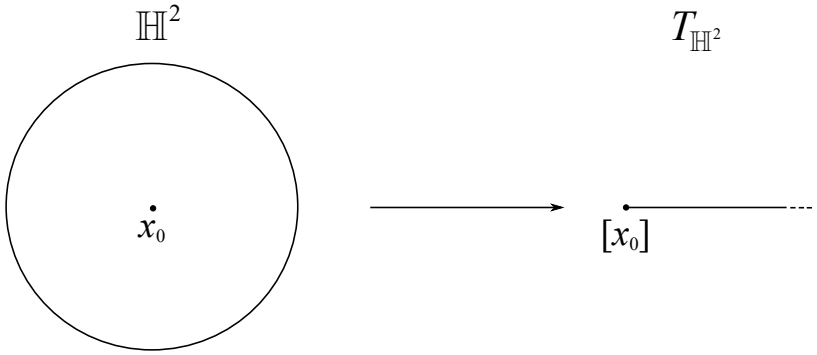


Figure 2.3: End-approximating tree of the hyperbolic plane

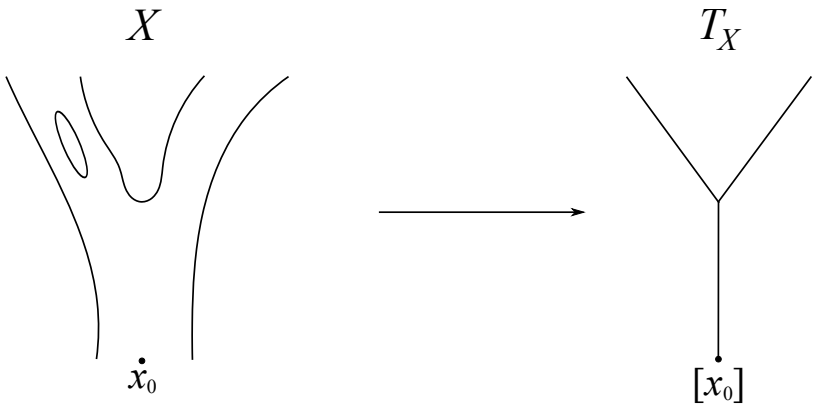


Figure 2.4: Approximating tree of a quasi-tree

2.3.1 $\mathbb{R}$-trees

In this subsection we will prove that if X is a quasi-tree then there exists some $C \geqslant 0$ such that the end-approximating map $f : X \to T_X$ is $(1, C)$-quasi-isometry. To

achieve this we must first prove a uniform version of Lemma 2.1.6 for quasi-trees, for which we will need Manning's bottleneck criterion.

Proposition 2.3.1. *A geodesic metric space (X, d) is a quasi-tree if and only if there exists $A \geqslant 0$ such that any $x_0, x_1, \ldots, x_n \in X$ satisfy*

$$(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0} - A.$$

Proof. Suppose X is a δ-hyperbolic quasi-tree, with bottleneck constant $\Delta \geqslant 0$. Let $x_0, x_1, \ldots, x_n \in X$ be arbitrary, and consider a geodesic $[x_1, x_n]$. By Lemma 2.1.10 there exists $z \in [x_1, x_n]$ such that $d(x_0, z) - 2\delta \leqslant (x_1, x_n)_{x_0}$.

Consider the path $[x_1, x_2] \cup \cdots \cup [x_{n-1}, x_n]$. By Corollary 2.1.31, the extended version of Manning's bottleneck criterion, for some $1 \leqslant i \leqslant n-1$ there exists $z' \in [x_i, x_{i+1}]$ such that $z' \in B(z, \Delta)$, and $(x_i, x_{i+1})_{x_0} \leqslant d(x_0, z')$ by Lemma 2.1.11. We can put this together to get that

$$(x_i, x_{i+1})_{x_0} \leqslant d(x_0, z') \leqslant d(x_0, z) + d(z, z') \leqslant (x_1, x_n)_{x_0} + \Delta + 2\delta.$$

Hence $(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1}(x_i, x_{i+1})_{x_0} - (\Delta + 2\delta)$. This concludes the proof in one direction.

Now suppose that (X, d) is a geodesic metric space, and that there exists $A \geqslant 0$ such that any $x_0, x_1, \ldots, x_n \in X$ satisfy $(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1}(x_i, x_{i+1})_{x_0} - A$. We want to show that X satisfies Manning's bottleneck criterion.

Let $x, y \in X$, with $[x, y]$ a geodesic between them with midpoint m, and let γ be any path from x to y. Choose $x = x_1, \ldots, x_n = y$ such that $x_i \in \gamma$ for all $1 \leqslant i \leqslant n$ and $d(x_i, x_{i+1}) \leqslant 2A$ for all $1 \leqslant i \leqslant n-1$.

Note that $(x_1, x_n)_m = (x, y)_m = 0$. We therefore have that

$$\min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_m - A \leqslant (x_1, x_n)_m \implies \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_m \leqslant A,$$

so for some $1 \leqslant i \leqslant n-1$ we have that

$$\begin{aligned}(x_i, x_{i+1})_m \leqslant A &\implies \frac{1}{2}(d(m, x_i) + d(m, x_{i+1}) - d(x_i, x_{i+1})) \leqslant A \\ &\implies d(m, x_i) + d(m, x_{i+1}) \leqslant 4A.\end{aligned}$$

Hence there exists $1 \leqslant i \leqslant n$ such that $d(m, x_i) \leqslant 2A$. As $x_i \in \gamma$, we have that there exists $z \in \gamma$ such that $d(m, z) \leqslant 2A$. By Manning's bottleneck criterion X is therefore a quasi-tree. □

The key difference here between the case of a quasi-tree and the case of, say, the hyperbolic plane $\mathbb{H}^2$, is that given any three points $x_0, x, y \in \mathbb{H}^2$, we can pick a large enough $n \in \mathbb{N}$ and a sequence of points $x = x_1, \ldots, x_n = y$ in $\mathbb{H}^2$ such that $\min_{1 \leqslant i \leqslant n-1}(x_i, x_{i+1})_{x_0}$ is arbitrarily close to $\min\{d(x_0, x), d(x_0, y)\}$. This is why $T_{\mathbb{H}^2}$ is the half-line, as illustrated by Figure 2.3.

In a quasi-tree X, however, given any sequence of points $x = x_1, \ldots, x_n = y$ in X, and assuming that $(x, y)_{x_0}$ is not approximately $\min\{d(x_0, x), d(x_0, y)\}$, there must be some i for which x_i and x_{i+1} lie in different branches of the quasi-tree. Here by branches we mean connected components of $X \backslash B(x_0, r)$, for r approximately $(x, y)_{x_0}$. This idea will be come up again in the proof of Lemma 2.5.3.

We can now combine Proposition 2.3.1 with the construction of the end-approximating tree to show that for every quasi-tree the end-approximating map is a $(1, C)$-quasi-isometry.

Proposition 2.3.2. *Let (X, d) be a quasi-tree with bottleneck constant $\Delta \geqslant 0$ and hyperbolicity constant $\delta \geqslant 0$. Then the end-approximating map $f : X \to T_X$ is a $(1, 2(\Delta + 2\delta))$-quasi-isometry.*

Proof. We first note that the end-approximating map is surjective, and hence is coarsely surjective. Now recall that for $[x], [y] \in T_X$, the distance between them is

$$d^*([x], [y]) = d'(x, y) = d(x_0, x) + d(x_0, y) - 2(x, y)'_{x_0},$$

where

$$(x, y)'_{x_0} = \sup_{S_{x,y}} \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0}.$$

We know that $(x, y)_{x_0} \leqslant (x, y)'_{x_0}$. By Proposition 2.3.1 we also obtain that

$$(x, y)'_{x_0} \leqslant (x, y)_{x_0} + \Delta + 2\delta.$$

Putting these inequalities together and substituting them into the definition of $d'(x, y)$ gives us that

$$d(x, y) - 2(\Delta + 2\delta) \leqslant d'(x, y) \leqslant d(x, y),$$

and so

$$d(x, y) - 2(\Delta + 2\delta) \leqslant d^*([x], [y]) \leqslant d(x, y).$$

Hence the function f is a $(1, 2(\Delta + 2\delta))$-quasi-isometry from X to T_X. □

Corollary 2.3.3. *For every quasi-tree X, there exists an $\mathbb{R}$-tree T and a constant $C \geqslant 0$ such that X is $(1, C)$-quasi-isometric to T.*

2.3.2 Simplicial trees

We note here that T_X will certainly not be a simplicial tree in general. This can be seen when X is an $\mathbb{R}$-tree, since then $T_X = X$, so if X is not simplicial then T_X will not be either (see Remark 2.5.7). Even in the relatively natural case that our quasi-tree X is a graph and the basepoint $x_0 \in X$ is a vertex, the $\mathbb{R}$-tree T_X may fail to be simplicial with edge lengths 1. However it is not hard to see that we can substitute our $\mathbb{R}$-tree for a simplicial tree with only an additive error in the quasi-isometry. The simplicial tree we construct here comes from the technique Manning uses in the proof of the bottleneck criterion [65].

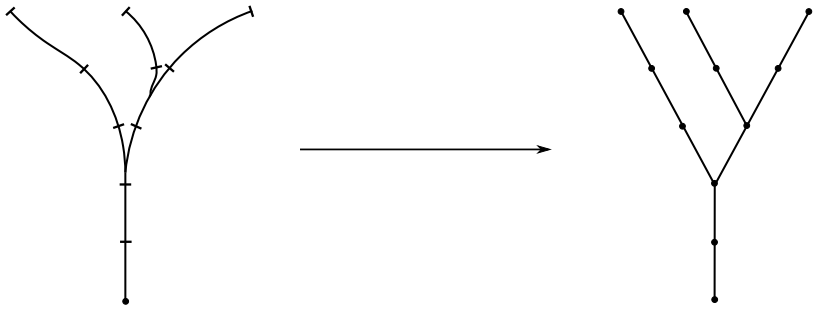


Figure 2.5: Constructing a simplicial tree from an $\mathbb{R}$-tree

Let T be an $\mathbb{R}$-tree. We begin our construction by picking a basepoint $x_0 \in T$, and letting $T_0 = \{x_0\}$. We then let $\Gamma_0 = (V_0, E_0)$ be a graph with $V_0 = \{v_{x_0}\}$ and $E_0 = \emptyset$.

Given $\Gamma_{k-1} = (V_{k-1}, E_{k-1})$, we construct a new graph Γ_k by adding the next layer of vertices. Let $T_k = S(x_0, k)$ be the sphere around x_0 in T of radius k, and then let $U_k = \{v_x : x \in T_k\}$. We let the vertex set of Γ_k be $V_k = V_{k-1} \cup U_k$, and the edge set be $E_k = E_{k-1} \cup \{\{v_x, v_y\} : v_x \in U_{k-1}, v_y \in U_k, x \in [x_0, y]\}$.

If Γ_{k-1} is a simplicial tree, then so is Γ_k. This is due to the fact that T is an $\mathbb{R}$-tree, so for every $y \in T_k$ there is a unique $x \in T_{k-1}$ such that $x \in [x_0, y]$, and this rules out the existence of any cycles. We now let $\Gamma = (V, E)$, where $V = \bigcup_k V_k$ and $E = \bigcup_k E_k$. The graph Γ is also a simplicial tree, as any cycle in Γ would have to exist in some subgraph Γ_k.

We define a map $\psi : T \to \Gamma$ as follows. For every $x \in T$, let $x' \in T$ be the unique point such that $x' \in [x_0, x]$ and $d_T(x_0, x') = \lfloor d_T(x_0, x) \rfloor$. We then set $\psi(x) = v_{x'}$.

Proposition 2.3.4. *Every $\mathbb{R}$-tree is $(1, 2)$-quasi-isometric to a simplicial tree.*

Proof. Let T be an $\mathbb{R}$-tree, and let Γ be the simplicial tree constructed above. We will show that the map $\psi : T \to \Gamma$ is a $(1, 2)$-quasi-isometry.

We will first show that ψ is coarsely surjective. Let $w \in \Gamma$. Then there exists a vertex $v \in \Gamma$ such that $d_\Gamma(w, v) \leqslant 1$. By the definition of Γ there exists some $x \in T$ such that $v_x = v$. We therefore have that $\psi(x) = v$, so we are done.

We will now show that ψ is coarsely bi-Lipschitz. Let $x, y \in T$. Suppose first that $x \in [x_0, y]$. Then $\psi(x) \in [v_{x_0}, \psi(y)]$, so $d_\Gamma(\psi(x), \psi(y)) = \lfloor d_T(x_0, y) \rfloor - \lfloor d_T(x_0, x) \rfloor$. As $d_T(x, y) = d_T(x_0, y) - d_T(x_0, x)$, this means that

$$d_T(x, y) - 1 \leqslant d_\Gamma(\psi(x), \psi(y)) \leqslant d_T(x, y) + 1.$$

Now suppose $x, y \in T$ are any two points. Let $z \in [x_0, x] \cap [x_0, y]$ be the unique point such that $d_T(x_0, z) = d_T(x_0, [x, y])$. Recall that $z' \in T$ is the unique point such that $z' \in [x_0, z]$ and $d_T(x_0, z') = \lfloor d_T(x_0, z) \rfloor$, so $[x_0, x] \cap [x_0, y] \cap S(x_0, d_T(x_0, z') + 1)$ is empty. this means that $v_{z'} \in [\psi(x), \psi(y)]$, and so

$$d_\Gamma(\psi(x), \psi(y)) = d_\Gamma(\psi(x), \psi(z)) + d_\Gamma(\psi(z), \psi(y)) = d_T(x, z') + d_T(z', y).$$

As

$$\begin{aligned} d_T(x, y) - 2 &= d_T(x, z) + d_T(z, y) - 2 \\ &\leqslant d_T(x, z') + d_T(z', y) \\ &\leqslant d_T(x, z) + d_T(z, y) + 2 \\ &= d_T(x, y) + 2, \end{aligned}$$

we can conclude that ψ is a (1,2)-quasi-isometry from the $\mathbb{R}$-tree T to the simplicial tree Γ. □

We noted in Section 2.1 that this means there is a (1,6)-quasi-isometry from Γ to T, although it is in fact possible to improve this to show that Γ is (1,2)-quasi-isometric to T. This can be proved via the map $\varphi : \Gamma \to T$, defined such that each edge $\{v_x, v_y\}$ is mapped isometrically to the geodesic $[x, y]$ in T. We provide a proof of this below, as φ will be used in the proof of Proposition 2.3.25.

Proposition 2.3.5. *For every $\mathbb{R}$-tree T there exists a simplicial tree that is $(1, 2)$-quasi-isometric to T.*

Proof. Let Γ and φ be as above. We will show that φ is a $(1, 2)$-quasi-isometry.

We will first show that φ is coarsely surjective. Let $x \in T$, and let $k = \lfloor d(x_0, x) \rfloor$. Then there exists some $y \in T_k$ such that $d_T(x, y) < 1$, so as y is the image of v_y under φ we therefore have that φ is coarsely surjective.

We will now show that φ is coarsely bi-Lipschitz. Note that it is immediate from the definition of Γ and φ that for any $x \in \Gamma$, the geodesic $[v_{x_0}, x]$ is mapped isometrically by φ onto the geodesic $[x_0, \varphi(x)]$ in T.

Let $x, y \in \Gamma$ be points that are not necessarily vertices, and consider the geodesic $[x, y]$. If $[x, y] \cap V = \emptyset$ then $[x, y]$ is contained in an edge of Γ, and therefore is mapped isometrically onto T by φ. Otherwise, we can write $[x, y]$ as

$$[x, y] = [x, v_1] \cup [v_1, v_2] \cup \cdots \cup [v_{n-1}, v_n] \cup [v_n, y],$$

where every $v_i \in V \cap [x, y]$, and each subsegment has length greater than 0 but no greater than 1, so each $[v_i, v_{i+1}]$ is an edge in Γ.

Suppose that $d_\Gamma(v_{x_0}, [x, y]) = d_\Gamma(v_{x_0}, x)$, so x is the nearest point to v_{x_0}. Then as Γ is a tree we have that $[x, y] \subset [v_{x_0}, y]$, so φ is an isometry on $[x, y]$. If $d_\Gamma(v_{x_0}, [x, y]) = d_\Gamma(v_{x_0}, y)$, then by the same reasoning φ is again an isometry on $[x, y]$.

If on the other hand $d_\Gamma(v_{x_0}, [x, y]) = d_\Gamma(v_{x_0}, v_i)$, then although $[x, v_i]$ and $[v_i, y]$ are mapped isometrically onto $[\varphi(x), \varphi(v_i)]$ and $[\varphi(v_i), \varphi(y)]$, we many have that the geodesics $[\varphi(v_{i-1}), \varphi(v_i)]$ and $[\varphi(v_i), \varphi(v_{i+1})]$ intersect at points other than $\varphi(v_i)$. Note that $\varphi(v_{i-1}) \neq \varphi(v_{i+1})$ as otherwise we would have $v_{i-1} = v_{i+1}$, which would contradict our description of the geodesic $[x, y]$. This means that we can write $[\varphi(x), \varphi(y)]$ as

$$[\varphi(x), \varphi(y)] = [\varphi(x), \varphi(v_{i-1})] \cup [\varphi(v_{i-1}), \varphi(v_{i+1})] \cup [\varphi(v_{i+1}), \varphi(y)].$$

Given we know that $[x, y] = [x, v_{i-1}] \cup [v_{i-1}, v_{i+1}] \cup [v_{i+1}, y]$, along with the fact that $d_\Gamma(x, v_{i-1}) = d_T(\varphi(x), \varphi(v_{i-1}))$ and $d_\Gamma(v_{i+1}, y) = d_T(\varphi(v_{i+1}), \varphi(y))$, it only remains to compare $d_T(\varphi(v_{i-1}), \varphi(v_{i+1}))$ with $d_\Gamma(v_{i-1}, v_{i+1})$.

We can note that $d_T(\varphi(v_{i-1}), \varphi(v_{i+1})) > 0$, and also that

$$d_T(\varphi(v_{i-1}), \varphi(v_{i+1})) \leqslant d_T(\varphi(v_{i-1}), \varphi(v_i)) + d_T(\varphi(v_i), \varphi(v_{i+1})) = 2.$$

As $d_\Gamma(v_{i-1}, v_{i+1}) = 2$, it immediately follows that

$$d_\Gamma(x, y) - 2 \leqslant d_T(\varphi(x), \varphi(y)) \leqslant d_\Gamma(x, y).$$

Therefore φ is a $(1, 2)$-quasi-isometry from the simplicial tree Γ to the $\mathbb{R}$-tree T. $\square$

We now combine Proposition 2.3.2 and Proposition 2.3.4 to get the main result of this section.

Theorem 2.3.6. *A geodesic metric space X is a quasi-tree if and only if it is $(1, C)$-quasi-isometric to a simplicial tree for some $C \geqslant 0$.*

Proof. If a geodesic metric space X is $(1, C)$-quasi-isometric to a simplicial tree then it is a quasi-tree by definition. Now let X be a quasi-tree. By Proposition 2.3.2, X is $(1, C)$-quasi-isometric to the $\mathbb{R}$-tree T_X for a constant $C = 2(\Delta + 2\delta)$, where X is δ-hyperbolic with bottleneck constant Δ. By Proposition 2.3.4 we have that T_X is $(1, 2)$-quasi-isometric to some simplicial tree Γ, so X is $(1, C + 4)$-quasi-isometric to Γ by Lemma 2.1.23. □

Remark 2.3.7. By Proposition 2.3.4 we could equally define quasi-trees as being the geodesic metric spaces that are $(1, C)$-quasi-isometric to $\mathbb{R}$-trees.

We can restate the above theorem in the following way.

Corollary 2.3.8. *A geodesic metric space is (L, C)-quasi-isometric to a simplicial tree for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to a simplicial tree for some $C' \geqslant 0$.*

Remark 2.3.9. Note that these will not usually be the same simplicial tree. This is easily illustrated by the fact that there are simplicial trees that are (L, C)-quasi-isometric for some $L > 1, C \geqslant 0$, but not $(1, C')$-quasi-isometric for any $C' \geqslant 0$.

Remark 2.3.10. As with Remark 2.3.7, we can easily apply Proposition 2.3.4 to substitute either or both of the simplicial trees in Corollary 2.3.8 for an $\mathbb{R}$-tree.

The statement of Corollary 2.3.8 was already known for the special case that the simplicial tree is $\mathbb{R}$, by a result of Manning.

Lemma 2.3.11. [66] *A graph is (L, C)-quasi-isometric to $\mathbb{R}$ for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to $\mathbb{R}$ for some $C' \geqslant 0$.*

One consequence of Theorem 2.3.6 is that it allows us to get rid of the requirement in Manning's result that the space in question is a graph.

Corollary 2.3.12. *A geodesic metric space is (L, C)-quasi-isometric to $\mathbb{R}$ for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to $\mathbb{R}$ for some $C' \geqslant 0$.*

Proof. If a geodesic metric space X is (L, C)-quasi-isometric to $\mathbb{R}$ then it is a quasi-tree, so by Theorem 2.3.6 X is $(1, C')$-quasi-isometric to a simplicial tree. This simplicial tree is quasi-isometric to $\mathbb{R}$, so by Lemma 2.3.11 it is $(1, C'')$-quasi-isometric to $\mathbb{R}$. Therefore X is $(1, C' + 2C'')$-quasi-isometric to $\mathbb{R}$ by Lemma 2.1.23. □

2.3.3 Quasi-actions

We note here that the improvement from Theorem 2.3.6 can also be applied to groups that quasi-act on simplicial trees.

Definition 2.3.13. Let (X, d) be a metric space, and let G be a group. A map $A : G \times X \to X$ is an (L, C)-*quasi-action* of G on X if there exist constants $L \geqslant 1$, $C \geqslant 0$ such that:

(1) For every $g \in G$ we have that the map $A(g, -) : X \to X$ is an (L, C)-quasi-isometry.

(2) For every $g, h \in G$ and $x \in X$ we have that $d(A(g, A(h, x)), A(gh, x)) \leqslant C$.

A quasi-action is *cobounded* if there exists a constant $C' \geqslant 0$ such that for every $x, y \in X$ there exists some $g \in G$ such that $d(A(g, x), y) \leqslant C'$.

A group having a cobounded quasi-action on a simplicial tree is a trivial property, as we could consider the isometric action of any group on the tree which consists of a single vertex. To obtain a non-trivial statement we therefore need a way of comparing quasi-actions.

Definition 2.3.14. Let (X, d_X) and (Y, d_Y) be metric spaces, and let $A_X : G \times X \to X$ and $A_Y : G \times Y \to Y$ be quasi-actions. A map $f : X \to Y$ is *coarsely equivariant* with respect to these quasi-actions if there exists $C \geqslant 0$ such that for every $x \in X$ and $g \in G$ we have that $d_Y(f(A_X(g, x)), A_Y(g, f(x))) \leqslant C$.

If f is also a quasi-isometry, then we call f a *quasi-conjugacy*. If there exists a quasi-conjugacy with respect to a pair of quasi-actions, then we say those quasi-actions are *quasi-conjugate*.

We check here that quasi-actions being quasi-conjugate is an equivalence relation, which is a standard result.

Lemma 2.3.15. *Let (X, d_X) and (Y, d_Y) be metric spaces, let $A_X : G \times X \to X$ and $A_Y : G \times Y \to Y$ be quasi-actions, and let $f : X \to Y$ be a quasi-conjugacy with respect to these quasi-actions. Then the quasi-inverse given by Lemma 2.1.22 is also a quasi-conjugacy with respect to these quasi-actions.*

Proof. Suppose A_X is an (L, C)-quasi-action for some $L \geqslant 1$, $C \geqslant 0$, and suppose f is an (L', C')-quasi-isometry for some $L' \geqslant 1$, $C' \geqslant 0$, and that in addition f is a quasi-conjugacy with constant $C'' \geqslant 0$. Let $f' : Y \to X$ be the function given by

Lemma 2.1.22, so f' is an $(L', 3L'C')$-quasi-isometry, and $d_X(x, f'(f(x)) \leqslant 2L'C'$ for every $x \in X$.

We need to check that f' is coarsely equivariant. Fix $y \in Y$, then let $x = f'(y)$. Note that $d_Y(f(x), y) \leqslant C'$ by definition of f'. For every $g \in G$ we have that

$$\begin{aligned}
d_X(f'(A_Y(g,y)), A_X(g,f'(y))) &\leqslant d_X(f'(A_Y(g,y)), f'(A_Y(g,f(x)))) \\
&\quad + d_X(f'(A_Y(g,f(x))), f'(f(A_X(g,x)))) \\
&\quad + d_X(f'(f(A_X(g,x))), A_X(g,x)) \\
&\quad + d_X(A_X(g,x), A_X(g,f'(y))) \\
&\leqslant L'd_Y(A_Y(g,y), A_Y(g,f(x))) + 3L'C' \\
&\quad + L'd_Y(A_Y(g,f(x)), f(A_X(g,x))) + 3L'C' \\
&\quad + 2L'C' \\
&\leqslant L'(Ld_Y(y,f(x)) + C) + L'C'' + 8L'C' \\
&\leqslant L'(LC' + C + C'' + 8C').
\end{aligned}$$

Hence f' is coarsely equivariant, so f' is a quasi-conjugacy with respect to these quasi-actions. □

Lemma 2.3.16. *Let (X, d_X), (Y, d_Y), and (Z, d_Z) be metric spaces, and let $A_X : G \times X \to X$, $A_Y : G \times Y \to Y$, and $A_Z : G \times Z \to Z$ be quasi-actions. Let $f : X \to Y$ be a quasi-conjugacy with respect to A_X and A_Y, and let $h : Y \to Z$ be a quasi-conjugacy with respect to A_Y and A_Z. Then $h \circ f : X \to Z$ is a quasi-conjugacy with respect to A_X and A_Z.*

Proof. Suppose h is an (L, C)-quasi-isometry for some $L \geqslant 1$, $C \geqslant 0$, and that in addition h is a quasi-conjugacy with constant $C' \geqslant 0$. Suppose f is a quasi-conjugacy with constant $C'' \geqslant 0$. We have that $h \circ f$ is a quasi-isometry by Lemma 2.1.23. It therefore remains to check that $h \circ f$ is coarsely equivariant. Let $x \in X$, then we have that

$$\begin{aligned}
d_Z(h \circ f(A_X(g,x)), A_Z(g, h \circ f(x))) &\leqslant d_Z(h \circ f(A_X(g,x)), h(A_Y(g,f(x)))) \\
&\quad + d_Z(h(A_Y(g,f(x))), A_Z(g, h \circ f(x))) \\
&\leqslant Ld_Y(f(A_X(g,x)), A_Z(g,f(x))) + C + C' \\
&\leqslant LC'' + C + C'.
\end{aligned}$$

Hence $h \circ f$ is coarsely equivariant, so $h \circ f$ is a quasi-conjugacy with respect to these quasi-actions. □

It is a result of Manning that every finitely generated group that has a cobounded quasi-action on a simplicial tree also admits a quasi-conjugate isometric action on a quasi-tree.

Proposition 2.3.17. [66] *Let G be a finitely generated group that has a cobounded quasi-action on a simplicial tree Γ. There exists a (possibly infinite) generating set S for G such that the Cayley graph $Cay(G,S)$ is quasi-isometric to Γ, and the isometric action of G on $Cay(G,S)$ is quasi-conjugate to the quasi-action of G on Γ.*

We therefore obtain the following.

Proposition 2.3.18. *If a finitely generated group admits a cobounded (L,C)-quasi-action on a simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1,C')$-quasi-action on a simplicial tree for some $C' \geqslant 0$.*

Proof. Suppose that G is a finitely generated group which has a cobounded quasi-action on a simplicial tree. Then by Proposition 2.3.17 there exists a generating set S of G such that $Cay(G,S)$ is a quasi-tree. By Theorem 2.3.6 we have that $Cay(G,S)$ is $(1,C')$-quasi-isometric to some simplicial tree Γ. We want to show that the isometric action of G on $Cay(G,S)$ induces a $(1,5C')$-quasi-action on Γ. Let d be the metric on $Cay(G,S)$, and d_Γ be the metric on Γ.

Let $\varphi : Cay(G,S) \to \Gamma$ be a $(1,C')$-quasi-isometry, and let $\psi : \Gamma \to Cay(G,S)$ be the $(1,3C')$-quasi-isometry given by Lemma 2.1.22. For every $g \in G$ we have that the map $\varphi \circ g \circ \psi : \Gamma \to \Gamma$ is a $(1,5C')$-quasi-isometry, by Lemma 2.1.23.

For $x \in \Gamma$ let $A(g,x) = \varphi \circ g \circ \psi(x)$. For every $g,h \in G$ and $x \in X$, we have that

$$\begin{aligned} d_\Gamma(A(g,A(h,x)),A(gh,x)) &= d_\Gamma(\varphi \circ g \circ \psi \circ \varphi \circ h \circ \psi(x), \varphi \circ g \circ h \circ \psi(x)) \\ &\leqslant d(\psi \circ \varphi \circ h \circ \psi(x), h \circ \psi(x)) + C'. \end{aligned}$$

By the definition of ψ from Lemma 2.1.22 we have that $d_\Gamma(\varphi(\psi(y)),y) \leqslant C'$ for every $y \in \Gamma$, so

$$d(\varphi \circ \psi \circ \varphi \circ h \circ \psi(x), \varphi \circ h \circ \psi(x)) \leqslant C',$$

and therefore

$$d(\psi \circ \varphi \circ h \circ \psi(x), h \circ \psi(x)) \leqslant 2C'.$$

We conclude that

$$d_\Gamma(A(g,A(h,x)),A(gh,x)) \leqslant 3C',$$

and so this is a $(1,5C')$-quasi-action.

We now want to check that this quasi-action is quasi-conjugate to the isometric action of G on $\mathrm{Cay}(G,S)$. We already have that $\varphi : Cay(G,S) \to \Gamma$ is a quasi-isometry, and we can also see that

$$\begin{aligned} d_\Gamma(\varphi \circ g(x), A(g,\varphi(x))) &= d_\Gamma(\varphi \circ g(x), \varphi \circ g \circ \psi \circ \varphi(x)) \\ &\leqslant d(x, \psi \circ \varphi(x)) + C' \\ &\leqslant d_\Gamma(\varphi(x), \varphi \circ \psi \circ \varphi(x)) + 2C' \\ &\leqslant 3C'. \end{aligned}$$

Hence this is a quasi-conjugacy. Proposition 2.3.17 tells us that the isometric action of G on $\mathrm{Cay}(G,S)$ is quasi-conjugate to the (L,C)-quasi-action on the original simplicial tree, so by Lemma 2.3.16 we have that these are both quasi-conjugate to the $(1,5C')$-quasi-action of G on Γ. □

As these actions are quasi-conjugate, several properties of the original quasi-action will be inherited by the $(1,C')$-quasi-action. For instance, the $(1,C')$-quasi-action will also be cobounded. This is a known lemma, however we prove it here for completeness.

Lemma 2.3.19. *Let (X,d_X) and (Y,d_Y) be metric spaces, let $A_X : G \times X \to X$ and $A_Y : G \times Y \to Y$ be quasi-actions, and let $f : X \to Y$ be a quasi-conjugacy with respect to these quasi-actions. Then A_X is cobounded if and only if A_Y is cobounded.*

Proof. Suppose f is an (L,C)-quasi-isometry for some $L \geqslant 1$, $C \geqslant 0$, and that in addition f is a quasi-conjugacy with constant $C' \geqslant 0$. Suppose A_Y is cobounded with constant $C'' \geqslant 0$. Let $x, x' \in X$. Then there exists $g \in G$ such that $d_Y(A_Y(g, f(x)), f(x')) \leqslant C''$. We therefore have that

$$\begin{aligned} d_Y(f(A_X(g,x)), f(x')) &\leqslant d_Y(f(A_X(g,x)), A_Y(g,f(x))) + d_Y(A_Y(g,f(x)), f(x')) \\ &\leqslant C' + C''. \end{aligned}$$

It follows that

$$d_X(A_X(g,x), x') \leqslant L d_Y(f(A_X(g,x)), f(x')) + LC \leqslant L(C + C' + C'').$$

This constant is independent of the choice of x and x', so A_X is cobounded.

If we instead suppose that A_X is cobounded, then we can apply the same reasoning using the quasi-conjugacy given by Lemma 2.3.15. □

Another known result is that quasi-conjugacy preserves some of the behaviours of specific group elements in the quasi-action.

Definition 2.3.20. Let (X, d) be a hyperbolic metric space, and let G be a group. Suppose $A : G \times X \to X$ is a quasi-action of G on X. Fix $x \in X$, and for $g \in G$ let $O(g, x) = \{A(g^n, x) : n \in \mathbb{Z}\}$. If $O(g, x)$ has finite diameter, then g quasi-acts *elliptically*. If the map $f : \mathbb{Z} \to O(g, x)$ defined by $f(n) = A(g^n, x)$ is a quasi-isometry, then g quasi-acts *loxodromically*.

Lemma 2.3.21. *Let (X, d) be a hyperbolic metric space, and let G be a group which has a quasi-action on X. The definition of $g \in G$ quasi-acting elliptically or loxodromically is independent of the choice of $x \in X$.*

Proof. Let $A : G \times X \to X$ be an (L, C)-quasi-action of G on X for some $L \geqslant 1$ and $C \geqslant 0$. We can observe that for every $n, m \in \mathbb{Z}$

$$\begin{aligned} d(A(g^n, y), A(g^m, y)) &\leqslant d(A(g^n, y), A(g^n, x)) + d(A(g^n, x), A(g^m, x)) \\ &\quad + d(A(g^m, x), A(g^m, y)) \\ &\leqslant d(A(g^n, x), A(g^m, x)) + 2Ld(x, y) + 2C, \end{aligned}$$

and we also have that

$$\begin{aligned} d(A(g^n, y), A(g^m, y)) &\geqslant d(A(g^n, x), A(g^m, x)) - d(A(g^m, x), A(g^m, y)) \\ &\quad - d(A(g^n, y), A(g^n, x)) \\ &\geqslant d(A(g^n, x), A(g^m, x)) - 2Ld(x, y) - 2C. \end{aligned}$$

Note that $d(x, y)$ is fixed, so it follows from this that $O(g, x)$ has finite diameter if and only if $O(g, y)$ has finite diameter, and that $O(g, x)$ is quasi-isometric to $\mathbb{Z}$ if and only if $O(g, y)$ is quasi-isometric to $\mathbb{Z}$. □

Remark 2.3.22. The concept of elliptic and loxodromic elements will appear again in Chapter 4, in the setting of isometric group actions.

Remark 2.3.23. If G has a quasi-action on a simplicial tree, then for every $g \in G$ we have that g either quasi-acts elliptically or quasi-acts loxodromically. The corresponding dichotomy also holds for isometric actions on quasi-trees [66, 20].

Lemma 2.3.24. *Let (X, d_X) and (Y, d_Y) be hyperbolic metric spaces, let $A_X : G \times X \to X$ and $A_Y : G \times Y \to Y$ be quasi-actions, and let $f : X \to Y$ be a quasi-conjugacy with respect to these quasi-actions. Then $g \in G$ quasi-acts elliptically (respectively loxodromically) on X if and only if g quasi-acts elliptically (respectively loxodromically) on Y.*

Proof. Suppose f is an (L, C)-quasi-isometry for some $L \geqslant 1$, $C \geqslant 0$, and that in addition f is a quasi-conjugacy with constant $C' \geqslant 0$. Let $x \in X$. We can observe that for every $n, m \in \mathbb{Z}$

$$\begin{aligned} d_Y(A_Y(g^n, f(x)), A_Y(g^m, f(x))) &\leqslant d_Y(A_Y(g^n, f(x)), f(A_X(g^n, x))) \\ &\quad + d_Y(f(A_X(g^n, x)), f(A_X(g^m, x))) \\ &\quad + d_Y(f(A_X(g^m, x)), A_Y(g^m, f(x))) \\ &\leqslant L d_X(A_X(g^n, x), A_X(g^m, x)) + C + 2C', \end{aligned}$$

and we also have that

$$\begin{aligned} d_Y(A_Y(g^n, f(x)), A_Y(g^m, f(x))) &\geqslant d_Y(f(A_X(g^n, x)), f(A_X(g^m, x))) \\ &\quad - d_Y(A_Y(g^n, f(x)), f(A_X(g^n, x))) \\ &\quad - d_Y(f(A_X(g^m, x)), A_Y(g^m, f(x))) \\ &\geqslant \frac{1}{L} d_X(A_X(g^n, x), A_X(g^m, x)) - C - 2C'. \end{aligned}$$

If g quasi-acts elliptically on Y, then the second inequality tells us that g also quasi-acts elliptically on X. If g quasi-acts loxodromically on Y, then the two inequalities together tell us that g also quasi-acts loxodromically on X.

If we instead suppose that g quasi-acts elliptically (respectively loxodromically) on X, then we can apply the same reasoning using the quasi-conjugacy given by Lemma 2.3.15. □

2.3.4 Countable and locally finite trees

We note that the definition of a simplicial tree allows a vertex to have uncountably many neighbours. We here give one criterion for when an $\mathbb{R}$-tree is $(1, C)$-quasi-isometric to a countable or locally finite simplicial tree, which will allow us to obtain analogues of Corollary 2.3.8 for countable and locally finite simplicial trees. We will also prove that every proper quasi-tree is $(1, C)$-quasi-isometric to a locally finite simplicial tree, which will be used at the end of Subsection 2.5.2.

Proposition 2.3.25. *Given an $\mathbb{R}$-tree T, we have the following:*

(1) If for some $r \geqslant 1$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has countably many connected components of diameter at least R, then T is $(1, C)$-quasi-isometric to a countable simplicial tree.

(2) If for some $r \geqslant 1$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T\backslash B(x,r)$ has finitely many connected components of diameter at least R, then T is $(1,C)$-quasi-isometric to a locally finite simplicial tree.

Proof. We prove only the first part, as the proof of the second part is analogous. Suppose that T is an $\mathbb{R}$-tree for which for some $r \geqslant 1$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T\backslash B(x,r)$ has countably many connected components of diameter at least R. By Proposition 2.3.5 we know that there exists a simplicial tree Γ with a $(1,2)$-quasi-isometry $\varphi : \Gamma \to T$. Here we take Γ and φ to be exactly the simplicial tree and quasi-isometry that we constructed in that proposition. Let Γ' be obtained from Γ by iteratively removing all leaves $n = \lceil r + R + 4 \rceil$ times. We then have that Γ' is a simplicial tree that is $(1,2n)$-quasi-isometric to Γ, and so is $(1,2n+4)$-quasi-isometric to T by Lemma 2.1.23. We want to show that Γ' is countable.

Let $v \in \Gamma'$ be a vertex. It suffices to show that its set of neighbours $S(v,1)$ in Γ' is countable. Let $w \in S(v,1)$, then w is also a neighbour of v in Γ, and moreover there exists $z_w \in \Gamma$ such that $d_\Gamma(v, z_w) = n$ and $w \in [v, z_w]$.

Recall that Γ is constructed with respect to a basepoint $x_0 \in T$, with v_{x_0} being the corresponding vertex in Γ such that $\varphi(v_{x_0}) = x_0$. There is at most one $\widehat{w} \in S(v,1)$ such that $\widehat{w} \in [v_{x_0}, v]$ in Γ, so $S(v,1)$ is countable if and only if $S(v,1)\backslash\{\widehat{w}\}$ is countable.

Note that if $w, w' \in S(v,1)\backslash\{\widehat{w}\}$ and $w \neq w'$, then $z_w \neq z_{w'}$ as Γ does not contain any cycles. We can also see that $d_\Gamma(z_w, z_{w'}) = 2n$, so $d_T(\varphi(z_w), \varphi(z_{w'})) \geqslant 2n - 2$. As $v \in [v_{x_0}, z_w]$, by the construction of φ in Proposition 2.3.5 we get that $d_T(\varphi(v), \varphi(z_w)) = n$. Similarly $d_T(\varphi(v), \varphi(z_{w'})) = n$. We therefore get that

$$(\varphi(z_w), \varphi(z_{w'}))_{\varphi(v)} \leqslant \frac{1}{2}(2n - (2n-2)) = 1,$$

so $\varphi(z_w)$ and $\varphi(z_{w'})$ lie in different connected components of $T\backslash B(\varphi(v),1)$. As $n \geqslant r$ and $r \geqslant 1$, we also have that $\varphi(z_w)$ and $\varphi(z_{w'})$ must lie in different connected components of $T\backslash B(\varphi(v),r)$. Note that each such connected component must have diameter at least $d_T(\varphi(v), \varphi(z_w)) - 1 \geqslant n - 3 > R$.

By our initial assumption there are at most countably many such connected components, so there are at most countably many $\varphi(z_w)$. As $w \neq w'$ implies that $\varphi(z_w) \neq \varphi(z_{w'})$, there are at most countably many $w \in S(v,1)\backslash\{\widehat{w}\}$, and so at most countably many $w \in S(v,1)$. Hence Γ' is a countable simplicial tree, which is $(1,2n+4)$-quasi-isometric to T. □

Remark 2.3.26. This is not true if we let $r = 0$. As an example, consider the comb space X formed by taking copies X_n of $[0, \infty)$ indexed by $n \in \mathbb{Z}_{\geqslant 0}$, then attaching them to $[0, 1]$ by identifying the 0 in X_0 with 0 in $[0, 1]$, and the 0 in X_n with $\frac{1}{n}$ in $[0, 1]$ for $n \in \mathbb{N}$. This is an $\mathbb{R}$-tree, and for any $x \in X$ we can see that $X \backslash \{x\}$ has at most three connected components, however by Proposition 2.3.27 below we can see that that X is not (L, C)-quasi-isometric to any locally finite simplicial tree.

Proposition 2.3.27. *Given an $\mathbb{R}$-tree T, we have the following:*

(1) If T is (L, C)-quasi-isometric to a countable simplicial tree, then for every $r \geqslant 0$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has countably many connected components of diameter at least R.

(2) If T is (L, C)-quasi-isometric to a locally finite simplicial tree, then for every $r \geqslant 0$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has finitely many connected components of diameter at least R.

Proof. We again prove only the first part, as the proof of the second part is analogous. Let T be an $\mathbb{R}$-tree, and suppose there exists a countable simplicial tree Γ with a $(1, C)$-quasi-isometry $\varphi : \Gamma \to T$. Suppose also that there exists some $r \geqslant 0$ such that for every $R \geqslant 0$ there exists $x \in T$ such that $T \backslash B(x, r)$ has uncountably many connected components of diameter at least R.

Given $r \geqslant 0$, we let $R = 3L + 5C$, and pick such an $x \in T$. By coarse surjectivity we know that there exists $v' \in \Gamma$ such that $d_T(\varphi(v'), x) \leqslant C$, and a vertex $v \in \Gamma$ such that $d_\Gamma(v, v') \leqslant 1$. Hence $d_T(\varphi(v), \varphi(v')) \leqslant L + C$, so $d_T(\varphi(v), x) \leqslant L + 2C$.

Let $\mathcal{C}_R$ be the set of connected components of $T \backslash B(x, r)$ of diameter greater than R, and for each $C \in \mathcal{C}_R$ pick $x_C \in C$ such that $r + \frac{1}{2}R \leqslant d_T(x, x_C) \leqslant r + R$, which must always exist as T is an $\mathbb{R}$-tree. Let $v_C \in \Gamma$ be a vertex such that $d_T(\varphi(v_C), x_C) \leqslant L + 2C$, then $d_T(\varphi(v), \varphi(v_C)) \leqslant r + R + 2(L + 2C)$. Therefore $d_\Gamma(v, v_C) \leqslant L(R + r + 2L + 5C)$.

Note that for $C \neq C'$ we have that

$$\begin{aligned} d_T(x_C, x_{C'}) \geqslant R &\implies d_T(\varphi(v_C), \varphi(v_{C'})) \geqslant R - 2(L + 2C) \\ &\implies d_\Gamma(v_C, v_{C'}) \geqslant \frac{1}{L}(R - 2L - 5C) > 0. \end{aligned}$$

This means that $v_C \neq v_{C'}$, so there are uncountably many such v_C's. Recall that when $m = L(R + r + 2L + 5C)$ we have that $d_\Gamma(v, v_C) \leqslant m$ for every $C \in \mathcal{C}_R$, so this means that there are uncountably many vertices in $B(v, m)$. This contradicts Γ being countable. Hence if T is $(1, C)$-quasi-isometric to a countable simplicial tree

we have that for every $r \geqslant 0$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has countably many connected components of diameter at least R. $\square$

We therefore obtain some equivalent criteria for when an $\mathbb{R}$-tree is quasi-isometric to a countable or locally finite simplicial tree, including a partial analogue to Corollary 2.3.8.

Corollary 2.3.28. *Given an $\mathbb{R}$-tree T, the following are equivalent:*

(1) T is (L, C)-quasi-isometric to a countable simplicial tree for some $L \geqslant 1, C \geqslant 0$.

(2) T is $(1, C')$-quasi-isometric to a countable simplicial tree for some $C' \geqslant 0$.

(3) For some $r \geqslant 1$ there exists $R > 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has countably many connected components of diameter at least R.

(4) For every $r \geqslant 0$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has countably many connected components of diameter at least R.

Corollary 2.3.29. *Given an $\mathbb{R}$-tree T, the following are equivalent:*

(1) T is (L, C)-quasi-isometric to a locally finite simplicial tree for some $L \geqslant 1, C \geqslant 0$.

(2) T is $(1, C')$-quasi-isometric to a locally finite simplicial tree for some $C' \geqslant 0$.

(3) For some $r \geqslant 1$ there exists $R > 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has finitely many connected components of diameter at least R.

(4) For every $r \geqslant 0$ there exists $R \geqslant 0$ such that for every $x \in T$ the space $T \backslash B(x, r)$ has finitely many connected components of diameter at least R.

We can then extend the first equivalences to geodesic metric spaces to get a full analogue of Corollary 2.3.8.

Corollary 2.3.30. *We have the following:*

(1) A geodesic metric space is (L, C)-quasi-isometric to a countable simplicial tree for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to a countable simplicial tree for some $C' \geqslant 0$.

(2) A geodesic metric space is (L, C)-quasi-isometric to a locally finite simplicial tree for some $L \geqslant 1, C \geqslant 0$ if and only if it is $(1, C')$-quasi-isometric to a locally finite simplicial tree for some $C' \geqslant 0$.

Proof. Let X be a geodesic metric space that is quasi-isometric to a countable simplicial tree. We then have that X is a quasi-tree, so there exists $C' \geqslant 0$ such that X is $(1, C')$-quasi-isometric to T_X. This means that T_X is an $\mathbb{R}$-tree that is quasi-isometric to a countable simplicial tree, so by Corollary 2.3.28 there exists $C'' \geqslant 0$ such that T_X is $(1, C'')$-quasi-isometric to a countable simplicial tree. We conclude that X is $(1, C' + 2C'')$-quasi-isometric to a countable simplicial tree by Lemma 2.1.23. The reverse direction is obvious, and the second part follows using the same method. □

This allows us to get analogues to Proposition 2.3.18 in the case of quasi-actions on countable or locally finite simplicial trees, using the same reasoning as in the proof of Proposition 2.3.18.

Corollary 2.3.31. *We have the following:*

(1) If a finitely generated group admits a cobounded (L, C)-quasi-action on a countable simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1, C')$-quasi-action on a countable simplicial tree for some $C' \geqslant 0$.

(2) If a finitely generated group admits a cobounded (L, C)-quasi-action on a locally finite simplicial tree for some $L \geqslant 1, C \geqslant 0$, then it admits a quasi-conjugate $(1, C')$-quasi-action on a locally finite simplicial tree for some $C' \geqslant 0$.

We conclude this section by noting that if our quasi-tree is a proper space then we can use Proposition 2.3.25 to obtain a $(1, C)$-quasi-isometry with a locally finite simplicial tree.

Corollary 2.3.32. *Every proper quasi-tree is $(1, C)$-quasi-isometric to a locally finite simplicial tree.*

Proof. We first show that every proper $\mathbb{R}$-tree has this property. Let T be an $\mathbb{R}$-tree, and suppose it is not $(1, C)$-quasi-isometric to a locally finite simplicial tree. By Proposition 2.3.25 we can find $x \in T$ such that $T \backslash B(x, 1)$ has infinitely many connected components of diameter greater than 2. Let $\mathcal{C}$ be the set of these connected components.

For every $C \in \mathcal{C}$ we can pick $x_C \in C$ such that $d(x, x_C) = 2$. Note that, as T is an $\mathbb{R}$-tree, for every $C, C' \in \mathcal{C}$ such that $C \neq C'$ we have that $d(x_C, x_{C'}) \geqslant 2$. We therefore have an infinite collection of points in $B(x, 2)$, no subsequence of which can converge. Hence $B(x, 2)$ is not compact, which implies that T is not proper.

Let X be a proper quasi-tree, then T_X is proper as the quotient map is continuous. As X is $(1, C)$-quasi-isometric to T_X, and T_X is $(1, C')$-quasi-isometric to a locally finite simplicial tree, the conclusion follows. □

2.4 Uniform and non-uniform tree approximation

We are now able to get a uniform version of Gromov's Tree Approximation Lemma in the case of quasi-trees, as an easy consequence of the results in Sections 2.2 and 2.3. In fact, we have an even stronger result, as our uniform approximation applies to any subset of our quasi-tree, as opposed to the finite subsets that were considered for general hyperbolic spaces. This result will be used in Chapter 4 to obtain a theorem about growth in acylindrically hyperbolic groups.

Proposition 2.4.1. *Let (X, d) be a δ-hyperbolic quasi-tree, with bottleneck constant $\Delta \geqslant 0$. Let $x_0 \in X$, and let $Z \subset X$. Let Y be a union of geodesic segments $\bigcup_{z\in Z}[x_0, z]$. Then there is an $\mathbb{R}$-tree T and a map $f : (Y, d) \to (T, d^*)$ such that:*

(1) For all $z \in Z$, the restriction of f to the geodesic segment $[x_0, z]$ is an isometry.

(2) For all $x, y \in Y$, we have that $d(x, y) - 2(\Delta + 2\delta) \leqslant d^*(f(x), f(y)) \leqslant d(x, y)$.

Proof. Consider the end-approximating tree T_X, and the end-approximating map $f : X \to T_X$. We now consider the restriction $f|_Y : Y \to T_X$. The fact that f is an isometry on every geodesic segment $[x_0, z]$ was noted in the first part of the proof of Proposition 2.2.8, so it is also true for $f|_Y$.

For the second part we simply apply Proposition 2.3.2. It was noted in the proof that, given X is a quasi-tree, we have that $d(x, y) - 2(\Delta + 2\delta) \leqslant d^*([x], [y]) \leqslant d(x, y)$ for every $x, y \in X$. Hence this is also true for $f|_Y$ when applied to any $x, y \in Y$. □

The converse to this is obviously true, as by taking our set Z to be the whole of X we would get a quasi-isometric embedding from X to an $\mathbb{R}$-tree, giving us that X is a quasi-tree by the following standard result.

Definition 2.4.2. A *subtree* T of a metric space X is a subspace of X that is an $\mathbb{R}$-tree under the shortest path metric d_T.

Example 2.4.3. Any path in a metric space that does not intersect itself is a subtree. For a graph, any subgraph without cycles is also a subtree.

Lemma 2.4.4. *Let X be a geodesic metric space. If X quasi-isometrically embeds into an $\mathbb{R}$-tree T, then X is quasi-isometric to a subtree of T.*

Proof. Let $f : X \to T$ be an (L, C)-quasi-isometric embedding, where T is an $\mathbb{R}$-tree. Assume that $C > 0$. Let $x_0 \in X$, and let T' be the union of geodesic segments

$\bigcup_{z\in X}[f(x_0), f(z)]$. The space T' is a path connected subset of an $\mathbb{R}$-tree, so is itself an $\mathbb{R}$-tree. We want to show that $f : X \to T'$ is a quasi-isometry.

As $f : X \to T'$ is an (L, C)-quasi-isometric embedding, we only need to check that for every $y \in T'$ there exists $x \in X$ such that $d'(f(x), y) \leqslant C$. Suppose this is not the case, so for some $y \in T'$ we have that $B(y, C) \cap f(X) = \emptyset$. By the definition of T', we can find $z \in X$ such that $y \in [f(x_0), f(z)]$.

As $B(y, C) \cap f(X) = \emptyset$, we have that $f([x_0, z])$ is disconnected in T', where $f(x_0)$ and $f(z)$ are in different connected components. We can group the connected components of $f([x_0, z])$ together depending on which component of $T'\backslash\{y\}$ they lie in.

For every $\varepsilon > 0$, we can therefore find $z_1, z_2 \in [x_0, z] \subset X$ such that $d(z_1, z_2) \leqslant \varepsilon$, and $f(z_1), f(z_2)$ lie in different connected components of $f([x_0, z])$, with $y \in [f(z_1), f(z_2)]$. This tells us that $2C \leqslant d'(f(z_1), f(z_2)) \leqslant Ld(z_1, z_2) + C \leqslant L\varepsilon + C$, which is a contradiction for $\varepsilon < \frac{C}{L}$. Hence $f : X \to T'$ is a quasi-isometry. □

Although it is sufficient by the above, the converse of Proposition 2.4.1 is actually a far stronger assumption than we need to prove that X is a quasi-tree. As it turns out, it is possible to show that a global quasi-isometry exists when we only assume uniform approximation for finite subsets. In particular, this means that having uniform tree approximation for finite subsets is equivalent to having uniform tree approximation for all subsets, so quasi-trees are exactly the geodesic spaces for which uniform tree approximation is possible.

In fact, to show that our space is a quasi-tree, we can even loosen our idea of uniform tree approximation for finite subsets to allow for a multiplicative error as well. This means that having this less strict version of tree approximation for all finite subsets is equivalent to having the stronger version for all finite subsets.

Proposition 2.4.5. *Let (X, d) be a geodesic metric space. The following are equivalent:*

(1) X is a quasi-tree.

(2) There exists a constant $C \geqslant 0$ such that for every finite subset Z of X, there exists an $\mathbb{R}$-tree (T, d^*) and a $(1, C)$-quasi-isometric embedding $f : (Z, d) \to (T, d^*)$.

(3) There exist constants $L \geqslant 1$ and $C \geqslant 0$ such that for every finite subset Z of X, there exists an $\mathbb{R}$-tree (T, d^*) and a (L, C)-quasi-isometric embedding $f : (Z, d) \to (T, d^*)$.

Proof. The first statement implies the second statement by Proposition 2.4.1. The second statement clearly implies the third. It remains to show that the third statement implies that X is a quasi-tree.

Now suppose that such constants $L \geqslant 1$ and $C \geqslant 0$ exist. Let $A = \max\{1, C\}$, so $A > 0$, and note that any (L, C)-quasi-isometric embedding is also a (L, A)-quasi-isometric embedding. We want to show that X satisfies the bottleneck criterion. Let $x, y \in X$, let $[x, y]$ be a geodesic between them, and let $m \in [x, y]$ be the midpoint. Let $\{a_0, \ldots, a_k\} \subset [x, y]$ be such that $a_0 = x$, $a_k = y$, and $d(a_i, a_{i+1}) \leqslant \frac{A}{L}$. Let γ be a path between x and y, and let $\{z_0, \ldots, z_n\} \subset \gamma$ be such that $z_0 = x$, $z_n = y$, and $d(z_i, z_{i+1}) \leqslant \frac{A}{L}$.

We let $Y = \{m, a_0, \ldots, a_k, z_1, \ldots, z_{n-1}\}$, and so by our assumption there exists exists an $\mathbb{R}$-tree T and a map $f : (Y, d) \to (T, d^*)$ such that $\frac{1}{L}d(a, b) - A \leqslant d^*(f(a), f(b)) \leqslant Ld(a, b) + A$ for every $a, b \in Y$.

Consider $[f(x), f(y)]$. As T is an $\mathbb{R}$-tree, there exists $m' \in [f(x), f(y)]$ such that $d^*(f(x), f(m)) = d^*(f(x), m') + d^*(m', f(m))$ and $d^*(f(m), f(y)) = d^*(f(m), m') + d^*(m', f(y))$. In other words, m' is the closest point in $[f(x), f(y)]$ to $f(m)$.

Consider the path formed by the concatenation $[f(a_0), f(a_1)] * \cdots * [f(a_{k-1}), f(a_k)]$. This is a path from $f(x)$ to $f(y)$ in T, so $m' \in [f(a_i), f(a_{i+1})]$ for some i. Recall that $d(a_i, a_{i+1}) \leqslant \frac{A}{L}$, so $d^*(f(a_i), f(a_{i+1})) \leqslant 2A$. Let $j \in \{1, \ldots, n\}$ be the least element such that $f(a_j) \in B(m', 2A)$, and let $l \in \{1, \ldots, n\}$ be the largest element such that $f(a_l) \in B(m', 2A)$. The above implies that these exist, and that $j \neq l$.

We note that $f(a_j)$ lies in the same connected component of $T \backslash \{m'\}$ as $f(x)$, as either $a_j = x$ or $f(a_{j-1})$ lies in the same connected component of $T \backslash B(m', 2A)$ as $f(x)$. By the same reasoning, $f(a_l)$ lies in the same connected component of $T \backslash \{m'\}$ as $f(y)$, so m' separates $f(a_j)$ and $f(a_l)$ in T, and therefore $m' \in [f(a_j), f(a_l)]$.

We also want to show that $m \in [a_j, a_l]$. Note that $m \in [a_i, a_{i+1}]$ for some i, and both $d^*(f(a_i), f(m)) \leqslant 2A$ and $d^*(f(a_{i+1}), f(m)) \leqslant 2A$. Either $f(m) = m'$, or $f(m)$ lies in a different connected component of $T \backslash \{m'\}$ to $f(x)$ and to $f(y)$. Either way, this implies that $i \geqslant j$ and $i + 1 \leqslant l$, so $m \in [a_j, a_l]$.

We therefore have that

$$\begin{aligned} d^*(f(a_j), f(a_l)) + 2d^*(m', f(m)) &= d^*(f(a_j), f(m)) + d^*(f(m), f(a_l)) \\ &\leqslant Ld(a_j, m) + Ld(m, a_l) + 2A \\ &= Ld(a_j, a_l) + 2A \\ &\leqslant L^2 d^*(f(a_j), f(a_l)) + L^2 A + 2A. \end{aligned}$$

This means that

$$\begin{aligned} d^*(m', f(m)) &\leqslant \frac{1}{2}((L^2-1)d^*(f(a_j), f(a_l)) + L^2A + 2A) \\ &\leqslant \frac{1}{2}(4(L^2-1)A + L^2A + 2A) \\ &= 3L^2A - A. \end{aligned}$$

Consider the path formed by the concatenation $[f(z_0), f(z_1)] * \cdots * [f(z_{n-1}), f(z_n)]$. This is a path from $f(x)$ to $f(y)$ in T, so as before $m' \in [f(z_i), f(z_{i+1})]$ for some i. Note that $d^*(f(z_i), f(z_{i+1})) \leqslant Ld(z_i, z_{i+1}) + A \leqslant 2A$, so for some i we have that $d^*(f(z_i), m') \leqslant A$.

We conclude that

$$\begin{aligned} d(m, z_i) &\leqslant Ld^*(f(m), f(z_i)) + A \\ &\leqslant L(d^*(f(m), m') + d^*(m', f(z_i))) + A \\ &\leqslant 3L^3A + A, \end{aligned}$$

so X satisfies the bottleneck criterion, and therefore is a quasi-tree. □

Gromov also proved an alternative version of tree approximation in hyperbolic spaces, in which the approximating tree is a subtree of the hyperbolic space.

Proposition 2.4.6. [43, p. 157–158][14] *Suppose X is a δ-hyperbolic geodesic metric space. There exists a function $h : \mathbb{N} \to [0, \infty)$ such that if $Y \subset X$ is a finite set of points with $|Y| = n$, then there exists a subtree $T \subset X$ such that $Y \subset T$ and for every $x, y \in T$ we have that $d_T(x, y) \leqslant d(x, y) + \delta h(n)$.*

It is clear from the definition of d_T that $d(x, y) \leqslant d_T(x, y)$. In this version of tree approximation the $\mathbb{R}$-tree T is constructed in the obvious way: let $Y = \{y_1, \ldots, y_n\}$, set $T_1 = [y_1, y_2]$, then set $T_2 = [y_3, t] \cup T_1$ where t is a closest point in T_1 to y_3, and so on. This process will terminate at $T = T_{n-1}$.

It is natural to ask if there is a uniform version of this type of tree approximation in the case of quasi-trees.

Question 1. *Let X be a quasi-tree. Does there exist a constant $C \geqslant 0$, depending only on X, such that for any set of finite points $Y \subset X$ we can construct a subtree $T \subset X$ such that $Y \subset T$, and for every $x, y \in T$ we have that $d_T(x, y) \leqslant d(x, y) + C$?*

It is immediately obvious that taking the same approach as in the proof of Proposition 2.4.6 will not work, as the following example shows.

Example 2.4.7. Consider the space $X = [0,1] \times \mathbb{R}$ with the Euclidean metric. It is clear that X is a quasi-tree. Let $n \in \mathbb{N}$ be odd, and let $Y = \{(0,0), (1,1), (0,2), \ldots, (0,n-1), (1,n)\}$. If a tree T is constructed using this ordering of the points, we get that $d((0,0),(1,n)) = \sqrt{1+n^2}$ and $d_T((0,0),(1,n)) = n\sqrt{2}$, so

$$\lim_{n\to\infty} \frac{d((0,0),(1,n))}{d_T((0,0),(1,n))} = \frac{1}{\sqrt{2}}.$$

In particular, we cannot have that $d_T((0,0),(1,n)) \leqslant d((0,0),(1,n)) + C$ for all odd $n \in \mathbb{N}$, whatever the choice of the constant C.

Although the naive approach does not work, in the above example the answer to Question 1 is actually yes, which we can easily see by letting $T = (\{\frac{1}{2}\} \times \mathbb{R}) \cup \bigcup_{i=1}^{n}([0,1] \times \{y_i\})$, and letting $C = 1$. We can however show that there are quasi-trees in which the answer to the above question will always be no, whichever way the subtree is constructed. The idea in the following example is that, over a large scale, it is not possible to keep the paths in a subtree tight to geodesics in all directions.

Example 2.4.8. Take the infinite 4-regular tree with a base vertex x_0 and edges of length 1. For all $n \in \mathbb{Z}_{\geqslant 0}$ replace the vertices in $S(x_0, n)$ with copies of the rectangle $[0,2] \times [0, 2^{n+1}]$. Attach the edges to these vertex rectangles at the midpoint of the sides, with one of the short sides oriented towards $[0,2]^2$ (see Figure 2.6 for an illustration). We will call this space X, and note that it is a quasi-tree under the piecewise Euclidean metric by Manning's bottleneck criterion.

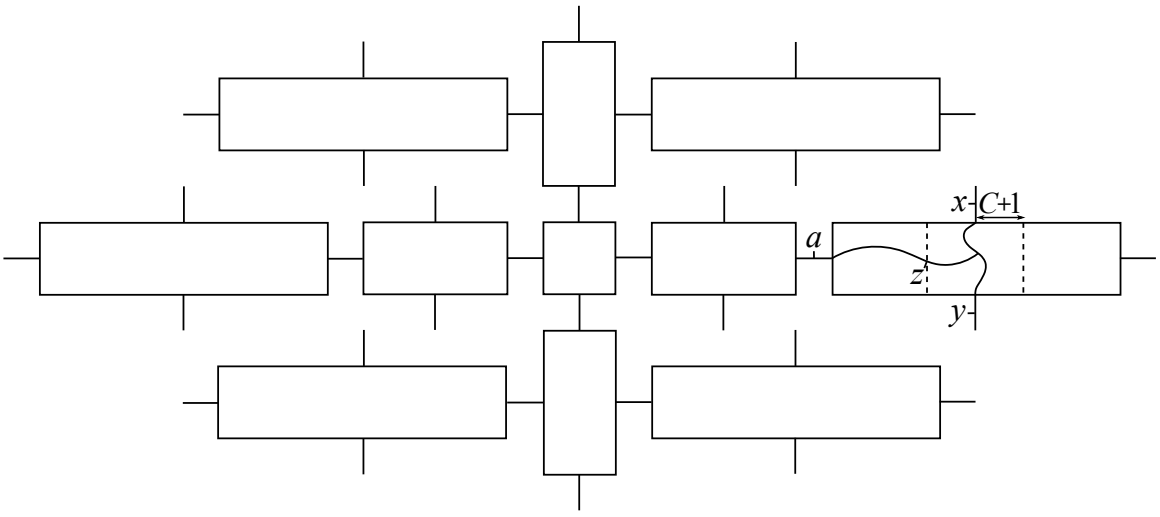


Figure 2.6: A counterexample to uniform subtree approximation

Let $C > 0$ be arbitrary. Let $n \in \mathbb{N}$ be such that $2^{n+1} > 2C + 3$, and let Y_n be the set of midpoints of the edges in X that correspond to the edges in $B(x_0, n)$ in the

4-regular tree. Suppose that T is a subspace of X such that $Y_n \subset T$, T is a tree under the shortest path metric d_T, and $d_T(x,y) \leqslant d(x,y) + C$ for every $x, y \in T$. Consider a rectangle in X of size $[0,2] \times [0, 2^{m+1}]$, where $2C + 2 < 2^{m+1} < 2^{n+1}$.

Let $x, y \in Y_n$ be the midpoints of the edges attached to the long sides of this rectangle, and let $a \in Y_n$ be the midpoint of the edge attached to the short side of this rectangle nearest to the central square. We will show that at least one of $d_T(a,x)$ or $d_T(a,y)$ must have large error in comparison to $d(a,x)$ or $d(a,y)$, respectively.

We first note that there must be a unique shortest path in T that connects x to y, and it must have length at most $d(x,y) + C = 3 + C$ by our assumption on T. Therefore this path must stay in the subrectangle $[0,2] \times [2^m - (C+1), 2^m + C + 1]$, as otherwise $d_T(x,y) \geqslant 1 + 2(C+1) = 3 + 2C > d(x,y) + C$.

As $2^m > C + 1$ this subrectangle is a proper subset of the original rectangle, and in particular does not contain the attaching point of the edge containing a. As T is a tree, this means that outside this subrectangle the shortest path in T between a and x must be identical to the shortest path in T between a and y. Let z be the point at which this path intersects $[0,2] \times \{2^m - (C+1)\}$, and suppose without loss of generality that $d(x,z) \geqslant d(y,z)$.

We now want to use this to compare $d_T(a,x)$ to $d(a,x)$. We first note that

$$d(a,z) + d(z,x) \leqslant d_T(a,z) + d_T(z,x) = d_T(a,x).$$

For $z' \in [0,2] \times \{2^m - (C+1)\}$ satisfying $d(x,z') \geqslant d(y,z')$, the sum $d(a,z') + d(z',x)$ will be minimal when $z' = (1, 2^m - (C+1))$, which is the central point of that side of the subrectangle. Substituting this in we get that

$$d_T(a,x) \geqslant d(a,z) + d(z,x) \geqslant 2^m - (C+1) + \frac{1}{2} + \sqrt{1 + (C+1)^2} + \frac{1}{2},$$

while $d(a,x) = 1 + \sqrt{1 + (2^m)^2}$, so

$$d_T(a,x) - d(a,x) \geqslant \sqrt{1 + (C+1)^2} - (C+1) + 2^m - \sqrt{1 + (2^m)^2}.$$

Note that

$$2^m - \sqrt{1 + (2^m)^2} \to 0 \text{ as } m \to \infty,$$

so taking $\varepsilon = \frac{\sqrt{1+(C+1)^2} - (C+1)}{2}$, we know that there exists some $M \in \mathbb{N}$ such that $2^M > C + 1$, and such that for all choices of m where $n > m \geqslant M$ we have that $d_T(a,x) \geqslant d(a,x) + \varepsilon$. This is true for any approximating tree $T \subset X$ of Y_n, and in fact M can be chosen uniformly, as it is not dependent on n.

We now want to use this to show that, if we choose large enough n, no subtree of X can approximate Y_n with error no greater than C. Suppose otherwise, so suppose there exists a constant $C > 0$ such that for any set of finite points $Y \subset X$ there exists an $\mathbb{R}$-tree T that is embedded in X such that $Y \subset T$, and such that for every $x, y \in T$ we have that $d_T(x, y) \leqslant d(x, y) + C$. Let $n \in \mathbb{N}$ be such that $\varepsilon(n - M - 1) > C$, and consider Y_n, with T the associated subtree of X.

Let x_1 be the midpoint of an edge adjacent to the central square $[0, 2]^2$. For x_i, where $1 \leqslant i \leqslant n - 2$, we choose $x_{i+1} \in Y_n$ such that x_i and x_{i+1} are midpoints of edges attached to the same rectangle, x_{i+1} is strictly further from the central square than x_i, and x_{i+1} is the midpoint of an edge attached to one of the long sides of the shared rectangle. We can think of this as moving outwards in the quasi-tree X, turning "left" or "right" each time we enter a rectangle. This gives us two possible options for x_{i+1}, we take the option that maximises $d_T(x_i, x_{i+1})$.

This means that for $i \geqslant M$ we have that $d_T(x_i, x_{i+1}) \geqslant d(x_i, x_{i+1}) + \varepsilon$, and we note that for all other i we still have that $d_T(x_i, x_{i+1}) \geqslant d(x_i, x_{i+1})$. By our construction of X, any path between x_1 and x_{n-1} must pass through every x_i, so we can see that $d_T(x_1, x_{n-1}) = d_T(x_1, x_2) + \cdots + d_T(x_{n-2}, x_{n-1})$. It is also clear from our construction of X that $d(x_1, x_{n-1}) = d(x_1, x_2) + \cdots + d(x_{n-2}, x_{n-1})$, so

$$d_T(x_1, x_{n-1}) \geqslant d(x_1, x_{n-1}) + \varepsilon(n - M - 1) > d(x_1, x_{n-1}) + C.$$

This is a contradiction, so no such constant $C > 0$ exists for the quasi-tree X.

Although the answer to Question 1 is no in general, we can still ask the following question.

Question 2. *Does there exists a characterisation of those quasi-trees for which the answer to Question 1 is yes?*

An example of such a quasi-tree would be one that mimics the construction in Example 2.4.8, but with the length of the rectangles being bounded. Here a uniform subtree can be constructed by picking a basepoint and taking the union of all infinite geodesic rays from this point. This example works firstly because the subtree that is constructed is a bounded distance from any point in the quasi-tree, however this could also be true for a subtree of the quasi-tree constructed in Example 2.4.8. The key difference is that here the only error when travelling within our subtree comes from when a path changes direction between going towards the basepoint and going away, and in this example that error is bounded. For any finite set of points in

our quasi-tree we can then simply add a geodesic between each point and its closest neighbour in our subtree, and this will satisfy the uniform error condition.

Although this method works for some examples, it is not a sufficient condition. Consider the quasi-tree as in Example 2.4.8, but this time instead of replacing the vertices with rectangles that simply get longer, replace them with rectangles of size $[0, \frac{1}{2^n}] \times [0, n]$, so they get longer and thinner as we move away from the base square. It is not hard to see that in this case a geodesic construction would not work: the "turning error" is not bounded, as some geodesics will diverge from each other as soon as they enter the vertex rectangles. This means that this error will increase as the length of the vertex rectangles increases. On the other hand, by instead taking a central line through these vertex rectangles it is possible to construct a subtree with the desired properties, as the vertex rectangles in the quasi-tree progressively get thinner, so the sum of the errors you get can be bounded.

This may suggest one approach to answering Question 2, which is to consider the error that would arise from attaching geodesics γ_1 and γ_2 at the point where they diverge in the hyperbolic sense, for example at approximately the distance given by $\limsup_{n\to\infty}(\gamma_1(n), \gamma_2(n))_{x_0}$. This may give control over the turning error, in which case if these attachings can be chosen in such a way that the cumulative error they create for any geodesic is bounded, then such a quasi-tree should give a positive answer to Question 1.

Whether such an approach could give a complete answer to Question 2 depends partly on the following question, which asks whether the existence of approximating subtrees for finite sets implies the existence of a universal approximating subtree.

Question 3. *Is a positive answer to Question 1 for a quasi-tree X equivalent to the existence of a subtree $T \subset X$ and a constant $C \geqslant 0$ such that for every $x, y \in T$ we have that $d_T(x, y) \leqslant d(x, y) + C$, and for every $z \in X$ we have that $d(z, T) \leqslant C$?*

2.5 Geometry of the end-approximating tree

We will now return to the end-approximating tree T_X. In Section 2.2 we showed that if X is a geodesic metric space then T_X is an $\mathbb{R}$-tree, and in Section 2.3 we showed that if X is a quasi-tree then T_X is $(1, C)$-quasi-isometric to X. In this section we will give more detail on the how the geometry of X relates to the geometry of T_X. In particular, we will show that when X is proper their spaces of ends are homeomorphic, and that when X is a quasi-tree their boundaries can be thought of as being isometric.

2.5.1 Path components

Here we will formalise the idea that T_X is constructed by collapsing X along the spheres $S(x_0, r)$, which was mentioned at the end of Section 2.2. In doing so we will look at the correspondence between the path components of $X \backslash B(x_0, r)$ and $T_X \backslash B([x_0], r)$, which will be relevant in Subsections 2.5.2 and 2.5.3.

For this section we will assume a fixed basepoint $x_0 \in X$. We begin with the following easy lemma.

Lemma 2.5.1. *Let (X, d) be a metric space, and let $x, y \in X$. Suppose that for every $\varepsilon > 0$ there exists a path $\gamma : [0,1] \to X$ from x to y such that $d(x_0, \gamma(t)) \geqslant \min\{d(x_0, x), d(x_0, y)\} - \varepsilon$ for all $t \in [0,1]$. Then $d^*([x], [y]) = |d(x_0, x) - d(x_0, y)|$.*

Proof. Recall that $d^*([x], [y]) = d'(x, y) = d(x_0, x) + d(x_0, y) - 2(x, y)'_{x_0}$, so if we can show that $(x, y)'_{x_0} = \min\{d(x_0, x), d(x_0, y)\}$ then we will be done. We already know that $(x, y)'_{x_0} \leqslant \min\{d(x_0, x), d(x_0, y)\}$, so we just need the opposite inequality.

Let $\varepsilon > 0$, and let $\gamma : [0,1] \to X$ be a path from x to y such that $d(x_0, \gamma(t)) \geqslant \min\{d(x_0, x), d(x_0, y)\} - \varepsilon$ for all $t \in [0,1]$. Let $\nu > 0$, and choose a sequence $0 = t_1 < t_2 < \cdots < t_n = 1$ such that $d(\gamma(t_i), \gamma(t_{i+1})) \leqslant \nu$ for every $1 \leqslant i \leqslant n-1$. Then

$$\begin{aligned}(\gamma(t_i), \gamma(t_{i+1}))_{x_0} &= \frac{1}{2}(d(x_0, \gamma(t_i)) + d(x_0, \gamma(t_{i+1})) - d(\gamma(t_i), \gamma(t_{i+1}))) \\ &\geqslant \frac{1}{2}(2\min\{d(x_0, x), d(x_0, y)\} - 2\varepsilon - \nu) \\ &= \min\{d(x_0, x), d(x_0, y)\} - \varepsilon - \frac{\nu}{2}.\end{aligned}$$

It follows that $(x, y)'_{x_0} \geqslant \min\{d(x_0, x), d(x_0, y)\} - \varepsilon - \frac{\nu}{2}$, so as $\varepsilon > 0$ and $\nu > 0$ were arbitrary we can conclude that $(x, y)'_{x_0} \geqslant \min\{d(x_0, x), d(x_0, y)\}$, and we therefore have equality. □

Corollary 2.5.2. *If X is a geodesic metric space such that for every $r \in [0, \infty)$ and every $x, y \in S(x_0, r)$ there exists a path γ between x and y such that $d(x_0, \gamma) \geqslant r$, then T_X is isometric to $[0, \infty)$.*

We would like to show that the reverse implication to Lemma 2.5.1 also holds. Note that the $\varepsilon > 0$ is necessary for this to be the case. For example, if we consider X to be an open ball around some x_0 in $\mathbb{R}^2$ with two points x and y added from its boundary, then it can be seen that those boundary points would still collapse together in T_X despite the lack of any path such that $d(x_0, \gamma(t)) \geqslant \min\{d(x_0, x), d(x_0, y)\}$.

To show the reverse implication to Lemma 2.5.1 we will need to look at the relationship between the path components of $X \backslash B(x_0, r)$ and $T_X \backslash B([x_0], r)$.

Lemma 2.5.3. *Let (X, d) be a geodesic metric space. Let $x, y \in X$, and let $r \in [0, \infty)$. Then x and y lie in the same path component of $X \backslash B(x_0, r)$ if and only if $([x], [y])^*_{[x_0]} > r$.*

Proof. Suppose x and y lie in the same path component of $X \backslash B(x_0, r)$. Then there exists $\delta > 0$ and a path $\gamma : [0, 1] \to X$ between them such that $d(x_0, \gamma(t)) \geqslant r + \delta$ for all $t \in [0, 1]$. The existence of such a δ follows from the fact that γ is continuous, as this implies that $\gamma([0, 1])$ is a compact subset of $X \backslash B(x_0, r)$. If there existed a sequence of points (x_n) in $\gamma([0, 1])$ such that $d(x_0, x_n) \to r$, then there would exist $x \in \gamma([0, 1])$ such that $d(x_0, x) = r$, a contradiction. Therefore we can apply the same argument as in the proof of Lemma 2.5.1 to get that $([x], [y])^*_{[x_0]} = (x, y)'_{x_0} \geqslant r + \delta > r$.

Now suppose x or y do not lie in the same path component of $X \backslash B(x_0, r)$. If x or y lie in $B(x_0, r)$ then it is immediate that

$$([x], [y])^*_{[x_0]} = (x, y)'_{x_0} \leqslant \min\{d(x_0, x), d(x_0, y)\} \leqslant r.$$

If x and y lie in different path components of $X \backslash B(x_0, r)$, then any geodesic $[x, y]$ must intersect $B(x_0, r)$. Let $z \in [x, y] \cap B(x_0, r)$. As $d(x_0, z) \leqslant r$, by Lemma 2.1.11 we have that $(x, y)_{x_0} \leqslant r$.

Let $x = x_1, \ldots, x_n = y$ be any sequence in $S_{x,y}$. As x and y lie in different path components of $X \backslash B(x_0, r)$, there must be some pair x_i, x_{i+1} such that $x_i \in B(x_0, r)$, or $x_{i+1} \in B(x_0, r)$, or x_i and x_{i+1} lie in different path components of $X \backslash B(x_0, r)$. Therefore $(x_i, x_{i+1})_{x_0} \leqslant r$, so as this sequence in $S_{x,y}$ was arbitrary we must have that $([x], [y])^*_{[x_0]} = (x, y)'_{x_0} \leqslant r$. □

Remark 2.5.4. This would not hold in general if we instead removed an open ball and considered when $([x], [y])^*_{[x_0]} \geqslant r$. To see this we can again look at the example of X being an open ball in $\mathbb{R}^2$ with two points added from its boundary.

Corollary 2.5.5. *Let (X, d) be a geodesic metric space. Let $x, y \in X$, and let $r \in [0, \infty)$. We then have that x and y lie in the same path component of $X \backslash B(x_0, r)$ if and only if $[x]$ and $[y]$ lie in the same path component of $T_X \backslash B([x_0], r)$.*

Proof. As T_X is an $\mathbb{R}$-tree, $[x]$ and $[y]$ lie in the same path component of $T_X \backslash B([x_0], r)$ if and only if $([x], [y])^*_{[x_0]} > r$, which by Lemma 2.5.3 is true if and only if x and y lie in the same path component of $X \backslash B(x_0, r)$. □

We can now give an alternative definition of the metric d^* from the geometry of X.

Corollary 2.5.6. *Let X be a geodesic metric space, and let $x, y \in X$. Then we have that $([x],[y])^*_{[x_0]} = \sup\{r \in [0,\infty) : x$ and y lie in the same path component of $X \backslash B(x_0, r)\}$.*

Proof. This follows from the fact that T_X is an $\mathbb{R}$-tree, so $([x],[y])^*_{[x_0]} = d^*([x_0],[z])$, where $[z] \in T_X$ is the unique point such that $[[x_0],[z]] = [[x_0],[x]] \cap [[x_0],[y]]$. We then simply have to notice that $d^*([x_0],[z]) = \sup\{r \in [0,\infty) : [x]$ and $[y]$ lie in the same path component of $T_X \backslash B([x_0], r)\}$, and apply Corollary 2.5.5. □

Remark 2.5.7. If X is an $\mathbb{R}$-tree then this confirms that $T_X = X$, as Corollary 2.5.6 tells us that $(x,y)_{x_0} = ([x],[y])^*_{[x_0]}$, so $d(x,y) = d^*([x],[y])$.

Remark 2.5.8. Recalling that $([x],[y])^*_{[x_0]} = (x,y)'_{x_0}$, we can use this along with Lemma 2.1.10 and Lemma 2.1.11 to give an alternative proof of the fact that when X is a quasi-tree the inequality $(x,y)_{x_0} \leqslant (x,y)'_{x_0} \leqslant (x,y)_{x_0} + \Delta + 2\delta$ holds, where $\Delta \geqslant 0$ is the bottleneck constant and $\delta \geqslant 0$ is the hyperbolicity constant.

We can now obtain our improvement of Lemma 2.5.1.

Corollary 2.5.9. *Let (X,d) be a geodesic metric space, and let $x, y \in X$. We have that $d^*([x],[y]) = |d(x_0,x) - d(x_0,y)|$ if and only if for every $\varepsilon > 0$ there exists a path $\gamma : [0,1] \to X$ from x to y such that $d(x_0, \gamma(t)) \geqslant \min\{d(x_0,x), d(x_0,y)\} - \varepsilon$ for all $t \in [0,1]$.*

Proof. If $d^*([x],[y]) = |d(x_0,x) - d(x_0,y)|$ then $([x],[y])^*_{[x_0]} = \min\{d(x_0,x), d(x_0,y)\}$, so the forward implication follows from Corollary 2.5.6. The reverse implication is Lemma 2.5.1. □

As a particular case of this, we can see that two points in X will collapse together in T_X if and only if they lie on the same sphere and there exist appropriate paths between them, as mentioned at the end of Section 2.2.

Corollary 2.5.10. *Let (X,d) be a geodesic metric space, and let $x, y \in X$. We have that $[x] = [y]$ if and only if there exists $r \in [0,\infty)$ such that $x, y \in S(x_0, r)$ and for every $\varepsilon > 0$ there exists a path γ between them in X such that $d(x_0, \gamma) \geqslant r - \varepsilon$.*

Proof. Suppose $[x] = [y]$. Then $d^*([x_0],[x]) = d^*([x_0],[y])$, so $d(x_0,x) = d(x_0,y)$, and so $d^*([x],[y]) = |d(x_0,x) - d(x_0,y)|$. Both directions now follow from Corollary 2.5.9. □

Remark 2.5.11. Suppose that two geodesic rays $\gamma_1 : [0, \infty) \to X$ and $\gamma_2 : [0, \infty) \to X$ based at x_0 have the property that for every $t \in [0, \infty)$ and $\varepsilon > 0$ the points $\gamma_1(t)$ and $\gamma_2(t)$ are in the same path component of $X \backslash B(x_0, t - \varepsilon)$. Then Corollary 2.5.10 means that $[\gamma_1(t)] = [\gamma_2(t)]$ for every $t \in [0, \infty)$, and in particular γ_1 and γ_2 are sent to the same geodesic ray in T_X.

2.5.2 Ends of T_X for proper metric spaces

In this subsection we will show that, when X is a proper metric space, we can think of T_X as X with each of its ends collapsed into one geodesic ray. This justifies our description of T_X as the end-approximating tree of a metric space. In particular, we will show that the space of ends of X is homeomorphic to the space of ends of T_X, and as a result is homeomorphic to the space of ends of a locally finite simplicial tree.

We will define the space of ends for a proper geodesic metric space, for the more general definition in topological spaces see [18, p. 144]. As before, we will assume a fixed basepoint $x_0 \in X$.

Definition 2.5.12. Let X be a proper metric space. A *proper ray* in X is a continuous map $\gamma : [0, \infty) \to X$ such that for every $r \in [0, \infty)$ the set $\gamma^{-1}(B(x_0, r))$ is bounded.

Remark 2.5.13. If we had picked an alternative basepoint $x_0' \in X$, then by the triangle inequality we would have that for every $r \in [0, \infty)$ there exists $r' \in [0, \infty)$ such that $B(x_0, r) \subset B(x_0', r')$. This tells us that the choice of basepoint does not matter in the above definition.

Definition 2.5.14. Let X be a proper metric space, and let γ_1 and γ_2 be proper rays in X. We say that γ_1 and γ_2 *converge to the same end* if for every $r \in [0, \infty)$ there exists $t \in [0, \infty)$ such that $\gamma_1([t, \infty))$ and $\gamma_2([t, \infty))$ are contained in the same path component of $X \backslash B(x_0, r)$. This defines an equivalence relation on the proper rays in X. We denote the equivalence class of a proper ray γ by $\mathrm{end}(\gamma)$.

Example 2.5.15. A bounded space has no ends, $\mathbb{R}$ has two ends, and $\mathbb{R}^2$ has one.

The following lemma can be found in [18]. We will mimic the idea of its proof for Lemma 2.5.21.

Lemma 2.5.16. [18, Lemma I.8.28] *Let X be a proper geodesic metric space. For every proper ray γ in X there exists a geodesic ray γ' based at $x_0 \in X$ such that $\gamma' \in \mathrm{end}(\gamma)$.*

The proof of the above lemma relies on the Arzelà-Ascoli Theorem, we state the necessary version of it here. For a stronger version, see [71, p. 290].

Theorem 2.5.17 (Arzelà-Ascoli Theorem)**.** *Let X and Y be metric spaces, and let $\mathcal{F}$ be a subset of $\{f : X \to Y : f \text{ continuous}\}$. If $\mathcal{F}$ is equicontinuous, and the set $\mathcal{F}_x = \{f(x) : f \in \mathcal{F}\}$ has compact closure in Y for every $x \in X$, then a subsequence of $\mathcal{F}$ converges pointwise to a continuous function $f : X \to Y$.*

The space of ends, and the topology on it, are defined as follows.

Definition 2.5.18. Let X be a proper metric space. The *space of ends* of X is the set of equivalence classes of proper ray in X, denoted $\mathrm{Ends}(X)$. A sequence of ends converges, $\mathrm{end}(\gamma_n) \to \mathrm{end}(\gamma)$, if for every $r \in [0, \infty)$ there exists an increasing sequence of real numbers $(N_n)_{n\in\mathbb{N}}$ such that $\gamma_n([N_n, \infty))$ and $\gamma([N_n, \infty))$ lie in the same path component of $X \backslash B(x_0, r)$ for all $n \in \mathbb{N}$ sufficiently large. The topology on $\mathrm{Ends}(X)$ is induced by this convergence.

If two proper metric spaces X and Y are quasi-isometric, then it is well known that this quasi-isometry induces a homeomorphism between $\mathrm{Ends}(X)$ and $\mathrm{Ends}(Y)$ [18, Proposition I.8.29]. This means that when X is a proper quasi-tree, we automatically get that $\mathrm{Ends}(X)$ and $\mathrm{Ends}(T_X)$ will be homeomorphic as a consequence of Proposition 2.3.2. We will show that this is in fact true for every proper metric space X.

Lemma 2.5.19. *Let (X, d) be a proper geodesic metric space. Let $f : X \to T_X$ be the end-approximating map. For every proper ray $\gamma : [0, \infty) \to X$, the map $f \circ \gamma : [0, \infty) \to T_X$ is a proper ray.*

Proof. We first recall that, as $d^*([x], [y]) \leqslant d(x, y)$, the function f is continuous. Therefore $f \circ \gamma$ is continuous. Now let $r \in [0, \infty)$, and note that $f^{-1}(B([x_0], r)) = B(x_0, r)$. As γ is a proper ray, $\gamma^{-1}(f^{-1}(B([x_0], r))) = \gamma^{-1}(B(x_0, r))$ is bounded, so we can conclude that $f \circ \gamma$ is a proper ray. □

Corollary 2.5.20. *Let (X, d) be a proper geodesic metric space. Let $f : X \to T_X$ be the end-approximating map. Let $\gamma_1 : [0, \infty) \to X$ and $\gamma_2 : [0, \infty) \to X$ be proper rays in X. Then $\gamma_2 \in \mathrm{end}(\gamma_1)$ if and only if $f \circ \gamma_2 \in \mathrm{end}(f \circ \gamma_1)$.*

Proof. We have that $\gamma_2 \in \mathrm{end}(\gamma_1)$ if and only if for every $r \in [0, \infty)$ there exists $t \in [0, \infty)$ such that $\gamma_1([t, \infty))$ and $\gamma_2([t, \infty))$ lie in the same path component of $X \backslash B(x_0, r)$. By Corollary 2.5.5 this is true if and only if for every $r \in [0, \infty)$ there exists $t \in [0, \infty)$ such that $f(\gamma_1([t, \infty)))$ and $f(\gamma_2([t, \infty)))$ lie in the same path component of $T_X \backslash B([x_0], r)$, and this is equivalent to saying that $f \circ \gamma_2 \in \mathrm{end}(f \circ \gamma_1)$. □

The proof of the following lemma uses the same ideas that are used in the proof of Lemma 2.5.16, as found in [18].

Lemma 2.5.21. *Let (X, d) be a proper geodesic metric space. Let $f : X \to T_X$ be the end-approximating map. Let $\gamma : [0, \infty) \to T_X$ be a proper ray in T_X. Then there exists a geodesic ray $\xi : [0, \infty) \to X$ such that $f \circ \xi \in \mathrm{end}(\gamma)$.*

Proof. We know by Lemma 2.5.16 that there exists a geodesic representative $\gamma' \in \mathrm{end}(\gamma)$ based at our chosen basepoint x_0. For every $n \in \mathbb{N}$ we choose $x_n \in f^{-1}(\gamma'(n))$, then define $\xi_n : [0, \infty) \to X$ by extending the geodesic $[x_0, x_n]$ such that ξ_n is constant on $[d(x_0, x_n), \infty)$. Note that this collection of functions is equicontinuous.

We can also note that for any $t \in [0, \infty)$ we have that $\{\xi_n(t) : n \in \mathbb{N}\} \subset B(x_0, t)$, so as X is proper the closure is compact. By the Arzelà-Ascoli Theorem there exists a subsequence $(\xi_{n_k})_{k\in\mathbb{N}}$ of $(\xi_n)_{n\in\mathbb{N}}$ that converges pointwise to a continuous function $\xi : [0, \infty) \to X$. We want to show that ξ is geodesic and that $f \circ \xi = \gamma'$. The fact that ξ is geodesic follows immediately from the fact that for every $t \in [0, \infty)$ we have that $d(x_0, \xi(t)) = \lim_{k\to\infty} d(x_0, \xi_{n_k}(t))$, and that $d(x_0, \xi_{n_k}(t)) = t$ for large enough k.

As f is continuous we also get that $f \circ \xi(t) = \lim_{k\to\infty} f \circ \xi_{n_k}(t)$. For $k \in \mathbb{N}$ such that $n_k \geqslant t$ we have that $\xi_{n_k}(t)$ lies on the geodesic $[x_0, x_{n_k}]$, where $x_{n_k} \in f^{-1}(\gamma'(n_k))$. Hence $f \circ \xi_{n_k}(t)$ lies on the geodesic $[[x_0], \gamma'(n)]$ in T_X. As $d^*([x_0], f \circ \xi_{n_k}(t)) = t$, and T_X is an $\mathbb{R}$-tree, we must have that $f \circ \xi_{n_k}(t) = \gamma'(t)$. Therefore $f \circ \xi(t) = \gamma'(t)$ for all $t \in [0, \infty)$, so $f \circ \xi = \gamma'$. We conclude that ξ is a geodesic ray such that $f \circ \xi \in \mathrm{end}(\gamma)$. □

Proposition 2.5.22. *Let (X, d) be a proper geodesic metric space. Let $f : X \to T_X$ be the end-approximating map, and let $F : \mathrm{Ends}(X) \to \mathrm{Ends}(T_X)$ be defined by $F(\mathrm{end}(\gamma)) = \mathrm{end}(f \circ \gamma)$ for every $\mathrm{end}(\gamma) \in \mathrm{Ends}(X)$. The map F is a homeomorphism.*

Proof. The fact that F is well-defined and injective comes from Corollary 2.5.20, and the fact that it is surjective comes from Lemma 2.5.21. It remains to show that F and F^{-1} are continuous.

Let $V \subset \mathrm{Ends}(T_X)$ be closed. We want to show that $F^{-1}(V)$ is closed. Let $(\mathrm{end}(\xi_n))_{n\in\mathbb{N}}$ be a sequence in $F^{-1}(V)$ such that $\mathrm{end}(\xi_n) \to \mathrm{end}(\xi)$, and choose these representatives to be geodesic rays. We want to show that $\mathrm{end}(\xi) \in F^{-1}(V)$.

We first note that $(\mathrm{end}(f \circ \xi_n))_{n\in\mathbb{N}}$ gives a sequence in V. We want to show that $\mathrm{end}(f \circ \xi_n) \to \mathrm{end}(f \circ \xi)$. Let $r \in [0, \infty)$. As $\mathrm{end}(\xi_n) \to \mathrm{end}(\xi)$, there exists an increasing sequence of natural numbers $(N_n)_{n\in\mathbb{N}}$ such that $\xi_n([N_n, \infty))$ and $\xi([N_n, \infty))$ lie in

the same path component of $X \backslash B(x_0, r)$ for large enough n. Clearly $f \circ \xi_n([N_n, \infty))$ and $f \circ \xi([N_n, \infty))$ lie in $T_X \backslash B([x_0], r)$ for large enough n, and by Corollary 2.5.5 they must lie in the same path component. Hence $\text{end}(f \circ \xi_n) \to \text{end}(f \circ \xi)$, so as V is closed we must have that $\text{end}(f \circ \xi) \in V$, so $\text{end}(\xi) \in F^{-1}(V)$, so $F^{-1}(V)$ is closed. Therefore F is continuous.

The argument for F^{-1} is almost identical. Let $U \subset \text{Ends}(X)$ be closed. We want to show that $F(U)$ is closed. Let $(\text{end}(\gamma_n))_{n \in \mathbb{N}}$ be a sequence in $F(U)$ such that $\text{end}(\gamma_n) \to \text{end}(\gamma)$, and choose these representatives to be geodesic rays. We want to show that $\text{end}(\gamma) \in F(U)$.

We pick geodesic rays $\xi_n, \xi : [0, \infty) \to X$ such that $\text{end}(f \circ \xi_n) = \text{end}(\gamma_n)$ and $\text{end}(f \circ \xi) = \text{end}(\gamma)$, and note that $(\text{end}(\xi_n))_{n \in \mathbb{N}}$ gives a sequence in U, as F is bijective. We want to show that $\text{end}(\xi_n) \to \text{end}(\xi)$. Let $r \in [0, \infty)$. As $\text{end}(f \circ \xi_n) \to \text{end}(f \circ \xi)$, there exists an increasing sequence of natural numbers $(N_n)_{n \in \mathbb{N}}$ such that $f \circ \xi_n([N_n, \infty))$ and $f \circ \xi([N_n, \infty))$ lie in the same path component of $T_X \backslash B([x_0], r)$ for large enough n.

Clearly $\xi_n([N_n, \infty))$ and $\xi([N_n, \infty))$ lie in $X \backslash B(x_0, r)$ for large enough n, and by Corollary 2.5.5 they must lie in the same path component. Hence $\text{end}(\xi_n) \to \text{end}(\xi)$, so as U is closed we must have that $\text{end}(\xi) \in U$, so $\text{end}(f \circ \xi) \in F(U)$, so $F(U)$ is closed. Therefore F^{-1} is continuous. □

We can therefore obtain the following result.

Proposition 2.5.23. *Let (X, d) be a proper geodesic metric space. Then there exists a locally finite simplicial tree T such that* $\text{Ends}(X)$ *is homeomorphic to* $\text{Ends}(T)$.

Proof. As the space of ends is preserved by quasi-isometry, it suffices to show that T_X is quasi-isometric to a locally finite simplicial tree. By Corollary 2.3.32, it therefore suffices to show that T_X is proper.

Let $r > 0$, and recall that for the end-approximating map $f : X \to T_X$ we have that $B([x_0], r) = f(B(x_0, r))$. As X is proper, $B(x_0, r)$ is compact, so as f is continuous we have that $B([x_0], r)$ is also compact.

Let $[y] \in T_X$, and let $R > 0$. There exists $r > 0$ such that $B([y], R) \subset B([x_0], r)$. Therefore $B([y], R)$ is a closed subset of a compact set, so is itself compact. We conclude that T_X is a proper metric space. □

2.5.3 Visual metrics on the boundaries of quasi-trees

For a tree, its space of ends is homeomorphic to its boundary. This is consequently also true for quasi-trees, as both the space of ends and the geodesic boundary are preserved up to homeomorphism by quasi-isometries. In the case that X is a proper quasi-tree, the result of Proposition 2.5.22 is a trivial consequence of this fact.

Not only is the topology on the boundaries preserved by quasi-isometries, some of the metric structure is too. Given a quasi-isometry between two hyperbolic geodesic metric spaces, the induced homeomorphism between their boundaries is known to be quasi-symmetric [13, 18, 61]. There is a family of metrics on these boundaries such that this is the case, known as the visual metrics.

When we have a $(1, C)$-quasi-isometry between spaces, it is known that the induced homeomorphism between the boundaries will in fact be a bi-Lipschitz map with respect to the visual metrics [13]. In particular, by how visual metrics are defined, this means that the pullback of a visual metric on one of the boundaries will be a visual metric on the other, so the homeomorphism can be instead viewed as an isometry.

In the case of $\mathbb{R}$-trees, given a choice of parameter, the visual metrics on the boundary with that parameter are all bi-Lipschitz equivalent to a visual metric which can be written down explicitly in a standard way. By Proposition 2.3.2 and the above we can choose a metric on the boundary of a quasi-tree X such that it is isometric to the boundary of T_X under this standard metric. Moreover, this metric on the boundary of X is a natural extension of the standard visual metric on the boundary of T_X.

Many of the definitions and facts about boundaries used here are taken from a survey paper on the topic by Kapovich and Benakli [52].

Convention. When working with metrics on boundaries of X, everything will be done with respect to a basepoint $x_0 \in X$. As in previous sections, this will be assumed throughout to have already been chosen. The reason the choice of basepoint does not matter will become clear later.

Definition 2.5.24. Let (X, d) be a hyperbolic metric space. A sequence (x_n) in X *converges to infinity* if $\lim_{i,j\to\infty}(x_i, x_j)_{x_0} = \infty$. We say that two such sequences are *equivalent*, and write $(x_n) \sim (y_n)$, if $\liminf_{i,j\to\infty}(x_i, y_j)_{x_0} = \infty$. The *sequential boundary* of X is the set of equivalence classes of sequences that converge to infinity, and is denoted by ∂X.

Notation 2.5.25. When $p \in \partial X$, and a sequence (x_n) converging to infinity in X belongs to the equivalence class p, we write $x_n \to p$.

As mentioned before, it is possible to put metrics on this boundary, although there is not a single canonical metric. To define the family of metrics we are interested in, we first note that the Gromov product can be extended to the sequential boundary.

Definition 2.5.26. Let X be a hyperbolic metric space. For $p, q \in \partial X$, the *Gromov product* of p and q is

$$(p,q)_{x_0} = \sup\{\liminf_{i,j\to\infty}(x_i, y_j)_{x_0} : x_i \to \xi, y_j \to \gamma\}.$$

Definition 2.5.27. Let X be a hyperbolic metric space, and let $a > 1$. A metric $d_{\partial X}$ on ∂X is a *visual metric* with respect to the *visual parameter* a if there exists $C \geqslant 1$ such that for every $p, q \in \partial X$

$$\frac{1}{C} a^{-(p,q)_{x_0}} \leqslant d_{\partial X}(p,q) \leqslant C a^{-(p,q)_{x_0}}.$$

It is known that for any hyperbolic space X there exists $a_0 > 1$ such that for every $a \in (1, a_0)$ the sequential boundary ∂X admits a visual metric with parameter a, see [18, Chapter III.H] or [44] for example. If X is an $\mathbb{R}$-tree then $a_0 = \infty$, and for every parameter $a > 1$ there is a standard visual metric given by $d_{\partial X}(p,q) = a^{-(p,q)_{x_0}}$.

Example 2.5.28. For a metric space of finite diameter, the sequential boundary is empty. For an n-regular simplicial tree the sequential boundary is a Cantor space, and for hyperbolic n-space with $n \geqslant 2$ it is an $(n-1)$-sphere.

There is another notion of boundary that applies when X is also a proper geodesic space.

Definition 2.5.29. Let (X, d) be a metric space, and let A and B be subsets of X. The *Hausdorff distance* between A and B is

$$d_H(A,B) = \inf\{\varepsilon \geqslant 0 : A \subset N(B,\varepsilon) \text{ and } B \subset N(A,\varepsilon)\},$$

where $N(A, \varepsilon)$ is the closed ε-neighbourhood of A.

Definition 2.5.30. Let (X, d) be a proper, hyperbolic, and geodesic metric space. We say that two geodesic rays $p : [0, \infty) \to X$ and $q : [0, \infty) \to X$ are equivalent, and write $p \sim q$, if $d_H(p, q)$ is finite. The *geodesic boundary* of X is the set of equivalence classes of geodesic rays in X, and is denoted by $\partial_\infty X$.

It is again possible to put a metric on this boundary, although as before there is not a single canonical metric.

Remark 2.5.31. For any such space X, and any $p, q \in \partial_\infty X$ that are not equal, there will exist a bi-infinite geodesic $\gamma : \mathbb{R} \to X$ from p to q, in the sense that $\gamma|_{[0,\infty)} \in q$ and $\widehat{\gamma}|_{[0,\infty)} \in p$ where $\widehat{\gamma}(t) = \gamma(-t)$. When X is an $\mathbb{R}$-tree this bi-infinite geodesic is unique.

Definition 2.5.32. Let (X, d) be a proper, hyperbolic, and geodesic metric space, and let $a > 1$. A metric $d_{\partial_\infty X}$ on $\partial_\infty X$ is a *visual metric* with respect to the *visual parameter* a if there exists $C \geqslant 1$ such that for every $p, q \in \partial_\infty X$, and every bi-infinite geodesic γ from p to q, we have that

$$\frac{1}{C} a^{-d(x_0,\gamma)} \leqslant d_{\partial_\infty X}(p, q) \leqslant C a^{-d(x_0,\gamma)}.$$

As with the sequential boundary, for any such space X, there exists $a_0 > 1$ such that for every $a \in (1, a_0)$ the geodesic boundary $\partial_\infty X$ admits a visual metric with parameter a. If X is an $\mathbb{R}$-tree then $a_0 = \infty$, and for every parameter $a > 1$ there is a standard visual metric given by $d_{\partial_\infty X}(p, q) = a^{-d(x_0,\gamma)}$ [52].

Remark 2.5.33. The visual metrics on the sequential and geodesic boundaries are defined in terms of a basepoint, however on either boundary two visual metrics with the same parameter but different basepoints will be bi-Lipschitz equivalent [52, 21].

It is well known that a quasi-isometry f between two hyperbolic spaces X and Y induces a canonical homeomorphism $\partial f : \partial X \to \partial Y$, and if X and Y are also geodesic then the same is true for $\partial_\infty f : \partial_\infty X \to \partial_\infty Y$. This homeomorphism has strong metric properties depending on the strength of the quasi-isometry, see [21, 52], however we will be interested in the following case.

Theorem 2.5.34. [13, Theorem 6.5] *Let X and Y be hyperbolic geodesic metric spaces, and let $f : X \to Y$ be a $(1, C)$-quasi-isometry. Let $d_{\partial X}$ and $d_{\partial Y}$ be visual metrics with parameter $a > 1$ on ∂X and ∂Y respectively. Then the induced map $\partial f : (\partial X, d_{\partial X}) \to (\partial Y, d_{\partial Y})$ is bi-Lipschitz.*

This may not be immediately obvious from reading Theorem 6.5 in [13], however we note that what they call a $(1, C)$-rough similarity is exactly a $(1, C)$-quasi-isometry, and moreover that if $f : X \to Y$ is a $(1, C)$-quasi-isometry then there exists $K \geqslant 0$ such that for every $x, y, x_0 \in X$ we have that $(x, y)_{x_0} - K \leqslant (f(x), f(y))_{f(x_0)} \leqslant (x, y)_{x_0} + K$. Their reasoning therefore gives us that ∂f is what they call a $(1, \lambda)$-snowflake map, which is precisely a bi-Lipschitz map.

Remark 2.5.35. As noted before, the family of visual metrics on a boundary with the same parameter is a bi-Lipschitz equivalence class, so if we define a metric d_X on ∂X by $d_{\partial X}(p,q) = d_{\partial Y}(\partial f(p), \partial f(q))$ then this would be a visual metric on ∂X, which would allow us to view ∂f as an isometry.

In a hyperbolic geodesic space X, it is well known that there is a natural homeomorphism between ∂X and $\partial_\infty X$, and moreover that visual metrics on ∂X and $\partial_\infty X$ with the same parameter are bi-Lipschitz equivalent, for example see [32, 45]. We therefore also get the following from [13].

Corollary 2.5.36. [13] *Let X and Y be hyperbolic geodesic metric spaces, and let $f : X \to Y$ be a $(1, C)$-quasi-isometry. Let $d_{\partial_\infty X}$ and $d_{\partial_\infty Y}$ be visual metrics with parameter $a > 1$ on $\partial_\infty X$ and $\partial_\infty Y$ respectively. Then the induced map $\partial_\infty f : (\partial_\infty X, d_{\partial_\infty X}) \to (\partial_\infty Y, d_{\partial_\infty Y})$ is bi-Lipschitz.*

Corollary 2.5.37. *Let X be a quasi-tree. For every $a > 1$ there exists a visual metric $d_{\partial X}$ on ∂X with visual parameter a such that $(\partial X, d_{\partial X})$ is isometric to $(\partial T_X, d_{\partial T_X})$, where $d_{\partial T_X}$ is the standard visual metric on ∂T_X with visual parameter a.*

Proof. Let $f : X \to T_X$ be the end-approximating map. Then Proposition 2.3.2 tells us that f is a $(1, C)$-quasi-isometry. Let $\partial f : \partial X \to \partial T_X$ be the induced bi-Lipschitz homeomorphism between the boundaries, and for $a > 1$ let $d_{\partial T_X}$ be the standard visual metric with parameter a on ∂T_X. Then $d_{\partial X}(p,q) = d_{\partial T_X}(\partial f(p), \partial f(q))$ is a visual metric on ∂X with visual parameter a, and ∂f acts as an isometry between the boundaries. □

Remark 2.5.38. The exact same statement holds if we replace T_X by Γ_X, the simplicial tree constructed from T_X in Subsection 2.3.2. We also get the same result for both T_X and Γ_X if we suppose that X is a proper quasi-tree, and then consider the geodesic boundaries.

In the case that the metric on ∂X is the pullback of the standard visual metric on ∂T_X (with respect to some parameter and basepoint), then we can in fact define this metric explicitly in a very natural way, using only the geometry of X. We start with the sequential boundary. Let $p, q \in \partial X$. Then for any $a > 1$ we define

$$d_{\partial X}(p,q) = d_{\partial T_X}(\partial f(p), \partial f(q)) = a^{-(\partial f(p), \partial f(p))^*_{[x_0]}}.$$

Let $x_i \to p$ and $y_j \to q$. The homeomorphism $\partial f : X \to T_X$ is defined such that $[x_n] \to \partial f(p)$ and $[y_n] \to \partial f(q)$. From the definition of the Gromov product on the

boundary, the geometry of an $\mathbb{R}$-tree, and Corollary 2.5.5 regarding path components we can see that

$$\begin{aligned}(\partial f(p), \partial f(q))^*_{[x_0]} &= \sup\{r \in [0,\infty) : \text{ for } [x_n] \to \partial f(p) \text{ and } [y_n] \to \partial f(q) \text{ there are} \\ &\quad \text{infinitely many terms in each sequence that lie in the same path} \\ &\quad \text{component of } T_X \backslash B([x_0], r)\} \\ &= \sup\{r \in [0,\infty) : \text{ for } x_n \to p \text{ and } y_n \to q \text{ there are infinitely} \\ &\quad \text{many terms in each sequence that lie in the same path} \\ &\quad \text{component of } X \backslash B(x_0, r)\}.\end{aligned}$$

Note that the metric on ∂X is given in terms of this value. In particular, it is clear that if X is an $\mathbb{R}$-tree then this simply returns the standard visual metric with parameter a, as expected given that f is the identity in this case.

We can do the same thing with the geodesic boundary. Let $p, q \in \partial_\infty X$. As f is an isometry when restricted to geodesics based at x_0, and T_X is an $\mathbb{R}$-tree, the image of the equivalence class p under $\partial_\infty f$ will be a single geodesic in T_X, with the same being true for the image of q. For any $a > 1$ we define

$$d_{\partial_\infty X}(p, q) = d_{\partial_\infty T_X}(\partial_\infty f(p), \partial_\infty f(q)) = a^{-d^*([x_0],[z])},$$

where $[[x_0], [z]] = p' \cap q'$, with the geodesics p' and q' being the images of the equivalence classes p and q. In particular we can see that

$$\begin{aligned}d^*([x_0], [z]) &= \inf\{r \in [0,\infty) : p' \backslash B([x_0], r) \text{ and } q' \backslash B([x_0], r) \text{ lie in different path} \\ &\quad \text{components of } T_X \backslash B([x_0], r)\} \\ &= \inf\{r \in [0,\infty) : r_p \backslash B(x_0, r) \text{ and } r_q \backslash B(x_0, r) \text{ lie in different path} \\ &\quad \text{components of } X \backslash B(x_0, r), \text{ for any choice of } r_p \in p \text{ and } r_q \in q\},\end{aligned}$$

and we note that $d_{\partial_\infty X}(p, q)$ is given in terms of this value. Again, if X is an $\mathbb{R}$-tree then this will give us the standard visual metric with parameter a, so these metrics on the boundaries of quasi-trees can be viewed as extensions of the standard visual metrics on the boundaries of $\mathbb{R}$-trees.

3

Product set growth

"Ford!" he said, "there's an infinite number of monkeys outside who want to talk to us about this script for Hamlet they've worked out."

– Douglas Adams, The Hitchhiker's Guide to the Galaxy

For a finite subset U of a group G, its *nth product set* is

$$U^n = \{u_1 \cdots u_n \ : \ u_1, \ldots, u_n \in U\}.$$

The topic of product set growth in infinite groups is about understanding how $|U^n|$ behaves as $n \to \infty$. The classical version of group growth in this setting is when G is taken to be finitely generated, and U is taken to be a ball in G with respect to the word metric induced by some finite generating set S. This will be the focus of Section 3.1, where we will give some background on the concepts of exponential and uniform exponential growth in groups.

For the remaining chapters of this thesis, however, we will be more interested in the case that U is an arbitrary finite subset of G. Some of the basics of this more general type of growth, as well as its similarities and differences with the case that U is a ball, will be explored in Section 3.2. We will then introduce the concept of uniform product set growth in Section 3.3, and say something about its relationship with uniform exponential growth. It will be there that we introduce the central question of the second half of this thesis, which Chapters 4 and 5 will be dedicated to answering.

In Section 3.4 we will give some of the history behind uniform product set growth, and then in Section 3.5 we will talk about a couple of consequences of a group or subgroup having this property. The remainder of this chapter will then be spent proving some general results about when certain subgroups can or cannot have uniform product set growth, which will be useful in later chapters.

3.1 Exponential and uniform exponential growth

In this section we will give an overview of the classical notion of group growth in infinite groups, with our focus being on groups with exponential growth. Note that all of the lemmas in this section are standard results that can be found in a number of textbooks on this subject. For a broader look at growth, which is not focused purely on the exponential side, we refer the reader to [64].

3.1.1 Growth of groups

Let G be a finitely generated group, and let S be a finite generating set of G. We will use $B_S(n)$ to denote the closed ball of radius n centred at the identity $e \in G$ under the word metric induced by S. In the notation of product sets, it is easy to note that when $U = B_S(1)$ we will have that $U^n = B_S(n)$. The growth of a group is characterised by the large scale growth of these balls as n increases. We will mainly be interested in groups where this growth is exponential.

Definition 3.1.1. Let G be a finitely generated group, and let S be a finite generating set of G. Let $B_S(n)$ denote the ball of radius n under the word metric induced by S. The *exponential growth rate of G with respect to S* is

$$\omega(G, S) = \lim_{n\to\infty} |B_S(n)|^{\frac{1}{n}}.$$

If $\omega(G, S) > 1$ for some (equivalently any) finite generating set S, then G has *exponential growth*.

Remark 3.1.2. This limit always exists, as $\log(|B_S(n)|)$ as a function of n is subadditive. That is, for all $n, m \in \mathbb{N}$ we have that $B_S(n+m) = B_S(n)B_S(m)$, so $|B_S(n+m)| \leqslant |B_S(n)||B_S(m)|$. Fekete's Lemma [36] tells us that $\lim_{n\to\infty} \frac{1}{n} \log(|B_S(n)|)$ exists, so $\omega(G, S)$ also exists.

Note that in the above definition, we said that we have $\omega(G, S) > 1$ for one finite generating set of G if and only if this is true for every finite generating set. This fact is easy to check.

Lemma 3.1.3. *Let G be a finitely generated group, and let S be a finite generating set of G. Suppose that $\omega(G, S) > 1$. Then for any other finite generating set S' of G we also have that $\omega(G, S') > 1$.*

Proof. Let $S = \{s_1, \ldots, s_k\}$, and let S' be any finite generating set of G. Now let $N \in \mathbb{N}$ be the maximum length of any $s_i \in S$ under the word metric induced by S'. Then $B_S(1) \subset B_{S'}(N)$, so for any $n \in \mathbb{N}$ we have that $B_S(n) \subset B_{S'}(nN)$. We therefore get that

$$\omega(G, S') = \lim_{n\to\infty} |B_{S'}(n)|^{\frac{1}{n}} = \lim_{n\to\infty} |B_{S'}(nN)|^{\frac{1}{nN}} \geqslant \lim_{n\to\infty} |B_S(n)|^{\frac{1}{nN}} = \omega(G, S)^{\frac{1}{N}} > 1.$$

□

Remark 3.1.4. This result is more generally a consequence of the fact that exponential growth is a quasi-isometry invariant of groups (see [30, p. 172], for example), and that the Cayley graphs we get from the same group with two different finite generating sets are quasi-isometric to each other.

The intuitive reason as to why $\omega(G, S)$ being strictly greater than one is referred to as exponential growth is given by the lemma below.

Lemma 3.1.5. *A finitely generated group G has exponential growth if and only if for every finite generating set S of G there exists $a > 1$ and $N \in \mathbb{N}$ such that $|B_S(n)| \geqslant a^n$ for every $n \geqslant N$.*

Proof. Let S be some finite generating set of G, and suppose such an a and N exist. Then

$$\begin{aligned}
\omega(G, S) &= \lim_{n\to\infty} |B_S(n)|^{\frac{1}{n}} \\
&\geqslant \lim_{n\to\infty} (a^n)^{\frac{1}{n}} \\
&= \lim_{n\to\infty} a \\
&= a \\
&> 1.
\end{aligned}$$

Now suppose that $\omega(G, S) > 1$ for some finite generating set S of G. Then we can choose $a > 1$ such that $\omega(G, S) > a$. Hence

$$\lim_{n\to\infty} |B_S(n)|^{\frac{1}{n}} > a,$$

so for some $N \in \mathbb{N}$ we have that $n \geqslant N$ implies

$$|B_S(n)|^{\frac{1}{n}} > a.$$

Hence

$$|B_S(n)| > a^n.$$

□

Remark 3.1.6. We note in particular that $\omega(G,S) > a$ implies that $|B_S(n)| > a^n$ for large enough n. In other words, not only does $\omega(G,S)$ tell us whether G has exponential growth or not, but it also gives us an explicit lower bound on how fast $|B_S(n)|$ grows for large n.

Exponential growth tends to be a relatively easy property to prove, due to the fact that it only needs to be checked for a single generating set. For example, it is not hard to see that all free groups on two or more generators have exponential growth, and abelian groups do not.

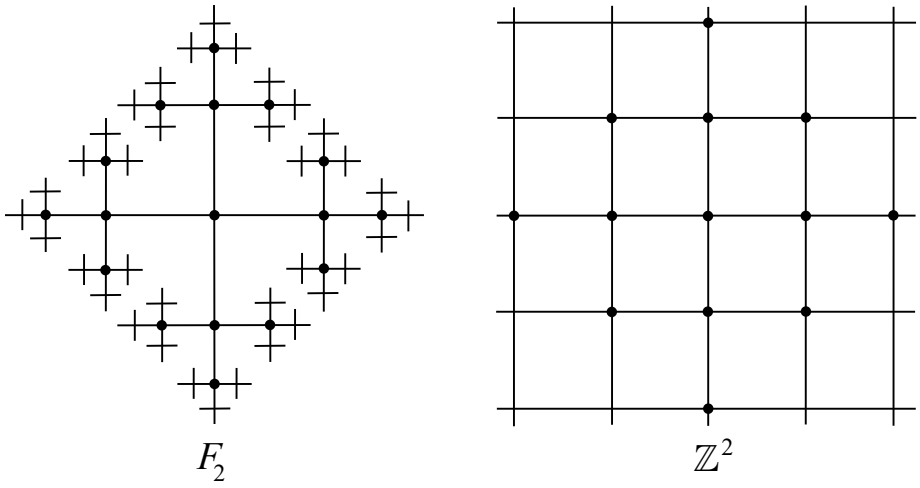


Figure 3.1: Balls of radius two in the free group and $\mathbb{Z}^2$

On the other hand, although having exponential growth tells us that all finite generating sets will exhibit some kind of exponential behaviour, there is not necessarily any control over their rates of growth, which may be arbitrarily close to one. If we wish to have this type of control, then we will need to prove something stronger.

Definition 3.1.7. Let G be a finitely generated group, and let

$$\mathcal{S} = \{S \subset G \ : \ S \text{ is a finite generating set of } G\}.$$

The *exponential growth rate of* G is

$$\omega(G) = \inf_{S \in \mathcal{S}} \omega(G,S).$$

If $\omega(G) > 1$ then G has *uniform exponential growth.*

In light of Remark 3.1.6, this is saying that a group G with exponential growth has uniform exponential growth if and only if there exists some $a > 1$ such that for every finite generating set S of G we have that $|B_S(n)| > a^n$ for large enough n. In other words, we have a common lower bound on the growth of balls in the group.

Clearly uniform exponential growth implies exponential growth. Most groups that are known to have exponential growth also have uniform exponential growth, including free groups, hyperbolic groups, and relatively hyperbolic groups [87]. Away from non-positively curved groups, it is known that every elementary amenable group that has exponential growth also has uniform exponential growth [72], and the same is true for any finitely generated subgroup of $\mathrm{GL}_n(\mathbb{C})$ that has exponential growth [34].

For a while it was an open question as to whether every group with exponential growth also had uniform exponential growth, however this was shown to be false by Wilson.

Theorem 3.1.8. [86] *There exists a finitely generated group with non-uniform exponential growth.*

This means that, unlike with exponential growth, we cannot check for uniform exponential growth by simply considering a single generating set. This, and the fact that it is unknown whether uniform exponential growth is a quasi-isometry invariant or not, makes it a much harder property to prove in most cases.

One important class of groups, which we will discuss in detail later, are the *acylindrically hyperbolic groups* (see Chapter 4). Briefly, these are groups that admit a certain type of action on hyperbolic spaces, and are a generalisation of hyperbolic and relatively hyperbolic groups. The class of acylindrically hyperbolic groups is wide, for example it includes right-angled Artin groups, most mapping class groups of surfaces, $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ for $n \geqslant 2$, and one-relator groups with at least three generators [73, 40]. Groups in this class do not have to be finitely generated in general. It is known that every finitely generated acylindrically hyperbolic group has exponential growth (see [49], for example), however the following question is still open.

Question 4. *Does every finitely generated acylindrically hyperbolic group have uniform exponential growth?*

This may well fail to be the case. A common method of proving uniform exponential growth is to find elements in balls of bounded size that generate free subsemigroups, and then apply the following result.

Lemma 3.1.9. *Let G be a finitely generated group. Suppose that there exists $N \in \mathbb{N}$ such that for every finite generating set S of G there exist $a, b \in B_S(N)$ such that a and b generate a free subsemigroup of G. Then G has uniform exponential growth.*

Proof. Let S be a finite generating set of G, and let $\{a, b\} \subset B_S(N)$ be elements that generate a free subsemigroup. For any $n \in \mathbb{N}$ we have that $|\{a, b\}^n| = 2^n$, and therefore $|B_S(nN)| = |B_S(N)^n| \geqslant 2^n$. We can then see that

$$\omega(G, S) = \lim_{n\to\infty} |B_S(n)|^{\frac{1}{n}} = \lim_{n\to\infty} |B_S(nN)|^{\frac{1}{nN}} \geqslant 2^{\frac{1}{N}}.$$

As N is independent of the choice of generating set S, we get that $\omega(G) \geqslant 2^{\frac{1}{N}} > 1$, and therefore G has uniform exponential growth. □

In acylindrically hyperbolic groups this is usually done by finding *loxodromic elements* (see Chapter 4), which in this case are simply the group elements with unbounded orbits in the space. However a result of Minasyan and Osin means that it is not possible to use this method for all acylindrically hyperbolic groups.

Theorem 3.1.10. [70] *There exists a finitely generated acylindrically hyperbolic group G such that for every $N \in \mathbb{N}$ there exists a finite generating set S of G such that no element of $B_S(N)$ is loxodromic.*

More generally, they proved that every countable family of countable acylindrically hyperbolic groups has a common finitely generated acylindrically hyperbolic quotient. This means that if the infimum of $\omega(G)$ over all finitely generated acylindrically hyperbolic groups is one, then their result would provide an example of an acylindrically hyperbolic group with non-uniform exponential growth, by the following standard result.

Lemma 3.1.11. *Let G be a finitely generated group, and let N be a normal subgroup of G. Then $\omega(G) \geqslant \omega(G/N)$.*

Proof. Let S be a finite generating set of G. Then $\overline{S} = \{sN : s \in S\}$ is a finite generating set of G/N, and $|B_S(n)| \geqslant |B_{\overline{S}}(n)|$, so the result follows by taking limits and infimums of both sides. □

On the other hand, uniform exponential growth is known for many examples of acylindrically hyperbolic groups, such as the previously mentioned mapping class groups [2]. More generally, in [1] it was shown that any virtually torsion free acylindrically hyperbolic group that is also a hierarchically hyperbolic group has uniform

exponential growth. In Section 3.4 we will introduce a type of growth that is stronger than uniform exponential growth, and in subsequent chapters we will show that it also holds for some of the examples above.

3.1.2 Growth of subgroups

A common theme in group theory is to try to understand a group by getting a better understanding of its subgroups, and the same is true for growth questions. In fact, much of the rest of this thesis will be focused on trying to understand the growth of subgroups for certain classes of groups.

A first easy observation is that subgroups can grow arbitrarily slowly compared to the group that they are contained in. Some simple examples include finite subgroups of infinite groups, or the fact that a group with non-exponential growth such as $\mathbb{Z}$ can be a subgroup of a group with exponential growth, such as the free group on two generators. If a subgroup grows quickly, however, we would expect the group to also grow quickly.

Lemma 3.1.12. *Let G be a finitely generated group, and let H be a finitely generated subgroup of G. If H has exponential growth, then so does G.*

Proof. Let S_G be a finite generating set of G, and S_H be a finite generating set of H. As H has exponential growth, we have that $\omega(H, S_H) > 1$. Let $S = S_G \cup S_H$. Then S is a finite generating set of G, and $B_{S_H}(n) \subset B_S(n)$. Hence $|B_{S_H}(n)| \leqslant |B_S(n)|$, so $\omega(G, S) \geqslant \omega(H, S_H) > 1$. □

On the other hand, if a group has a subgroup with uniform exponential growth, then it does not necessarily follow that the group itself has uniform exponential growth. In fact, Wilson's example of a group with non-uniform exponential growth contains a free subgroup [86]. For a subgroup to pass uniform exponential growth to the group, then for any generating set of the group we have to be able to find generators of the subgroup within a ball of bounded radius, as in Lemma 3.1.9. One situation in which this is always true is when the subgroup in question has finite index.

Proposition 3.1.13. [80, Proposition 3.3] *Let G be a finitely generated group, and let H be a finitely generated finite index subgroup of G. If H has uniform exponential growth, then so does G.*

Remark 3.1.14. It is still an open question as to whether uniform exponential growth passes from the group to its finite index subgroups.

Another situation in which uniform exponential growth can pass to a supergroup is when we are considering direct products. These will be an important example for us later in this chapter. We begin by considering exponential growth in the factors.

Lemma 3.1.15. *Let $H \leqslant G_1 \times \cdots \times G_m$ be a finitely generated group, and let H_i be the projection of H to G_i. Let S be a finite generating set of H, and let S_i be the projection of S to H_i. Then*

$$\sup_{1 \leqslant i \leqslant m} \omega(H_i, S_i) \leqslant \omega(H, S) \leqslant \omega(H_1, S_1) \cdots \omega(H_m, S_m).$$

Proof. We first note that $B_{S_i}(n)$ is the projection of $B_S(n)$ to H_i, so $|B_{S_i}(n)| \leqslant |B_S(n)|$. We also have that $B_S(n) \subset B_{S_1}(n) \times \cdots \times B_{S_m}(n)$, which means that $|B_S(n)| \leqslant |B_{S_1}(n)| \cdots |B_{S_m}(n)|$. Taking the limit of the $\frac{1}{n}$th power of these inequalities, we get our result. □

Corollary 3.1.16. *Let $H \leqslant G_1 \times \cdots \times G_m$ be a finitely generated group, and let H_i be the projection of H to G_i. If any factor H_i has exponential growth, then H has exponential growth. If no factor has exponential growth, then H does not have exponential growth.*

We can now use Lemma 3.1.15 to get a similar result for uniform exponential growth in direct products, although this time we restrict ourselves to considering the whole of the direct product, rather than all of its subgroups. This is because the generating sets of a subgroup of a direct product cannot necessarily give us a complete picture of the generating sets of its projections.

Lemma 3.1.17. *Let $G = G_1 \times \cdots \times G_m$ be a direct product of finitely generated groups. Then*

$$\sup_{1 \leqslant i \leqslant m} \omega(G_i) \leqslant \omega(G) \leqslant \omega(G_1) \cdots \omega(G_m).$$

Proof. Let $\mathcal{S} = \{S \subset G \ : \ S \text{ is a finite generating set of } G\}$. Then Lemma 3.1.15 gives us that

$$\begin{aligned}
\omega(G) &= \inf_{S \in \mathcal{S}} \omega(G, S) \\
&\leqslant \inf_{S \in \mathcal{S}} (\omega(G_1, S_1) \cdots \omega(G_m, S_m)) \\
&= \inf_{S_1} \omega(G_1, S_1) \cdots \inf_{S_m} \omega(G_m, S_m) \\
&= \omega(G_1) \cdots \omega(G_m).
\end{aligned}$$

Here we used that for any collection of finite generating sets S_i of G_i, we can let $S = S_1 \times \cdots \times S_m$, and this will be a finite generating set of G such that S_i is the projection of S to G_i.

As every finite generating set of G projects to a finite generating set of G_i, we get that $\inf_{S_i} \omega(G_i, S_i) \leqslant \inf_{S \in \mathcal{S}} \omega(G, S)$ for every i, so

$$\sup_{1 \leqslant i \leqslant m} \omega(G_i) = \sup_{1 \leqslant i \leqslant m} (\inf_{S_i} \omega(G_i, S_i)) \leqslant \inf_{S \in \mathcal{S}} \omega(G, S) = \omega(G).$$

□

Corollary 3.1.18. *Let $G = G_1 \times \cdots \times G_m$ be a direct product of finitely generated groups. If any factor G_i has uniform exponential growth, then G has uniform exponential growth. If no factor has uniform exponential growth, then G does not have uniform exponential growth.*

Remark 3.1.19. Suppose $H \leqslant G_1 \times \cdots \times G_m$ is a finitely generated group, and let H_i be the projection of H to G_i. Suppose that some H_i has uniform exponential growth. The final part of the proof of Lemma 3.1.17 still applies here, so $\sup_{1 \leqslant i \leqslant m} \omega(H_i) \leqslant \omega(H)$, and therefore H has uniform exponential growth.

We will revisit factors of direct products in Section 3.6, but for now we go back to the more general case of subgroups with uniform exponential growth. We know that for every such subgroup H, we have that $\omega(H) > 1$. However, much as with the generating sets for a group with exponential growth, there is no guarantee in general that there will be a lower bound on $\omega(H)$ for every subgroup with uniform exponential growth.

Definition 3.1.20. Let G be a group, and let $\mathcal{H}$ be a collection of finitely generated subgroups of G. We say that the collection $\mathcal{H}$ has *uniform uniform exponential growth* if $\inf_{H \in \mathcal{H}} \omega(H) > 1$.

This is an even harder question to answer than just uniform exponential growth, however there are groups for which it is known, including hyperbolic groups [5], linear groups [17], and mapping class groups [62].

Another method for trying to understand the subgroups of a group is to find some kind of dichotomy for those subgroups. One of the best known examples of this is the Tits alternative.

Definition 3.1.21. A group satisfies the *Tits alternative* if every one of its subgroups is either virtually solvable, or contains a free subgroup on at least two generators.

This was most famously shown by Tits for finitely generated linear groups in [84]. Framing this in terms of growth, we know that groups containing free subgroups have exponential growth. On the other hand, there are many virtually solvable groups that also have exponential growth, despite containing no free subgroups. This means that, although the Tits alternative does give us a dichotomy of subgroups, it does not in general give us a dichotomy for their growth.

On the other hand, there are many groups that satisfy the Tits alternative such that their virtually solvable subgroups are actually contained in a smaller subclass, such as the virtually abelian groups. In this case none of these subgroups would have exponential growth (more specifically they would have polynomial growth). Examples of groups that satisfy this version of the Tits alternative include hyperbolic groups [43, 41] and mapping class groups [50, 68].

We will now leave exponential and uniform exponential growth as background motivation, and consider the more general question of product set growth.

3.2 Growth of arbitrary subsets

Recall that if U is a subset of a group G, then we use U^n to denote the product set $\{u_1 \cdots u_n : u_1, \ldots, u_n \in U\} \subset G$. In the last section we considered the case that G is finitely generated, S is a finite generating set of G, and $U = B_S(1)$. We will now consider what we can say when U is an arbitrary finite subset of G. In particular, this means that we are dropping the assumption that U is symmetric, as well as the assumption that it contains the identity element.

The loss of these assumptions means that the growth of these sets is generally less well behaved than the growth of balls. For instance, if U does not contain the identity element then U^n is not necessarily a subset of U^{n+1}.

Example 3.2.1. If $U = \{a\}$ is a subset of a group, and $a \neq e$, then $U \not\subset U^2$.

Remark 3.2.2. If U is symmetric, however, then we do at least have that $U^n \subset U^{n+2}$, as in this case U^2 contains the identity element.

On the other hand, there are some properties that are retained, such as the fact that growth is monotonic.

Lemma 3.2.3. *Let U be a finite subset of a group G. Then $|U^n| \leqslant |U^{n+1}|$ for every $n \in \mathbb{N}$.*

Proof. For any $u \in U$ and $v, w \in U^n$, we have that $uv = uw$ if and only if $v = w$, so $|U^n| = |uU^n|$, and $uU^n \subset U^{n+1}$. □

Another property worth noting is that the growth of subsets remains stable under conjugation, as $|(gUg^{-1})^n| = |gU^ng^{-1}| = |U^n|$. We also retain from the growth of balls that $\log(|U^n|)$ as a function of n is subadditive.

Lemma 3.2.4. *Let U be a finite subset of a group G. Then $|U^{n+m}| \leqslant |U^n||U^m|$ for every $n, m \in \mathbb{N}$.*

Proof. Any $u \in U^{n+m}$ can be written as $u = vw$ for some $v \in U^n, w \in U^m$. □

Corollary 3.2.5. *Let U be a finite subset of a group G. Then $|U^{nm}| \leqslant |U^n|^m$ for every $n, m \in \mathbb{N}$.*

One of the most important properties when looking at the growth of balls was the fact that the asymptotic behaviour of this growth was determined by the group generated by the ball in question. When considering the growth of sets, we can see that this is still the case when U is symmetric.

Lemma 3.2.6. *Let U be a finite symmetric subset of a group G. Then $|U^n| \leqslant |B_U(n)| \leqslant 2|U^n|$.*

Proof. The lower bound follows from the fact that $U^n \subset B_U(n)$, recalling that $B_U(n) = U \cup U^{-1} \cup \{e\}$, where e is the identity element of G, and U^{-1} is the set of inverses of the elements of U.

For the upper bound, recall that U being symmetric means that $U^n \subset U^{n+2}$ for every $n \in \mathbb{N}$. Therefore

$$B_U(n) = \{e\} \cup U \cup \cdots \cup U^n = U^{n-1} \cup U^n$$

for any $n \in \mathbb{N}$. This means that $|B_U(n)| \leqslant |U^{n-1}| + |U^n|$, so as $|U^{n-1}| \leqslant |U^n|$ we get that $|B_U(n)| \leqslant 2|U^n|$. □

On the other hand, when we drop the assumption of symmetry we find that two sets contained in the same minimal subgroup do not necessarily share the same asymptotic growth. Several simple examples of this are given below. We can see this as the non-symmetric sets telling us something about the growth of the subsemigroup that they are contained in, rather than the subgroup. The behaviour of subsemigroups is not as well understood in general as the behaviour of subgroups.

Example 3.2.7. Let $G = \mathbb{Z}$, $U = \{1\}$, and $V = \{1, 2\}$. Then $\langle U \rangle = \langle V \rangle = \mathbb{Z}$, however $U^n = \{n\}$ and $V^n = \{n, \ldots, 2n\}$. Therefore $|U^n| = 1$, while $|V^n| = n + 1 \to \infty$.

Example 3.2.8. Let $G = \mathbb{Z} \times \mathbb{Z}_2$, $U = \{(1, 0), (0, 1)\}$, and $V = \{(1, 0), (1, 1)\}$. Then $\langle U \rangle = \langle V \rangle = \mathbb{Z} \times \mathbb{Z}_2$, however $U^n = \{(0, n \mod 2), (1, n - 1 \mod 2), \ldots, (n, 0)\}$ and $V^n = \{(n, 0), (n, 1)\}$. Therefore $|U^n| = n + 1 \to \infty$, while $|V^n| = 2$.

Example 3.2.9. Let $G = \mathbb{Z}^2$, $U = \{(1, 0), (0, 1)\}$, and $V = U \cup U^{-1}$. Then $\langle U \rangle = \langle V \rangle = \mathbb{Z}^2$, however $U^n = \{(p, q) : p, q \geqslant 0 \text{ and } p+q = n\}$ and $V^n = \{(p, q) : |p|+|q| \leqslant n \text{ and } p + q \equiv n \mod 2\}$. Therefore $|U^n| = n + 1$, while $|V^n| = (n + 1)^2$.

For a general group G and a general finite subset $U \subset G$, the most we can say about $|U^n|$ in terms of $|U|$ is that $|U| \leqslant |U^n| \leqslant |U|^n$. We can see that the upper bound is possible by taking G to be the free group on generators a and b, and considering the case where $U = \{a, b\}$. We can see that the lower bound is possible by some of the examples above.

We can also see that the asymptotic behaviour of the growth of $|U^n|$ is bounded above by the growth of $\langle U \rangle$, as the inequality $|U^n| \leqslant |B_U(n)|$ proved in Lemma 3.2.6 was not dependent on symmetry. This in turn means that the growth of $|U^n|$ is also bounded above by the growth of G, if G is finitely generated.

In terms of lower bounds, in the case that our set U is a ball we can note when $\langle U \rangle$ is infinite the growth of $|U^n|$ is strictly monotonic, rather than just monotonic. This is easy to show using the fact that $B_S(n) \subset B_S(n + 1)$, as then $|B_S(n)| = |B_S(n + 1)|$ implies that $B_S(n) = B_S(n+1)$, so there are no words of length greater than n, which would mean that the group $\langle U \rangle$ is finite. This implication is not true for arbitrary finite subsets, as we have already noted that a set with bounded growth can generate an infinite group. We can however show that if a set does have unbounded growth, then that growth will also be strictly monotonic.

Lemma 3.2.10. *Let U be a finite subset of a group G. Then $|U^n| \to \infty$ if and only if $|U^{n+1}| > |U^n|$ for every $n \in \mathbb{N}$.*

Proof. As $1 \leqslant |U| \leqslant |U^2| \leqslant \ldots$ we have that

$$\begin{aligned} |U^n| \not\to \infty &\iff \text{there exists } M \in \mathbb{N} \text{ such that } |U^n| \leqslant M \text{ for all } n \in \mathbb{N} \\ &\iff \text{there exists } m \in \mathbb{N} \text{ such that } |U^m| = |U^{m+n}| \text{ for all } n \in \mathbb{N} \\ &\iff \text{there exists } m \in \mathbb{N} \text{ such that } |U^m| = |U^{m+1}|. \end{aligned}$$

The reverse implication of the final equivalence comes from the fact that if $|U^m| = |U^{m+1}|$, then for every $u \in U$ we will have that $|uU^m| = |U^m| = |U^{m+1}|$. As $uU^m \subset U^{m+1}$, we therefore get that $uU^m = U^{m+1}$.

Hence we can rewrite any element of U^{m+n} so that it starts with u, so $U^{m+n} \subset uU^{m+n-1} \subset U^{m+n}$, so $|U^{m+n}| = |uU^{m+n-1}| = |U^{m+n-1}|$. By induction we therefore have that $|U^{m+n}| = |U^m|$.

We can now conclude that

$$|U^n| \to \infty \iff |U^{n+1}| > |U^n| \text{ for all } n \in \mathbb{N}$$

□

Remark 3.2.11. This tells us that if $|U^n| \to \infty$, then we get a lower bound $|U^n| \geqslant |U| + (n-1) \geqslant n$ for every $n \in \mathbb{N}$.

As we mentioned before, when U is a ball we know that $|U^n| \to \infty$ whenever $\langle U \rangle$ is infinite. Lemma 3.2.6 tells us that this implication is also true when U is symmetric. If U is an arbitrary finite set, however, then we cannot do much better than to say that $|U^n| \to \infty$ whenever $\langle U \rangle$ is not virtually cyclic. The method of the following proof was suggested by Jan Steinebrunner.

Lemma 3.2.12. *Let U be a finite subset of a group G. If U is not contained in a virtually cyclic subgroup, then $|U^n| \to \infty$.*

Proof. Suppose $|U^n| \not\to \infty$. Then we know from the proof of Lemma 3.2.10 that there exists $m \in \mathbb{N}$ such that $|U^m| = |U^{m+n}|$ for all $n \in \mathbb{N}$. In particular this means that $uU^m = U^{m+n} = U^m u$ for any $u \in U^n$, and so we also have that $U^m u^{-1} = u^{-1} U^m$ for any $u \in U^n$.

Pick some $u \in U$. Let $H = \bigcup_{n \in \mathbb{Z}} u^n U^m$. We want to show that H is a virtually cyclic subgroup of G that contains U. First we check that it is a subgroup. Let $x, y \in H$. We want to show that $xy \in H$.

We can write $x = u^n v$ and $y = u^k w$ where $n, k \in \mathbb{Z}$ and $v, w \in U^m$. Then $xy = u^n v u^k w$. Note that $vu^k \in U^m u^k = u^k U^m$, so $vu^k = u^k v'$ for some $v' \in U^m$. Now $v'w \in U^{2m} = u^m U^m$, so $v'w = u^m w'$ for some $w' \in U^m$. Hence $xy = u^n v u^k w = u^{n+k} v' w = u^{n+k+m} w' \in u^{n+k+m} U^m \subset H$. Therefore H is closed under multiplication.

We now want to show that H is closed under inverses. Let $x \in H$. Then we can write $x = u^n v$ for some $n \in \mathbb{Z}$ and $v \in U^m$. Hence $x^{-1} = v^{-1} u^{-n}$. Note that $vU^m = U^{2m}$, so for some $w \in U^m$ we have that $vw = u^{2m}$. Therefore $v^{-1} = w u^{-2m} \in U^m u^{-2m} = u^{-2m} U^m \subset H$. We also have that $u^{-n} \in u^{-(n+m)} U^m \subset H$, so as H is

closed under multiplication we can conclude that $x^{-1} = v^{-1}u^{-n} \in H$. Hence H is closed under inverses, so H is a subgroup of G.

It is easy to see that H is virtually cyclic, as H can also be written as $\bigcup_{v \in U^m} \langle u \rangle v$, and U^m is a finite set. It now only remains to check that H contains U.

We can see that $u = u^{-(m-1)}u^m \in u^{-(m-1)}U^m \subset H$. Recall that for any $v \in U$ and $n \in \mathbb{Z}$ we have that $u^n U^m = v^n U^m$, so H can also be written as $\bigcup_{n \in \mathbb{Z}} v^n U^m$. Hence by the same reasoning as for u, we have that $v \in H$ for any $v \in U$. Therefore $U \subset H$, so U is contained in a virtually cyclic subgroup. □

Remark 3.2.13. In the case that $|U^n| \not\to \infty$, we get that for any $u \in U$ the cyclic group $\langle u \rangle$ is a finite index subgroup of $\langle U \rangle$. This means that either every element of U is torsion or none of them are, and these cases correspond exactly to $\langle U \rangle$ being either finite or virtually $\mathbb{Z}$. On the other hand, by Example 3.2.7, we could still have that $|U^n| \to \infty$ while U is contained in a virtually cyclic subgroup. In particular, we now know that if U contains both a torsion and a non-torsion element then we must have that $|U^n| \to \infty$, as in Example 3.2.8. The revere implication does not hold, however, as our previous examples have shown.

3.3 Uniform product set growth

We will continue to consider the growth of finite subsets for the remainder of this thesis. In particular we will be interested in getting a specific lower bound on $|U^n|$, where this lower bound is dependent on the size of U.

Definition 3.3.1. Let G be an infinite finitely generated group. We say that G has *uniform product set growth* if there exist $\alpha, \beta > 0$ such that for every finite (symmetric) generating set U of G, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Remark 3.3.2. If the group G in question was finite, then this would be trivially satisfied by letting $\alpha = \frac{1}{|G|}$. For this reason we will generally only be interested in determining this property for infinite groups.

It is not hard to check that uniform product set growth implies uniform exponential growth, even in the weaker form where we only ask that the inequality holds for symmetric generating sets.

Lemma 3.3.3. *Let G be a finitely generated group. Suppose there exist constants $\alpha, \beta > 0$ such that for every finite symmetric generating set U of G we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. Then either G is finite with $|G| \leqslant \frac{1}{\alpha}$, or G has uniform exponential growth.*

Proof. Suppose that G is finite, then G is a finite generating set of itself, and $G^n = G$ for every $n \in \mathbb{N}$. Therefore $|G| \geqslant (\alpha|G|)^{\beta n}$ for every $n \in \mathbb{N}$, so $|G| \leqslant \frac{1}{\alpha}$.

Now suppose that G is infinite, and let S be a finite generating set of G. We want to calculate a lower bound on $\omega(G,S) = \lim_{n\to\infty} \frac{1}{n}\log(|B_S(n)|)$. We begin by recalling that $B_S(n) = B_S(1)^n$, so for every $n \in \mathbb{N}$ we have that $|B_S(n)| \geqslant (\alpha|B_S(1)|)^{\beta n}$. If $|B_S(1)| \leqslant \frac{1}{\alpha}$ then our equation would tell us nothing about $|B_S(n)|$, however to calculate the limit it is sufficient to consider a subsequence, so we just have to start this subsequence with a large enough set.

For any $m \in \mathbb{N}$ we also have that $B_S(m)$ is also a finite generating set of G, and $B_S(m)^n = B_S(mn)$ for all $n \in \mathbb{N}$, so $|B_S(mn)| \geqslant (\alpha|B_S(m)|)^{\beta n}$. We can also see that $|B_S(m)| \geqslant m - 1 + |S|$, so choosing $m = \lceil \frac{1}{\alpha} \rceil + 1$ we get that

$$\alpha|B_S(m)| \geqslant \alpha(\lceil \frac{1}{\alpha} \rceil + |S|) \geqslant 1 + \alpha|S| > 1.$$

With this choice of m we now get that

$$|B_S(mn)|^{\frac{1}{mn}} \geqslant (\alpha|B_S(m)|)^{\frac{\beta}{m}} \geqslant (1+\alpha|S|)^{\frac{\beta}{\lceil 1/\alpha \rceil + 1}}.$$

Hence we have that

$$\omega(G,S) = \lim_{n\to\infty} |B_S(n)|^{\frac{1}{n}} = \lim_{n\to\infty} |B_S(mn)|^{\frac{1}{mn}} \geqslant (1+\alpha|S|)^{\frac{\beta}{\lceil 1/\alpha \rceil + 1}},$$

so as S was arbitrary, and $|S| \geqslant 1$, we get that

$$\omega(G) = \inf_{S\in\mathcal{S}} \omega(G,S) \geqslant (1+\alpha)^{\frac{\beta}{\lceil 1/\alpha \rceil + 1}} > 1.$$

□

Remark 3.3.4. If some finite generating set U of an infinite group satisfies $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, then this is enough to imply that the group has exponential growth. A set U contained in a group with subexponential growth can only satisfy this inequality in the trivial case that $|U| \leqslant \frac{1}{\alpha}$.

In fact, it is not hard to see that uniform product set growth is actually a strictly stronger property than uniform exponential growth. This is simply because uniform exponential growth only requires a common lower bound on the exponential growth rates of generating sets, whereas uniform product set growth requires that the lower bound increases as the size of the generating set increases. An example of this difference will be given in Section 3.6, where we consider a direct product where one of the factors has uniform exponential growth and the other does not.

We can now introduce the question that will be the main focus of the rest of this thesis.

Question 5. *For a group G, do there exist constants $\alpha, \beta > 0$ and a class of subgroups $\mathcal{H}$ such that for every finite (symmetric) $U \subset G$ where $\langle U \rangle \notin \mathcal{H}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$?*

Ideally we would like to choose $\mathcal{H}$ in such a way as to produce a dichotomy for the infinite finitely generated subgroups of G. That is, for any infinite finitely generated $H \leqslant G$, either $H \in \mathcal{H}$ and does not have uniform product set growth for any choice of constants, or every finite set U generating H satisfies $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. A possibly weaker question would be to allow α and β to vary based on the subgroup in question, however in this thesis we will only be interested in the stronger case, as given by Question 5.

One thing to note here is that we have allowed for two versions of this question, one where we consider all finite subsets, and one where we only consider the finite subsets that are also symmetric. In the non-symmetric case a positive answer to Question 5 would not only tell us something about the growth of subgroups of G, it would also tell us something about the growth of subsemigroups. On the other hand, the lack of symmetry may make Question 5 much harder to answer. We therefore also allow for the slightly weaker version of this question, where symmetry is required. In either case the aim will still be to get a dichotomy of subgroups.

When thinking about how to divide our subgroups, one fact that we already know from Lemma 3.3.3 is that uniform product set growth implies uniform exponential growth. Given this, we can think of Question 5 as being analogous to the question of uniform uniform exponential growth for a collection of subgroups. In both cases we are asking if every subgroup in our collection has the same growth with the same associated constants. In fact, a positive answer to Question 5 tells us that the finitely generated subgroups that do not lie in $\mathcal{H}$ will have uniform uniform exponential growth. This immediately follows from Lemma 3.3.3, as the exponential growth rate of such a subgroup is only dependent on α and β.

We can also compare Question 5 to the question of finding free subgroups or free subsemigroups of a group, especially as the attempt to find a dichotomy of subgroups may remind us of the Tits alternative. We cannot in general expect the division of subgroups to be the same in each dichotomy, given that uniform product set growth is strictly stronger than uniform exponential growth, and the existence of a free subgroup only implies exponential growth. On the other hand, a potential approach to proving Question 5 could be to show that there exist some $N \in \mathbb{N}$, $\alpha, \beta > 0$ such that for a given set U, the set U^N contains generators of a free subsemigroup of rank $(\alpha|U|)^{\beta}$. We will see this idea appear later, in Section 4.3. This is similar to the

usual method of proving uniform exponential growth, which we recall is to show that there exists some $N \in \mathbb{N}$ such that for any generating set S, the ball $B_S(N)$ contains generators of a free subsemigroup. When these generators in a bounded ball give us a free subgroup, this is sometimes referred to as a quantitative version of the Tits alternative.

The rest of this chapter will be focused around uniform product set growth, especially in relation to Question 5. The next section gives a brief survey of previous results in the area, then Section 3.5 gives some consequences of uniform product set growth. We will finish this chapter with Section 3.6, by discussing some cases in which we can determine whether a subgroup has uniform product set growth or not. This will be helpful in later chapters, as we try to answer Question 5 for subgroups of mapping class groups.

3.4 Known results for product set growth

In this section we will go over some known results about product set growth in groups, specifically those that relate in some way to Question 5. In the cases that these results are about groups that do not appear elsewhere in this thesis, we will usually not give definitions for them, and instead refer to the relevant papers.

The origins of Question 5 came from the study of growth in finite and abelian groups, where there is often interest in estimating the size of $|U^2|$ and $|U^3|$. One example of this is given by Helfgott in [46], where it was shown that there exist $c, \delta > 0$ such that for certain $U \subset SL_2(\mathbb{Z}/p\mathbb{Z})$ with p prime we have that $|U^3| \geqslant c|U|^{1+\delta}$. Some of the techniques used by Helfgott in this paper were adapted to infinite groups by Chang, who showed the following.

Theorem 3.4.1. [22] *There exist constants $c, \delta > 0$ such that for any finite subset U of $SL_2(\mathbb{C})$, at least one of the following must hold:*

(1) $\langle U \rangle$ is metabelian or finite.

(2) $|U^3| \geqslant c|U|^{1+\delta}$.

Remark 3.4.2. In Chang's paper the first case is that U is contained in a virtually abelian subgroup, however in [19] it was shown that this condition should be replaced by a finite or metabelian subgroup.

As $F_2 \leqslant SL_2(\mathbb{C})$, the same result applies to any finite $U \subset F_2$ such that U is not contained in a cyclic subgroup.

Corollary 3.4.3. [22] *There exist constants $c, \delta > 0$ such that for any finite subset U of F_2, at least one of the following must hold:*

(1) $\langle U \rangle$ is cyclic.

(2) $|U^3| \geqslant c|U|^{1+\delta}$.

Chang believed that for free groups there should be a direct combinatorial proof of this fact. The first proof in this direction was given by Razborov, in the wider case of virtually free groups.

Theorem 3.4.4. [74] *Let G be a virtually free group. Then there exists a constant $d > 0$ such that for every finite subset U of G, at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually cyclic.

(2) $|U^3| \geqslant \frac{|U|^2}{\log(|U|)^d}$.

Both of these results were then further improved on by Safin in the case of free groups. In particular, Safin was able to completely answer Question 5 in this case.

Theorem 3.4.5. [76] *Let G be a free group. Then for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is isomorphic to $\mathbb{Z}$.

(2) $|U^n| \geqslant (\frac{1}{372}|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Remark 3.4.6. In Safin's paper this result is stated slightly differently. The first difference is that in [76], all of the cyclic subgroups are ruled out. However as any free group is torsion-free, these cyclic subgroups are only the trivial group and the subgroups isomorphic to $\mathbb{Z}$, and the trivial case is taken care of by the fact that any set containing a single element clearly satisfies the second inequality. The other difference is that the inequality in [76] is expressed in the form $|U^n| \geqslant c_n(|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for some constants $c_n > 0$, however a simple inspection of the argument tells us that we can rewrite it in the above form.

This result gives us the dichotomy of subgroups that we are looking for, as we know that $\mathbb{Z}$ does not have exponential growth, and so cannot have uniform product set growth. By an example of Safin, we can also see that this order of growth is the best that we can expect.

Example 3.4.7. [76] Let F be a free group, and let $a, b \in F$ be elements that do not commute. Then for some $k \in \mathbb{N}$ let $U = \{a, b, b^2, \dots, b^{k-1}\}$. We can see that $|U| = k$, and $|U^2| = 4k - 4 \leqslant 4k$. Note that $|U^n| = |U^{n-2}U^2| \leqslant |U^{n-2}||U^2| \leqslant |U^{n-2}|4k$ for $n \geqslant 3$. By induction we therefore get that $|U^{2n}| \leqslant (4k)^n$ and $|U^{2n+1}| \leqslant 4^n k^{n+1} \leqslant (4k)^{n+1}$. Hence we can see that $|U^n| \leqslant (4k)^{\lfloor \frac{n+1}{2} \rfloor} = (4|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Another interesting feature of Safin's paper was the method used to prove Theorem 3.4.5, as it turned out to be highly generalisable to other classes of groups. The first notable example of this was given by Button, who was able to prove a tripling result for free products.

Theorem 3.4.8. [19] *Let $G * H$ be a free product. Then for every finite $U \subset G * H$ at least one of the following must hold:*

(1) U is conjugate into one of the factors.

(2) $\langle U \rangle$ is infinite cyclic or infinite dihedral.

(3) $|U^3| \geqslant \frac{1}{7776}|U|^2$.

Button asked if this type of tripling applies to hyperbolic groups. This question was answered recently by Delzant and Steenbock, who were able to greatly generalise the methods of Safin and Button to extend Safin's result to hyperbolic groups.

Theorem 3.4.9. [31] *Let G be a hyperbolic group. Then there exists a constant $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Remark 3.4.10. The original statement of this theorem ruled out all virtually cyclic groups, however it is known that the size of finite subgroups in hyperbolic groups is bounded, see [12] for example. So long as α is chosen in such a way that $\frac{1}{\alpha}$ is larger than this bound, the above statement holds.

An immediate corollary of this is uniform uniform exponential growth for subgroups of hyperbolic groups, which was also proved in [16].

Corollary 3.4.11. [16, 31] *Every finitely generated subgroup of a hyperbolic group that is not virtually cyclic has uniform exponential growth. Moreover, they have uniform uniform exponential growth, with the positive lower bound depending only on the hyperbolic group.*

In fact, Delzant and Steenbock generalised these methods even further than this, to the extent that they were able to apply them to all groups with an acylindrical action on a hyperbolic space. A group action on a metric space (X, d) is said to be *acylindrical* if it is by isometries, and if for every $R \geqslant 0$ there exist $N > 0$ and $L > 0$ such that for every $x, y \in X$ with $d(x, y) \geqslant L$ we have that $|\{g \in G : d(x, gx) \leqslant R \text{ and } d(y, gy) \leqslant R\}| \leqslant N$. Intuitively, this is saying that if two points in the space are sufficiently far apart, only finitely many elements of the group can almost stabilise both of them. We will discuss this definition in more detail in Section 4.1.

As mentioned in Section 3.1, there is a wide class of groups with interesting acylindrical actions on hyperbolic spaces. The subgroup structure of such groups may be much more complicated than the subgroup structure of a hyperbolic group, so it is not generally sufficient to simply rule out the virtually cyclic subgroups when looking for uniform product set growth. In particular, if all of our information about a group comes from its acylindrical action, then it is hard to say anything about subsets whose image under the orbit map have small diameter.

Delzant and Steenbock ruled out such subsets by using a certain notion of displacement. If (X, d) is a hyperbolic space that a groups G acts on by isometries, and $U \subset G$ is finite, pick $x_0 \in X$ such that $\sum_{u \in U} d(x_0, ux_0)$ is (almost) minimised, then define the *displacement* of U to be $\lambda_0(U) = \max_{u \in U} d(x_0, ux_0)$ (for a more detailed definition see Definition 4.2.1). Their statement for acylindrically hyperbolic groups is then as follows.

Theorem 3.4.12. [31] *Let G be a group acting acylindrically on a hyperbolic space. Then there exist constants $K > 0$ and $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K \log_2(2|U|)$.

(3) $|U^n| \geqslant \left(\frac{\alpha}{\log_2^6(2|U|)}|U|\right)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Remark 3.4.13. The original version of this theorem also ruled out all virtually cyclic groups, however Proposition 4.4 in [31] tells us that any set that generates a finite subgroup must satisfy at least one of $|U| \leqslant \frac{1}{\alpha}$ or $\lambda_0(U) < K \log_2(2|U|)$.

In the case that the space being acted on is a simplicial tree, Delzant and Steenbock obtained another version of this theorem, where the logarithm terms do not appear. In Section 4.2 we will show that this also holds when the space being acted on is a quasi-tree.

Theorem 3.4.14. [31] *Let T be a simplicial tree. Let G be a group that acts on T by isometries such that for some $L > 0$ the pointwise stabiliser of every geodesic of length at least L is trivial. Then there exist constants $K > 0$ and $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K$.

(3) $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Remark 3.4.15. This is a version of acylindricity for simplicial trees, as defined by Sela in [77], which implies the usual notion of acylindricity [69]. Note that in [31] the definition of a simplicial tree allows the tree to have a constant edge length ρ such that $\rho \neq 1$.

As mentioned previously, the second case in the above two theorems can be thought of as capturing the subsets that keep certain points within a bounded distance. In Chapter 4 we will show that this condition can be overcome by considering the sets U such that U^k contains a loxodromic element, for some bounded $k \in \mathbb{N}$.

We end this section by referring the reader to a couple of recent papers on this topic. The first is by Coulon and Steenbock, which contains results on product set growth in Burnside groups.

Theorem 3.4.16. [27] *Let G be a non-cyclic torsion-free hyperbolic group, and for every $k \in \mathbb{N}$ let $G_k = \langle g^k : g \in G \rangle$. Then there exist $k_0 \in \mathbb{N}$ and $\alpha > 0$ such that for every odd $k \geqslant k_0$, and every finite $U \subset G/G_k$, we have that $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.*

Remark 3.4.17. The original statement of this result ruled out the case where $\langle U \rangle$ is a finite cyclic group, however as any cyclic group in G/G_k has order at most k, the statement above holds so long as $\alpha \leqslant \frac{1}{k}$.

Coulon and Steenbock were also able to obtain the following result for groups with acylindrical actions on hyperbolic spaces, which can be compared with Theorem 3.4.12. Here instead of displacement they use the *l^∞-energy* of a set U of isometries of a space X, which is given by $\lambda(U) = \inf_{x \in X} \sup_{u \in U} d(x, ux)$.

Remark 3.4.18. What Coulon and Steenbock refer to as l^∞-energy is also sometimes called the displacement of a set, see [16] for example.

Theorem 3.4.19. [27] *Let G be a group acting acylindrically on a hyperbolic space. Then there exist constants $C > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\lambda(U) \leqslant C$.

(2) $|U^n| \geqslant \left(\frac{1}{C\lambda(U)}|U|\right)^{\lfloor \frac{n+1}{2} \rfloor}$ *for every* $n \in \mathbb{N}$.

The other recent paper is by Cui, Jiang, and Yang, who proved uniform uniform exponential growth for certain subgroups of relatively hyperbolic groups. The lower bound they got on the exponential growth rate depended on the size of the generating set, which came from a product set growth type result. We leave the definition of an elementary subgroup in this scenario to their paper.

Theorem 3.4.20. [28] *Let G be a non-elementary relatively hyperbolic group. Then there exist constants $\alpha, \beta > 0$ such that for every finite symmetric $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ *is elementary.*

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ *for every* $n \in \mathbb{N}$.

Remark 3.4.21. This is not explicitly stated in their main theorem, however it can be seen in the final stages of their proof, where they show that for some $k, N_0 > 0$ there exists $T \subset B_U(nk)$ such that $|T| \geqslant \frac{1}{2N_0}|U|$, and $\langle T \rangle$ is a free group of rank $|T|$. We can then replace $|B_U(n)|$ with $|U^n|$ using Lemma 3.2.6.

3.5 Applications of uniform product set growth

In this section we will briefly give some properties that are implied by uniform product set growth. These are in addition to uniform exponential growth, which was discussed in Section 3.3. It was also mentioned in the introduction that linking growth rates with the size of the generating set in question is a key step in proofs about growth rates being well-ordered, which can be seen in [38] and [37]. This will come up again at the end of Section 5.2.

3.5.1 Helfgott type growth

As mentioned in Section 3.4, the study of product set growth has often previously focussed on estimating the size of double and triple products of sets. Here we give a link between uniform product set growth and the growth of triple products.

Definition 3.5.1. [19] An infinite group G has *Helfgott type growth* if there exist $c, \delta > 0$ such that any finite $U \subset G$ such that $\langle U \rangle$ is not virtually nilpotent satisfies $|U^3| \geqslant c|U|^{1+\delta}$.

It is known that if the size of the triple product is bounded, then the size of every higher product can also be bounded. The following is not explicitly stated in [46], however it can be directly obtained from the work there.

Lemma 3.5.2. [46] *Let G be a group, and let $U \subset G$ be finite. If $|U^3| \leqslant K|U|$ for some $K > 0$, then $|U^n| \leqslant (8K^3)^{n-2}|U|$ for every $n \geqslant 3$.*

Proof. Let $\overline{U} = U \cup U^{-1}$. It was shown in [46] that

$$\frac{|\overline{U}^n|}{|U|} \leqslant \left(\frac{|\overline{U}^3|}{|U|}\right)^{n-2}$$

for $n \geqslant 3$. Clearly $|U^n| \leqslant |\overline{U}^n|$, so it remains to bound $|\overline{U}^3|/|U|$ above by a function of $|U^3|/|U|$. By considering all possible terms in $\overline{U}^3$, we can see that

$$|\overline{U}^3| \leqslant 2(|UUU| + |UUU^{-1}| + |UU^{-1}U| + |U^{-1}UU|).$$

Here we use that $UU^{-1}U^{-1} = (UUU^{-1})^{-1}$, so they have the same cardinality. Similar logic applies to the other products.

In [46] it was shown that the cardinality of each of these products is bounded above by $|U|(|U^3|/|U|)^3$, and therefore

$$\frac{|\overline{U}^3|}{|U|} \leqslant 8\left(\frac{|U^3|}{|U|}\right)^3.$$

Combining these inequalities, we see that

$$\frac{|U^n|}{|U|} \leqslant \left(8\left(\frac{|U^3|}{|U|}\right)^3\right)^{n-2}.$$

If $|U^3| \leqslant K|U|$ for some $K \geqslant 1$, then substituting this into the above inequality gives us the result that we were looking for. □

When uniform product set growth is given in the form $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$, as with many of the results in the previous section, we automatically get Helfgott type growth by simply taking $c = \alpha^2$ and $\delta = 1$. However Lemma 3.5.2 allows us to also obtain Helfgott type growth in the case of the more general type of uniform product set growth, that is when $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. This observation, and the following proof, were given by Jack Button.

Proposition 3.5.3. *Let G be a group, and let $U \subset G$ be finite. If there exist $\alpha, \beta > 0$ such that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, then $|U^3| \geqslant (\frac{\alpha^\beta}{8})^{\frac{1}{3}}|U|^{1+\frac{\beta}{6}}$.*

Proof. This is trivially true if $|U| = 1$, as in this case our inequality tells us that $\alpha^\beta \leqslant 1$. We therefore assume that $|U| \geqslant 2$, and suppose that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, but $|U^3| < (\frac{\alpha^\beta}{8})^{\frac{1}{3}}|U|^{1+\frac{\beta}{6}}$. We can now take $(\frac{\alpha^\beta}{8})^{\frac{1}{3}}|U|^{\frac{\beta}{6}}$ to be our K from Lemma 3.5.2, and obtain

$$(\alpha|U|)^{\beta n} \leqslant |U^n| \leqslant (8K^3)^{n-2}|U| = (\alpha^\beta|U|^{\frac{\beta}{2}})^{n-2}|U| = \frac{(\alpha^\beta|U|^{\frac{\beta}{2}})^n|U|}{(\alpha^\beta|U|^{\frac{\beta}{2}})^2}$$

for every $n \geqslant 3$. We therefore have that

$$|U|^{\frac{\beta n}{2}} \leqslant \frac{|U|}{\alpha^{2\beta}|U|^\beta}$$

for every $n \geqslant 3$. However, as $|U| \geqslant 2$, we also have that $|U|^{\frac{\beta n}{2}}$ tends to infinity. This contradicts the existence of a fixed upper bound, as given in the above inequality, and hence we must have that $|U^3| \geqslant (\frac{\alpha^\beta}{8})^{\frac{1}{3}}|U|^{1+\frac{\beta}{6}}$. □

3.5.2 Approximate groups

One other possible motivation for studying product set growth is its link with approximate groups, as defined in [83].

Definition 3.5.4. Let G be a group, and let $k \geqslant 1$. A finite symmetric subset $U \subset G$ is a *k-approximate group* if there exists a finite subset $X \subset G$ such that $|X| \leqslant k$ and $U^2 \subset XU$.

It is standard that approximate groups have small tripling, which leads us to the following observation of Button.

Proposition 3.5.5. [19] *Let G be a group, and let $k \geqslant 1$. Suppose U is a k-approximate group in G satisfying $|U^3| \geqslant (\alpha|U|)^2$ for some $\alpha > 0$. Then $|U| \leqslant (\frac{k}{\alpha})^2$.*

Proof. Let $X \subset G$ be such that $|X| \leqslant k$ and $U^2 \subset XU$. Then $U^3 \subset XU^2 \subset X^2U$, so

$$|U^3| \leqslant |X^2U| \leqslant |X|^2|U| \leqslant k^2|U|.$$

We therefore have that

$$(\alpha|U|)^2 \leqslant |U^3| \leqslant k^2|U| \implies |U| \leqslant \left(\frac{k}{\alpha}\right)^2.$$

□

Using the same reasoning as Button, we can obtain a slightly different version of this result for sets satisfying the type of uniform product set growth inequality discussed in Section 3.3.

Proposition 3.5.6. *Let G be a group, and let $k \geqslant 1$, $\alpha, \beta > 0$. Suppose U is a k-approximate group in G satisfying $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. Then $|U| \leqslant \frac{k^{\lceil 2/\beta \rceil - 1}}{\alpha^2}$.*

Proof. Let $X \subset G$ be such that $|X| \leqslant k$ and $U^2 \subset XU$. Then $U^n \subset X^{n-1}U$ for every $n \in \mathbb{N}$, so

$$|U^n| \leqslant |X^{n-1}U| \leqslant |X|^{n-1}|U| \leqslant k^{n-1}|U|$$

for every $n \in \mathbb{N}$.

Let $n = \lceil 2/\beta \rceil$, then we have that

$$(\alpha|U|)^{\beta(\lceil 2/\beta \rceil)} \leqslant |U^{\lceil 2/\beta \rceil}| \leqslant k^{\lceil 2/\beta \rceil - 1}|U|.$$

Note that if $\alpha|U| \geqslant 1$ then the above inequality and the fact that $2 \leqslant \beta(\lceil 2/\beta \rceil)$ tell us that $(\alpha|U|)^2 \leqslant |U^{\lceil 2/\beta \rceil}|$. If $\alpha|U| < 1$ then the same conclusion follows immediately from the observation that $|U^{\lceil 2/\beta \rceil}| \geqslant 1$. We can therefore conclude that

$$(\alpha|U|)^2 \leqslant k^{\lceil 2/\beta \rceil - 1}|U| \implies |U| \leqslant \frac{k^{\lceil 2/\beta \rceil - 1}}{\alpha^2}.$$

□

3.6 Obstructions and sufficient conditions

In the final section of this chapter, we prove a collection of results about when certain classes of subgroups can or cannot have uniform product set growth. We will begin by showing that any group with infinite centre cannot have uniform product set growth. Our focus will then switch to subgroups of direct products, as there are cases where

the uniform product set growth property can pass from the factors to the subgroup, and cases where certain factors make uniform product set growth impossible. We will also show that uniform product set growth passes to finite index supergroups. These tools will be vital in Chapter 5, when attempting to answer Question 5 for right-angled Artin groups and mapping class groups.

3.6.1 Infinite order central subgroups

So far the only subgroups we have completely ruled out when attempting to find those with uniform product set growth have been the subgroups that do not have uniform exponential growth (see Lemma 3.3.3). This is enough to get us a dichotomy in cases like Theorem 3.4.9, where the only subgroups that are not shown to have uniform product set growth are the virtually cyclic ones, as these are exactly the subgroups of hyperbolic groups that do not have uniform exponential growth.

On the other hand, a subgroup having uniform exponential growth is not generally enough to guarantee that it has uniform product set growth. One particular obstacle is when the group in question has an infinite central subgroup. The intuitive reason as to why such subgroups are a problem is that we can pick a generating set U of our group G such that an arbitrarily large proportion of it is in the centre $Z(G)$. As $Z(G)$ is an abelian subgroup, which does not have exponential growth, this part of U has minimal impact on the long term growth of U^n. This means that we cannot link the growth of U^n with the size of U in a meaningful way.

We can in fact show that there is a slightly more general obstruction than just infinite central subgroups. We will first consider any infinite normal subgroup without exponential growth, where this subgroup is a union of finite conjugacy classes. The obstruction given by central subgroups was initially pointed out by Jack Button, and the extension below came from discussions with Ric Wade.

Proposition 3.6.1. *Let G be a finitely generated group, and suppose that it has an infinite normal subgroup H, where H does not have exponential growth, and for every $h \in H$ the conjugacy class of h in G is finite. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. Let G and H be as above, and note that G must be infinite as it has an infinite subgroup. Fix constants $\alpha, \beta > 0$, and let U be some finite generating set of G. As H is infinite, we can pick a finite set $V' \subset H$ such that $|V'| > \frac{1}{\alpha}|U|^{\frac{1}{\beta}}$. Let V be the union of conjugacy classes $\bigcup_{v \in V'} \{gvg^{-1} : g \in G\}$. As H is a normal subgroup

of G, and conjugacy classes are finite, we have that V is a finite subset of H, and $|V| \geqslant |V'| > \frac{1}{\alpha}|U|^{\frac{1}{\beta}}$.

Let $W = U \cup V$. Then W is a finite generating set of G. Suppose that $|W^n| \geqslant (\alpha|W|)^{\beta n}$ for every $n \in \mathbb{N}$. Note that $|W| \geqslant |V|$, so $|W^n| \geqslant (\alpha|V|)^{\beta n}$.

We can also note that, as V is closed under conjugation by elements of G, we have that $W^n = U^n \cup U^{n-1}V \cup \cdots \cup UV^{n-1} \cup V^n$. We can see this by considering $w_1, \ldots, w_n \in W$, and supposing that $w_k \notin U$. Then $w_k \in V$, and $w_1 \cdots w_n = w_1 \cdots w_{k-1} w_{k+1} \cdots w_n w'_k$, where $w'_k = (w_{k+1} \cdots w_n)^{-1} w_k w_{k+1} \cdots w_n$. As this is a conjugate of $w_k \in V$, we have that $w'_k \in V$. We can therefore say that $|W^n| \leqslant (n+1)|U^n||V^n| \leqslant (n+1)|U|^n|V^n|$.

By our assumption, we therefore have that $|V^n| \geqslant \frac{(\alpha|V|)^{\beta n}}{(n+1)|U|^n}$. Hence

$$\begin{aligned}
\omega(\langle V \rangle, V) &= \lim_{n\to\infty} |B_V(n)|^{\frac{1}{n}} \\
&\geqslant \lim_{n\to\infty} |V^n|^{\frac{1}{n}} \\
&\geqslant \lim_{n\to\infty} \frac{(\alpha|V|)^{\beta}}{(n+1)^{\frac{1}{n}}|U|} \\
&= \frac{(\alpha|V|)^{\beta}}{|U|} \\
&> 1.
\end{aligned}$$

Therefore $\langle V \rangle$ has exponential growth. However H does not have exponential growth, and $\langle V \rangle \leqslant H$, so this is a contradiction. Hence $|W^n| < (\alpha|W|)^{\beta n}$ for some $n \in \mathbb{N}$. □

Proposition 3.6.2. *Let G be a finitely generated group, and suppose that the centre of G is infinite. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. This does not necessarily follow directly from Proposition 3.6.1, as even a finitely presented group can have a centre that is not finitely generated [75]. However the same proof structure works in this case, where we replace H by $Z(G)$, and note that the conjugacy classes have size one. The only difference is that the contradiction comes from the fact that any subgroup of $Z(G)$ will be an abelian group, so will not have exponential growth. □

An important example of where the conditions in Proposition 3.6.2 are satisfied is when we have a direct product with a free abelian group.

Corollary 3.6.3. *Let G be a finitely generated subgroup of $H \times \mathbb{Z}^m$, where the projective map from G to H is not injective. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. As the projective map is not injective, we can find $g \in H$ and $a, b \in \mathbb{Z}^m$ such that $a \neq b$ and $(g, a), (g, b) \in G$. Hence $(e, a^{-1}b) \in G$, and so $\{(e, (a^{-1}b)^k) : k \in \mathbb{Z}\}$ is an infinite subgroup of $Z(G)$. The conclusion then follows by applying Proposition 3.6.2. □

3.6.2 Direct products with non-exponential factors

In Corollary 3.6.3 we showed that certain subgroups of direct products with $\mathbb{Z}$ do not have uniform product set growth. This is clearly not necessarily true for all subgroups of direct products.

Example 3.6.4. Consider $F_2 \times \mathbb{Z}$, where $F_2 = \langle a, b \rangle$. Let $H = \langle (a, 1), (b, 1) \rangle$. Then H is a subgroup of $F_2 \times \mathbb{Z}$ with non-trivial projection to $\mathbb{Z}$, however for each $g \in F_2$ there exists a unique $n \in \mathbb{Z}$ such that $(g, n) \in H$. In other words $H \cong F_2$, and so H has uniform product set growth by Theorem 3.4.5.

We would like to understand when direct products and their subgroups have uniform product set growth, and when they do not. We begin by considering a direct product where one of the factors does not have uniform exponential growth.

Proposition 3.6.5. *Let $H_1 \times H_2$ be a direct product of groups such that H_2 is infinite and does not have uniform exponential growth. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of $H_1 \times H_2$ such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. Fix constants $\alpha, \beta > 0$. Suppose that for every finite generating set U of $H_1 \times H_2$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. We will show that this gives a contradiction to H_2 not having uniform exponential growth.

We fix a finite $V_1 \subset H_1$ such that $\langle V_1 \rangle = H_1$, and then let $V_2 \subset H_2$ be finite with $\langle V_2 \rangle = H_2$. Let $B = B_{V_1}(1) \times B_{V_2}(1)$, then we have that $\langle B \rangle = H_1 \times H_2$. We note that $B^n = B_{V_1}(n) \times B_{V_2}(n)$ for every $n \in \mathbb{N}$, so $|B^n| = |B_{V_1}(n)||B_{V_2}(n)|$.

As B generates $H_1 \times H_2$, by our assumption $|B^n| \geqslant (\alpha|B|)^{\beta n}$ for every $n \in \mathbb{N}$, and every possible choice of V_2. This means that for every $n \in \mathbb{N}$ we have that

$$\begin{aligned}
|B^n| \geqslant (\alpha|B|)^{\beta n} &\implies |B_{V_1}(n)||B_{V_2}(n)| \geqslant (\alpha|B_{V_1}(1)||B_{V_2}(1)|)^{\beta n} \\
&\implies \lim_{n\to\infty} (|B_{V_1}(n)||B_{V_2}(n)|)^{\frac{1}{n}} \geqslant (\alpha|B_{V_1}(1)||B_{V_2}(1)|)^{\beta} \\
&\implies \omega(H_1, V_1)\omega(H_2, V_2) \geqslant (\alpha|B_{V_2}(1)|)^{\beta} \\
&\implies \omega(H_2, V_2) \geqslant \frac{(\alpha|B_{V_2}(1)|)^{\beta}}{\omega(H_1, V_1)}.
\end{aligned}$$

In particular, this means that if $|B_{V_2}(1)| \geqslant \frac{(2\omega(H_1,V_1))^{\frac{1}{\beta}}}{\alpha}$, then $\omega(H_2, V_2) \geqslant 2$. We let $k = \lceil \frac{(2\omega(H_1,V_1))^{\frac{1}{\beta}}}{\alpha} \rceil$, recalling that V_1 was fixed.

For an arbitrary finite $V_2 \subset H_2$ such that $\langle V_2 \rangle = H_2$, we let $\overline{V_2} = B_{V_2}(k)$. Note that this is also a finite generating set of H_2, with the additional property that $|B_{\overline{V_2}}(1)| = |B_{V_2}(k)| \geqslant k$. We therefore have that $\omega(H_2, \overline{V_2}) \geqslant 2$, and so

$$2 \leqslant \omega(H_2, \overline{V_2}) = \lim_{n\to\infty} |B_{\overline{V_2}}(n)|^{\frac{1}{n}} = \lim_{n\to\infty} |B_{V_2}(nk)|^{\frac{1}{n}} = \omega(H_2, V_2)^k,$$

so $\omega(H_2, V_2) \geqslant 2^{\frac{1}{k}}$. Recalling that k was fixed, we therefore have that $\omega(H_2) \geqslant 2^{\frac{1}{k}} > 1$, so H_2 has uniform exponential growth. This is a contradiction, so there must exist some finite generating set U of $H_1 \times H_2$ such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$. $\square$

The key to the above proof was being able to say something about the growth of every ball in H_2. This will usually not be possible using a subgroup of such a direct product, as it may not be able to give us sufficient information about H_2 to determine whether it has uniform exponential growth or not. On the other hand, if we instead consider a direct product where one of the factors does not have exponential growth, a similar idea can still work.

For a direct product $H_1 \times H_2$, consider the projection map $\varphi : H_1 \times H_2 \to H_1$. For $U \subset H_1 \times H_2$, we can also consider the restriction $\varphi|_{\langle U \rangle}$. This map is a group homomorphism. In Example 3.6.4, this homomorphism is injective, allowing us to see $\langle U \rangle$ as a subgroup of H_1, and effectively allowing us to ignore the direct product.

More generally, the fact that is it a homomorphism means that such a projection will always be k-to-1, where k is the cardinality of $\ker(\varphi|_{\langle U \rangle})$. So long as this cardinality is finite, it turns out that uniform product set growth for $\langle U \rangle$ can be inherited from its projection to H_1.

Lemma 3.6.6. *Let H_1 and H_2 be groups, and let $U \subset H_1 \times H_2$ be finite. Suppose that for some $\alpha, \beta > 0$ the projection U_1 of U to H_1 satisfies $|U_1^n| \geqslant (\alpha|U_1|)^{\beta n}$ for*

every $n \in \mathbb{N}$, and that the projective map from $\langle U \rangle$ to H_1 is k-to-1 for some $k \in \mathbb{N}$. Then $|U^n| \geqslant (\frac{\alpha}{k}|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Proof. As the projective map is k-to-1 we have that $|U| = k|U_1|$, so combined with the fact that $|U^n| \geqslant |U_1^n|$, we obtain our inequality. □

One consequence of this is that if one of the factors in a direct product is finite, then the growth of any subgroup is entirely determined by its growth in the projection to the other factors. This is why the assumption of H_2 being infinite was necessary in Proposition 3.6.5.

Lemma 3.6.7. *Let H_1 and H_2 be groups, where H_2 is finite. Let $G \leqslant H_1 \times H_2$ be finitely generated. Then G has uniform product set growth if and only if its projection to H_1 has uniform product set growth.*

Proof. Let G_1 be the projection of G to H_1. The projective map is at most $|H_2|$-to-1, so if G_1 has uniform product set growth then by Lemma 3.6.6 so does G.

Now suppose that G_1 does not have uniform product set growth, so given $\alpha, \beta > 0$ we can pick a finite generating set U_1 of G_1 such that $|U_1^n| < (\alpha|U_1|)^{\beta n}$ for some $n \in \mathbb{N}$. Let $U = \{(u, s) \in G : u \in U_1\}$, and note that $\langle U \rangle \leqslant G$.

We claim that U is a finite generating set of G. Let $(a, b) \in G$, and we proceed by induction on the length of a with respect to U_1. When a has length one then it is trivially true that $(a, b) \in B_U(1)$. Now suppose that for any $(a, b) \in G$ where a has length $k - 1$ we have that $(a, b) \in B_U(k - 1)$, and consider $(a, b) \in G$ where a has length k.

Write $a = u_1 \cdots u_k$ where $u_1^{\pm 1}, \ldots, u_k^{\pm 1} \in U_1$, and choose $(u_k, s_k) \in B_U(1)$. Then $(au_k^{-1}, bs_k^{-1}) \in G$ and au_k^{-1} has length $k-1$, so $(au_k^{-1}, bs_k^{-1}) \in B_U(k-1)$, and therefore $(a, b) = (au_k^{-1}, bs_k^{-1})(u_k, s_k) \in B_U(k)$. This means that $G \leqslant \langle U \rangle$, so we have equality.

We know that there exists $n \in \mathbb{N}$ such that $|U_1^n| < (\alpha|U_1|)^{\beta n}$. We can also see that $|U_1^n| \leqslant |U^n| \leqslant |H_2||U_1^n|$. Hence

$$|U^n| \leqslant |H_2||U_1^n| < |H_2|(\alpha|U_1|)^{\beta n} \leqslant (\alpha|H_2|^{\frac{1}{\beta n}}|U|)^{\beta n} \leqslant (\alpha|H_2|^{\frac{1}{\beta}}|U|)^{\beta n}.$$

So given any $\alpha', \beta > 0$, we pick $\alpha = \frac{\alpha'}{|H_2|^{1/\beta}}$. We can then choose $U_1 \subset G_1$ and $n \in \mathbb{N}$ as above, and we will have that $|U^n| < (\alpha'|U|)^{\beta n}$. Therefore G does not have uniform product set growth. □

Remark 3.6.8. Although G_1 having uniform product set growth implies that G has uniform product set growth, note that the constant α changes depending on the size of the finite factor.

On the other hand, if the cardinality of the projection to the other factors is infinite, then this can form an obstruction to uniform product set growth.

Proposition 3.6.9. *Let H_1 and H_2 be groups, where H_2 does not have exponential growth. Let $G \leqslant H_1 \times H_2$ be finitely generated, such that the projective map from G to H_1 is infinite-to-1. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. Fix constants $\alpha, \beta > 0$. Suppose that for every finite generating set U of G, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$. We will show that this gives a contradiction to H_2 not having exponential growth.

Let G_{H_1} and G_{H_2} be the projections of G to H_1 and H_2 respectively. Note that as the projection of G to G_{H_1} is infinite-to-1, for every $g \in G_{H_1}$ there exists infinitely many $h \in H_2$ such that $(g, h) \in G$.

We fix a finite $V \subset G$ such that $\langle V \rangle = G$. Let V_{H_1} and V_{H_2} be its projections to H_1 and H_2 respectively. As V_{H_2} is finite we can write it as $V_{H_2} = \{v_1, \dots, v_m\}$. Pick $g \in V_{H_1}$, then for some $k \in \mathbb{N}$ choose $\{\overline{v}_1, \dots, \overline{v}_k\} \subset H_2$ such that $\{v_1, \dots, v_m\} \cap \{\overline{v}_1, \dots, \overline{v}_k\} = \emptyset$ and $(g, \overline{v}_i) \in G$ for every $i \in \{1, \dots, k\}$. Let $\overline{V} = V \cup \{(g, \overline{v}_i) : i \in \{1, \dots, k\}\}$. Then $\langle \overline{V} \rangle = \langle V \rangle = G$, where $\overline{V}_{H_1} = V_{H_1}$ and $|\overline{V}_{H_2}| \geqslant k$. Note that the choice of k here was arbitrary.

As $\overline{V}$ generates G, by our assumption $|\overline{V}^n| \geqslant (\alpha|\overline{V}|)^{\beta n}$ for every $n \in \mathbb{N}$. Note that

$$|\overline{V}^n_{H_2}| \leqslant |\overline{V}^n| \leqslant |\overline{V}^n_{H_1}||\overline{V}^n_{H_2}| = |V^n_{H_1}||\overline{V}^n_{H_2}| \leqslant |B_{V_{H_1}}(n)||B_{\overline{V}_{H_2}}(n)|.$$

This means that for every $n \in \mathbb{N}$ we have that

$$\begin{aligned}
|\overline{V}^n| \geqslant (\alpha|\overline{V}|)^{\beta n} &\implies |B_{V_{H_1}}(n)||B_{\overline{V}_{H_2}}(n)| \geqslant (\alpha|\overline{V}_{H_2}|)^{\beta n} \\
&\implies \lim_{n\to\infty} (|B_{V_{H_1}}(n)||B_{\overline{V}_{H_2}}(n)|)^{\frac{1}{n}} \geqslant (\alpha|\overline{V}_{H_2}|)^{\beta} \\
&\implies \omega(G_{H_1}, V_{H_2})\omega(G_{H_1}, \overline{V}_{H_2}) \geqslant (\alpha|\overline{V}_{H_2}|)^{\beta} \\
&\implies \omega(G_{H_2}, \overline{V}_{H_2}) \geqslant \frac{(\alpha|\overline{V}_{H_2}|)^{\beta}}{\omega(G_{H_1}, V_{H_1})}.
\end{aligned}$$

In particular, if we picked $\overline{V}$ such that $|\overline{V}_{H_2}| \geqslant \frac{(2\omega(G_{H_1}, V_{H_1}))^{\frac{1}{\beta}}}{\alpha}$ and $\overline{V}_{H_1} = V_{H_1}$, then we would get that $\omega(G_{H_2}, \overline{V}_{H_2}) \geqslant 2$. We have already shown that such a choice of $\overline{V}$ is possible, so this tells us that G_{H_2} has exponential growth. This is a contradiction, as G_{H_2} is a subgroup of H_2, which does not have exponential growth. We conclude that there must exist some finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$. □

In Chapter 5 we will want to use Lemma 3.6.6 and Proposition 3.6.9 to help answer Question 5 for certain groups. Recall that to fully answer this question we would like to obtain a dichotomy of subgroups that either do or do not have uniform product set growth, for some common $\alpha, \beta > 0$. One potential issue with using the above results to achieve this is that, as noted in Remark 3.6.8, the constants we get from Lemma 3.6.6 are dependent on the cardinality of the projection. Fortunately for us, in the case that H_2 has a a finite index torsion-free subgroup, we find that the cardinality of this projection is bounded by the index of this subgroup. This will be useful when the overall group we are considering has a finite index torsion-free subgroup, as then every subgroup will have a finite index torsion-free subgroup, with the same index.

Lemma 3.6.10. *Let H_1 and H_2 be groups, where H_2 has a torsion-free subgroup H_2' of index $d \in \mathbb{N}$. Let $G \subset H_1 \times H_2$ be finitely generated, such that the projective map from G to H_1 is k-to-1. Then either $k \leqslant d$ or k is infinite.*

Proof. Let G_{H_1} and G_{H_2} be the projections of G to H_1 and H_2 respectively. Suppose that the projective map is k-to-1 for some k such that $k > d$, or is infinite. We will show that it must be infinite.

The kernel of the projective map has cardinality k, so by the pigeonhole principle there must exist $a, b \in G_{H_2}$ such that $a \neq b$, $(e, a), (e, b) \in G$, and $a, b \in hH_2' \subset H_2$. That is, they are in the same coset of H_2' in H_2.

This means that there exist $n_a, n_b \in H_2'$ such that $n_a \neq n_b$ and $a = hn_a$, $b = hn_b$, so $a^{-1}b = n_a^{-1}n_b \neq e$. Therefore $(e, n_a^{-1}n_b) \in G$, and has infinite order as H_2' is torsion-free. In particular this is in the kernel of the projective map, so the kernel must have infinite cardinality. □

In the case that H_2 itself is torsion-free, this statement can be simplified slightly.

Corollary 3.6.11. *Let H_1 and H_2 be groups, where H_2 is torsion-free. Let $G \subset H_1 \times H_2$ be finitely generated. Then the projective map from G to H_1 is either injective or infinite-to-1.*

We are therefore also able to get a simpler version of Proposition 3.6.9 in this case.

Corollary 3.6.12. *Let H_1, H_2 be groups where H_2 is torsion-free and does not have exponential growth. Let $G \leqslant H_1 \times H_2$ be such that G is finitely generated, and the projection of G to H_1 is not injective. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

We therefore re-obtain Corollary 3.6.3 as a particular case of Corollary 3.6.12.

3.6.3 Direct products of groups with uniform product set growth

In the previous subsection we dealt with the case where we have a subgroup of a direct product where one of the factors does not have exponential growth. We will now consider a very different case, which is a subgroup of a direct product where all of the factors have uniform product set growth. We will show that the growth of the factors passes to the growth of the subgroup, so long as the number of factors is bounded. The first thing to note here is that if we have a subset of a direct product, we can always relate the size of the original set to the size of one of its projections.

Lemma 3.6.13. *Let $G = G_1 \times \cdots \times G_m$ be a group, and let $U \subset G$ be finite, with U_i the projection of U to the factor G_i. Then* $\max\{|U_1|, \ldots, |U_m|\} \geqslant |U|^{\frac{1}{m}}$.

Proof. We proceed by induction, noting that the base case where $m = 1$ is trivial. We now assume that it is true for $m = k$, and consider the case where $m = k + 1$.

Let $u, v \in U$, and let $u_i, v_i \in U_i$ be their projections to the factor G_i. Let $\sim$ be an equivalence relation on U such that $u \sim v$ if and only if $u_{k+1} = v_{k+1}$. Let p be the number of such equivalence classes, and note that $p = |U_{k+1}|$. There must be an equivalence class of size at least $\frac{|U|}{p}$, we call this class V.

Let V' be the projection of V to $G_1 \times \cdots \times G_k$, and note that $|V'| = |V|$ and $V'_i = V_i$ for every $i \in \{1, \ldots, k\}$. By assumption we therefore have that $\max\{|V_1|, \ldots, |V_k|\} = \max\{|V'_1|, \ldots, |V'_k|\} \geqslant |V'|^{\frac{1}{k}} = |V|^{\frac{1}{k}}$.

We now observe that

$$\begin{aligned}
\max\{|U_1|, \ldots, |U_{k+1}|\} &= \max\{\max\{|U_1|, \ldots, |U_k|\}, |U_{k+1}|\} \\
&\geqslant \max\{\max\{|V_1|, \ldots, |V_k|\}, |U_{k+1}|\} \\
&\geqslant \max\{|V|^{\frac{1}{k}}, |U_{k+1}|\} \\
&\geqslant \max\{(|U|/p)^{\frac{1}{k}}, p\} \\
&\geqslant |U|^{\frac{1}{k+1}},
\end{aligned}$$

which concludes our induction. □

This almost immediately gives us that if all of the projections of a subset U of a direct product satisfy the uniform product set growth inequality, then so will U, with a change of constants that depends on the number of factors in the direct product.

Corollary 3.6.14. *Let $G = G_1 \times \cdots \times G_m$ be a group, and let $U \subset G$ be finite, with U_i the projection of U to the factor G_i. Suppose there exist $\alpha, \beta > 0$ such that $|U_i^n| \geqslant (\alpha|U_i|)^{\beta n}$ for every $n \in \mathbb{N}$ and $i \in \{1, \ldots, m\}$. Then $|U^n| \geqslant (\alpha^m|U|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$.*

Proof. By Lemma 3.6.13 we choose U_i such that $|U_i| \geqslant |U|^{\frac{1}{m}}$. We can then conclude that $|U^n| \geqslant |U_i^n| \geqslant (\alpha|U_i|)^{\beta n} \geqslant (\alpha|U|^{\frac{1}{m}})^{\beta n} = (\alpha^m|U|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$. □

The above result tell us that if we know about the growth of the projections of a specific subset, then we can say something about the growth of the whole subset. This naturally means that if we can say something general about the growth of subsets of the factor groups, then we can also say something about the growth of subsets of the direct product.

Definition 3.6.15. We say a group is *contained non-trivially in a direct product* when it is a subgroup of a direct product, and none of its projections to any of the factors groups are trivial.

Corollary 3.6.16. *Let $G = G_1 \times \cdots \times G_m$ be a group, where each G_i is infinite. Suppose there exist $\alpha, \beta > 0$ such that for every $i \in \{1, \ldots, m\}$, and every finite $V_i \subset G_i$ such that $\langle V_i \rangle$ has (uniform) exponential growth, we have that $|V_i^n| \geqslant (\alpha|V_i|)^{\beta n}$ for every $n \in \mathbb{N}$. Then for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ does not have (uniform) exponential growth.

(2) $\langle U \rangle$ is contained non-trivially in a direct product of the form $H_1 \times H_2$, where H_2 does not have (uniform) exponential growth.

(3) $|U^n| \geqslant (\alpha^m|U|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$.

Proof. Suppose $U \subset G$ is such that $\langle U \rangle$ has (uniform) exponential growth, and is not contained non-trivially in a direct product $H_1 \times H_2$ where H_2 does not have (uniform) exponential growth. Then every projection U_i to G_i is either trivial, or $\langle U_i \rangle$ has (uniform) exponential growth. Suppose that k of the projections are non-trivial, and note that $k \geqslant 1$. Then we can regard U as a subset of $G_{i_1} \times \cdots \times G_{i_k}$, where each projection is non-trivial, and hence each projection generates a subgroup with (uniform) exponential growth. Therefore $|U^n| \geqslant (\alpha^k|U|)^{\frac{\beta n}{k}}$ for every $n \in \mathbb{N}$ by Corollary 3.6.14, so as $k \leqslant m$ we get that $|U^n| \geqslant (\alpha^m|U|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$. □

Remark 3.6.17. The case where any G_i is finite can be reduced to the above case by Lemma 3.6.7.

Theorem 3.4.9 tells us that the assumptions of Corollary 3.6.16 are satisfied if the G_i are hyperbolic groups, however in this case we also have control over the size of finite subgroups, so we do not need to rule them out when looking for uniform product set growth.

Corollary 3.6.18. *Let $G = G_1 \times \cdots \times G_m$ be a group such that each G_i is hyperbolic. There exists a constant $\alpha > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\langle U \rangle$ is contained non-trivially in a direct product of the form $H_1 \times H_2$, where H_2 is virtually $\mathbb{Z}$.

(3) $|U^n| \geqslant (\alpha|U|)^{\frac{n}{2m}}$ for every $n \in \mathbb{N}$.

Proof. The statement of Theorem 3.4.9 says that for each hyperbolic group G_i there exists $\alpha_i > 0$ such that for every finite $V_i \subset G_i$ where $\langle V_i \rangle$ is not virtually $\mathbb{Z}$, we have that $|V_i^n| \geqslant (\alpha_i|V_i|)^{\frac{n}{2}}$ for every $n \in \mathbb{N}$. Suppose $U \subset G$ is such that $\langle U \rangle$ is not virtually $\mathbb{Z}$, and is not contained non-trivially in a direct product $H_1 \times H_2$ where H_2 is virtually $\mathbb{Z}$. Let $\alpha = \min\{\alpha_1, \ldots, \alpha_m\}^m$, then the conclusion follows from Corollary 3.6.14. □

This also suggests a possibly more general technique for finding uniform product set growth in groups that contain direct products with a bounded number of factors. If we can prove something about subsets that are not contained non-trivially in direct products, then we can carry this over to subsets that are, so long as all of the projections have the correct properties.

It is important, however, to remark here that Corollary 3.6.16 and Corollary 3.6.18 do not give us a dichotomy of subgroups, as demonstrated by Example 3.6.4. We might be tempted to try to use the work in the previous subsection to get such a dichotomy, however in the general case this is not possible, as using finite-to-1 projections to get down to the factors with uniform product set growth will only work if we have some control over the size of the projections.

On the other hand, we know that we do have some control over these projections in the case of virtually torsion-free groups by Lemma 3.6.10, and so in this case we can get a dichotomy.

Proposition 3.6.19. *Let $G = G_1 \times \cdots \times G_m$ be a direct product of virtually torsion-free groups. Suppose there exist $\alpha, \beta > 0$ such that for every $i \in \{1, \ldots, m\}$, and every*

finite $V_i \subset G_i$ such that $\langle V_i \rangle$ has exponential growth, we have that $|V_i^n| \geqslant (\alpha |V_i|)^{\beta n}$ for every $n \in \mathbb{N}$. Then there exist $\alpha', \beta' > 0$ such that for every finite $U \subset G$ at least one of the following must hold:

(1) $\langle U \rangle$ is contained in a direct product of the form $H_1 \times H_2$, where H_2 does not have exponential growth and the projection of $\langle U \rangle$ to H_1 is infinite-to-1.

(2) $|U^n| \geqslant (\alpha' |U|)^{\beta' n}$ for every $n \in \mathbb{N}$.

Proof. For every G_i, there exists a torsion-free subgroup G_i' of finite index. Let that index be $d_i \in \mathbb{N}$. Then $G_1' \times \cdots \times G_m'$ is a torsion-free subgroup of G with index $d = d_1 \dots d_m$. Hence every subgroup of G has a torsion-free subgroup with index at most d. Note that $d \geqslant d_i$ for every i.

Let $U \subset G$ be finite. Assume that if $\langle U \rangle$ is contained a direct product of the form $H_1 \times H_2$, where H_2 does not have exponential growth, then the projection of $\langle U \rangle$ to H_1 is finite-to-1. By Lemma 3.6.10, this projection would be at most d-to-1. We note that this allows us to rule out the case that $\langle U \rangle$ is infinite and does not have exponential growth, as otherwise we could see $\langle U \rangle$ as a subgroup of $\{e\} \times \langle U \rangle$, where the projection to $\{e\}$ is infinite-to-1.

Suppose that $\langle U \rangle$ is finite. As it has a torsion-free subgroup of index at most d, this means that the trivial subgroup has index at most d in $\langle U \rangle$. Hence $|U| \leqslant |\langle U \rangle| \leqslant d$, so as long as we ensure that $\alpha \leqslant \frac{1}{d}$ we will have that U satisfies $|U^n| \geqslant (\alpha |U|)^{\beta n}$ for every $n \in \mathbb{N}$.

We can therefore assume that $\langle U \rangle$ is infinite. Let U_i be the projection of U to G_i, so $\langle U \rangle \leqslant \langle U_1 \rangle \times \cdots \times \langle U_m \rangle$. We rewrite the U_i as $V_1, \dots, V_k, W_1, \dots, W_{m-k}$, where the $\langle V_i \rangle$ have exponential growth, and the $\langle W_i \rangle$ do not. Note that the set $\{V_1, \dots, V_k\}$ is non-empty, as otherwise $\langle U \rangle$ would not have exponential growth by Lemma 3.1.15. On the other hand, $\{W_1, \dots, W_{m-k}\}$ may be empty.

Let U' be the projection of U to $\langle V_1 \rangle \times \cdots \times \langle V_k \rangle$. The projection of U' to each of these factors is simply V_i, and as each $\langle V_i \rangle$ has exponential growth, the properties of G_i tell us that $|V_i^n| \geqslant (\alpha |V_i|)^{\beta n}$ for every $n \in \mathbb{N}$. Noting that $k \leqslant m$, we have that $|(U')^n| \geqslant (\alpha^m |U'|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$ by Corollary 3.6.14.

We know that $\langle W_1 \rangle \times \cdots \times \langle W_{m-k} \rangle$ has a torsion-free subgroup of index at most d, so the projection of $\langle U \rangle$ to $\langle V_1 \rangle \times \cdots \times \langle V_k \rangle$ is at most d-to-1. It therefore follows that $|U^n| \geqslant (\frac{\alpha^m}{d} |U|)^{\frac{\beta n}{m}}$ for every $n \in \mathbb{N}$ by Lemma 3.6.6.

Combining this with the finite case, we see that if we let $\alpha' = \min\{\frac{\alpha^m}{d}, \frac{1}{d}\}$ and $\beta' = \frac{\beta}{m}$, then any finite $U \subset G$ that satisfies our initial assumption will also satisfy $|U^n| \geqslant (\alpha' |U|)^{\beta' n}$ for every $n \in \mathbb{N}$. $\square$

Remark 3.6.20. This is a dichotomy of subgroups by Proposition 3.6.9. Note that in the collection of groups we rule out, H_2 will always be infinite, as if it were finite then it would not be possible for the projection to H_1 to be infinite-to-1.

The method used in the above proof is very similar to the methods we will use in Chapter 5, where we will be considering right-angled Artin groups and mapping class groups, which are virtually torsion-free and have bounds on the number of factors in any subgroup that is a direct product.

We note here that, as with Corollary 3.6.16, the above result has a corollary for products of hyperbolic groups. We again use the fact that Theorem 3.4.9 tells us that hyperbolic groups satisfy the hypotheses of Proposition 3.6.19, although this time we need to add the assumption that these hyperbolic groups are virtually torsion-free.

Corollary 3.6.21. *Let $G = G_1 \times \cdots \times G_m$ be a direct product of virtually torsion-free hyperbolic groups. There exist $\alpha, \beta > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is contained a direct product of the form $H_1 \times H_2$, where H_2 is virtually $\mathbb{Z}$ and the projection of $\langle U \rangle$ to H_1 is infinite-to-1.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Remark 3.6.22. Whether all hyperbolic groups are virtually torsion-free or not is an open question, and is equivalent to the problem of whether all hyperbolic groups are residually finite [53].

3.6.4 Passing to finite index supergroups

We now move away from direct products, and consider another way that uniform product set growth can be inherited. It is well known that uniform exponential growth passes to finite index supergroups [80, Proposition 3.3]. This can be shown by checking that for any finite generating set of our supergroup, there exists a finite generating set of our finite index subgroup where each generator has bounded length. We can adapt this method to show that uniform product set growth also passes to finite index supergroups, at least in the case of symmetric sets.

Lemma 3.6.23. [64, Proposition 2.3] *Let G be a finitely generated group, and let U be a finite generating set of G. Let $H \leqslant G$ such that $[G : H] = d$. There exists a set of representatives of the right cosets of H such that each representative has length at most $d - 1$ in the word metric induced by U.*

Proposition 3.6.24. *Let G be a finitely generated group, and let U be a finite generating set of G. Let $H \leqslant G$ such that $[G : H] = d$. There exists a finite generating set V of H such that $|V| \geqslant \frac{1}{d}|U|$ and $V \subset B_U(2d-1)$.*

Proof. Let $\{g_1, \ldots, g_d\} \subset G$ be a set of representatives of the right cosets of H such that each representative has length at most $d-1$ in the word metric induced by U, which we know exists by Lemma 3.6.23. Proposition 2.1 in [64] tells us that the set of elements $g_i u g_j^{-1}$ such that $u \in U$ or $u^{-1} \in U$, and g_j is the right coset representative of $g_i u$, is a finite generating set for H. We call this set V, and note that for $u, v \in U$ such that $u \neq v$, we have that $g_1 u g_i^{-1} = g_1 v g_j^{-1}$ implies that $i \neq j$. This means that for any $d+1$ elements in U, we can construct at least two different elements in V, and so forth. Therefore $|V| \geqslant \frac{1}{d}|U|$. We also note that each $g_i u g_j^{-1}$ has length at most $2d-1$ in the word metric induced by U, so $V \subset B_U(2d-1)$. □

Proposition 3.6.25. *Let G be a finitely generated group, and let U be a finite symmetric generating set of G. Let $H \leqslant G$ such that $[G : H] = d$. Suppose that there exist $\alpha, \beta > 0$ such that for any finite symmetric generating set V of H, we have that $|V^n| \geqslant (\alpha|V|)^{\beta n}$ for every $n \in \mathbb{N}$. Let $m = 2d-1$. Then*

$$|U^n| \geqslant \left(\frac{\alpha}{2^{\frac{m}{\beta}} d}|U|\right)^{\frac{\beta n}{m}}$$

for every $n \in \mathbb{N}$.

Proof. We know from Proposition 3.6.24 that there exists a finite symmetric generating set V of H such that $|V| \geqslant \frac{1}{d}|U|$ and $V \subset B_U(m)$. Let $n \in \mathbb{N}$. Using Lemma 3.2.6 we know that $|B_U(n)| \leqslant 2|U^n|$, as U is symmetric.

For every $n \in \mathbb{N}$ we have that

$$\begin{aligned}(2|U^n|)^m &\geqslant |B_U(n)|^m \geqslant |B_U(n)^m| = |B_U(m)^n| \\ &\geqslant |V^n| \geqslant (\alpha|V|)^{\beta n} \geqslant \left(\frac{\alpha}{d}|U|\right)^{\beta n}.\end{aligned}$$

Therefore

$$|U^n| \geqslant \frac{1}{2}\left(\frac{\alpha}{d}|U|\right)^{\frac{\beta n}{m}} = \left(\frac{\alpha}{2^{\frac{m}{\beta n}} d}|U|\right)^{\frac{\beta n}{m}}$$

It now only remains to note that $2^{\frac{1}{n}} \leqslant 2$ for every $n \in \mathbb{N}$, and the conclusion follows. □

4

Growth in acylindrically hyperbolic groups

Algebra is but written geometry and geometry is but figured algebra.

– Attributed to Sophie Germain

The focus of this chapter will be to prove some results about product set growth for certain subsets of acylindrically hyperbolic groups. This will provide us with the tools that we will need in Chapter 5 when trying to answer Question 5 for subgroups of mapping class groups.

We will first cover the basics of acylindrical actions on hyperbolic spaces in Section 4.1. Then, in Section 4.2, we will use the uniform version of tree approximation given by Proposition 2.4.1 to get a generalisation of Theorem 3.4.14 for groups acting acylindrically on quasi-trees. As noted in the introduction, it was shown by Balasubramanya that every acylindrically hyperbolic group has an acylindrical action on a quasi-tree. Hence this generalisation to quasi-trees can be seen as being complementary to Delzant and Steenbock's statement for groups acting acylindrically on hyperbolic spaces, as stated in Theorem 3.4.12.

We will also use Section 4.2 to give a concrete method by which we can use this generalisation to prove uniform product set growth for certain subgroups of acylindrically hyperbolic groups, namely by quickly finding loxodromic elements in its action on a quasi-tree. In Section 4.3 we will show that this method can also be used when the acylindrical action is on a hyperbolic space, which is a more practical setting for many groups.

4.1 Acylindrical actions and loxodromic elements

We begin by recalling the definition of an acylindrical action.

Definition 4.1.1. An isometric action by a group G on a metric space (X, d) is *acylindrical* if for every $R \geqslant 0$, there exist $N > 0$ and $L > 0$ such that for every $x, y \in X$ with $d(x, y) \geqslant L$ we have that

$$|\{g \in G : d(x, gx) \leqslant R \text{ and } d(y, gy) \leqslant R\}| \leqslant N.$$

In other words, when we let X be a geodesic space, we can see an acylindrical action as one where only a bounded number of group elements can almost preserve any long cylinder in X. This is a generalisation of how a finitely generated group acts on its Cayley graph, where proper discontinuity means that only a bounded number of group elements can almost preserve a point in the graph. In an acylindrical action this is not the case, in particular a point could have a stabiliser of infinite cardinality. We instead require a restriction on the size of the stabiliser of any two points that are sufficiently far apart.

When the space being acted on is hyperbolic, we can also use the following alternative terminology from [29].

Definition 4.1.2. We say that an isometric action of a group G on a δ-hyperbolic geodesic metric space is (κ_0, N_0)-*acylindrical* for $\kappa_0 \geqslant \delta$ and $N_0 \geqslant 1$ if for every $x, y \in X$ with $d(x, y) \geqslant \kappa_0$ we have that

$$|\{g \in G : d(x, gx) \leqslant 100\delta \text{ and } d(y, gy) \leqslant 100\delta\}| \leqslant N_0.$$

Remark 4.1.3. If an action is acylindrical on a δ-hyperbolic geodesic metric space in the sense of Definition 4.1.1, then by taking $R = 100\delta$ we have that the action is (L, N)-acylindrical. In fact, by Proposition 5.31 in [29], if we let $\delta > 0$ then the two notions are equivalent.

Convention. Every hyperbolic space we work with in the remaining chapters will be assumed to be geodesic.

It turns out that every group admits an acylindrical action on a hyperbolic space, simply by considering the trivial action on a point. To have a meaningful class of groups to consider we therefore need another condition on the action.

Definition 4.1.4. An acylindrical action by a group G on a hyperbolic geodesic metric space X is *non-elementary* if G is not virtually cyclic and all orbits are unbounded.

Definition 4.1.5. [73] A group is *acylindrically hyperbolic* if it admits a non-elementary acylindrical action on a hyperbolic geodesic metric space.

Example 4.1.6. The following are all example of acylindrically hyperbolic groups (for more details see the appendix of [73]):

(1) Hyperbolic groups that are not virtually cyclic.

(2) Relatively hyperbolic groups that are not virtually cyclic and have proper peripheral subgroups.

(3) Non-finite mapping class groups of surfaces without boundary [67, 15].

(4) Right-angled Artin groups that are not cyclic, and do not decompose as a direct product of non-trivial groups (a similar result also holds for graph products) [69].

(5) $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ for $n \geqslant 2$ [10, 40].

(6) Groups acting properly on proper CAT(0) spaces, that are not virtually cyclic and contain a rank one element [73, 81].

(7) One-relator groups with at least three generators [69].

One well-known example is that of free products acting on their Bass-Serre tree, which we will use in Chapter 5.

Lemma 4.1.7. *Let G and H be groups. Their free product $G * H$ acts acylindrically on its associated Bass-Serre tree.*

Proof. The Bass-Serre tree for a free product is the graph where the vertices are cosets of G and H, and for every $g \in G * H$ the cosets gG and gH have an edge between them. The action of $G * H$ on this tree is by left multiplication. For more details on this see [78].

The stabiliser of an edge in this tree is therefore of the form $gGg^{-1} \cap gHg^{-1} = g(G \cap H)g^{-1}$ for some $g \in G * H$, so we can conclude that these stabilisers are trivial. By Lemma 5.2 in [69] this implies that the action of $G * H$ on its Bass-Serre tree is acylindrical. □

Remark 4.1.8. This action has unbounded orbits so long as neither G nor H is trivial, and if additionally at least one of G or H has order at least three, then $G * H$ is not virtually cyclic.

Some non-examples of acylindrical hyperbolicity are provided by the following result.

Theorem 4.1.9. [73] *Suppose G is an acylindrically hyperbolic group. Then G is non-amenable, and if $G = G_1 \times G_2$ then exactly one factor is a finite group.*

In Section 4.2 we will be interested in the acylindrically hyperbolic groups that have an acylindrical action on a quasi-tree. By a result of Balasubramanya, this is in fact the entire class of acylindrically hyperbolic groups.

Theorem 4.1.10. [6] *Every acylindrically hyperbolic group admits a non-elementary acylindrical action on a quasi-tree.*

This theorem means that any result for groups acting acylindrically on quasi-trees can automatically be applied to every acylindrically hyperbolic group. This helps motivate Section 4.2.

Later in this chapter we will show that we can say something about the growth of a subset of an acylindrically hyperbolic group if that subset contains a certain type of element. By type of element we mean that if a group has an acylindrical action on a hyperbolic space, we can classify the group elements by how they behave in this action.

Definition 4.1.11. Let G be a group acting by isometries on a metric space X. For $g \in G$, the *stable translation length* of g is

$$\tau(g) = \lim_{n\to\infty} \frac{d(x, g^n x)}{n},$$

where $x \in X$.

Remark 4.1.12. It is clear from the triangle inequality that $\tau(g)$ is not dependent on the choice of $x \in X$.

Definition 4.1.13. Let G be a group acting by isometries on a metric space X. We say that $g \in G$ is an *elliptic* element if some (equivalently any) orbit of g in X has bounded diameter.

Definition 4.1.14. Let G be a group acting by isometries on a metric space X. We say that $g \in G$ is a *loxodromic* element if $\tau(g) > 0$.

These two types of element are naturally mutually exclusive, as for an elliptic element g we have that $\tau(g) = 0$. A result of Bowditch tells us that these two types classify all elements in an acylindrical action on a hyperbolic space.

Proposition 4.1.15. [15] *Let G be a group acting by isometries on a hyperbolic space X. Every element $g \in G$ is either elliptic or loxodromic.*

Another important concept later in this chapter will be that of independent loxodromic elements. To define this, it is easier to use an equivalent definition of loxodromic elements in the case of an action on a hyperbolic space.

Definition 4.1.16. Let G be a group acting by isometries on a metric space X. We say that $g \in G$ is a *loxodromic* element if it fixes exactly two points on the geodesic boundary $\partial_\infty X$. We call these points $\mathrm{Fix}(g)$.

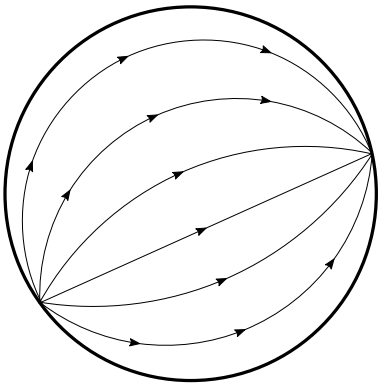


Figure 4.1: A loxodromic element acting on the hyperbolic plane

Remark 4.1.17. Another alternative definition is given by Definition 2.3.20, in the more general case of a quasi-action.

We can now define what it means for two loxodromic elements to be independent.

Definition 4.1.18. Two loxodromic elements g and h are *independent* if $\mathrm{Fix}(g) \cap \mathrm{Fix}(h) = \emptyset$.

We finish this section with the following notation.

Definition 4.1.19. Let G be a group acting acylindrically on a hyperbolic space X, and let g be a loxodromic element in this action. Denote by $E(g)$ the maximal virtually cyclic subgroup containing g.

Remark 4.1.20. This maximal virtually cyclic subgroup exists by Lemma 6.5 in [29], and is exactly the stabiliser of $\mathrm{Fix}(g)$ in G.

4.2 Growth from actions on quasi-trees

In this section we will be interested in a theorem of Delzant and Steenbock, previously stated as Theorem 3.4.12, where they give a lower bound on the growth of U^n for certain subsets of acylindrically hyperbolic groups [31]. As a direct result of Proposition 2.4.1 we will be able to improve this theorem in the case of groups acting acylindrically on quasi-trees. By Theorem 4.1.10 this improvement will apply to every acylindrically hyperbolic group, as every such group admits an acylindrical action on a quasi-tree.

Before we can state the result that we wish to improve, we must first give a few definitions.

Definition 4.2.1. [31] Let U be a finite set of isometries of a δ-hyperbolic geodesic metric space (X, d), with $\delta > 0$. The *normalised ℓ^1-energy* of U is

$$E(U) = \inf_{x \in X} \frac{1}{|U|} \sum_{u \in U} d(x, ux).$$

We then fix a base point $x_0 \in X$ such that it minimises $E(U)$ up to δ, so

$$\frac{1}{|U|} \sum_{u \in U} d(x_0, ux_0) \leqslant E(U) + \delta.$$

The *displacement* of U is then defined to be

$$\lambda_0(U) = \max_{u \in U} d(x_0, ux_0).$$

Remark 4.2.2. Another version of displacement which is sometimes used is to instead take the infimum of $\max_{u \in U} d(x, ux)$ over any $x \in X$, which is clearly a lower bound for $\lambda_0(U)$. For a comparison of such quantities see [16].

The result that Delzant and Steenbock obtained for groups acting acylindrically on hyperbolic spaces is the following.

Theorem 4.2.3. [31] *Let G be a group acting (κ_0, N_0)-acylindrically on a δ-hyperbolic geodesic metric space. Assume $\kappa_0 \geqslant \delta$ and $N_0 \geqslant 1$. There exist constants $\alpha = \alpha(\delta, \kappa_0, N_0) > 0$ and $K = K(\kappa_0) > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K \log_2(2|U|)$.

(3) $|U^n| \geqslant \left(\frac{\alpha}{\log_2^6(2|U|)}|U|\right)^{\lfloor \frac{n+1}{2} \rfloor}$ *for every* $n \in \mathbb{N}$.

Remark 4.2.4. We note here that the action in Theorem 4.2.3 is not assumed to be non-elementary, and as stated before every group G acts acylindrically on some hyperbolic space. However, if the action in question is elementary, then any finite subset U of G must satisfy that $\langle U \rangle$ is virtually $\mathbb{Z}$, or $\lambda_0(U) < K \log_2(2|U|)$, or be such that $|U| \leqslant \frac{1}{\alpha}$, by Proposition 4.4 in [31].

We would like to re-state this theorem for groups acting acylindrically on quasi-trees, but with the logarithm terms removed. To do this we will first check that the step in Delzant and Steenbock's proof in which the logarithm terms are introduced can be improved in the case of quasi-trees, and then combine this with the rest of their results. We will then consider the theorem we obtain in the context of Question 5, which means that we would like to understand which subsets have large enough displacement that our theorem will allow us to say something about their product set growth.

4.2.1 The reduction lemma for quasi-trees

The logarithm terms in Theorem 4.2.3 are introduced in a key step that Delzant and Steenbock call the reduction lemma, which we state below.

Notation 4.2.5. By $(Ux_0, Vx_0)_{x_0} \leqslant R$ we mean that $(ux_0, vx_0)_{x_0} \leqslant R$ for every $u \in U, v \in V$, as in [31].

The constant κ_0 used below is chosen such that $\kappa_0 \geqslant \delta$, and will generally be taken from the acylindrical action on the space, as in the statement of Theorem 4.2.3.

Lemma 4.2.6 (Reduction lemma)**.** [31] *Let U be a finite set of isometries of a δ-hyperbolic geodesic metric space X, with $\delta > 0$. If at most $\frac{1}{4}$ of the isometries $u \in U$ have displacement $d(x_0, ux_0) \leqslant 10^{10}\kappa_0 \log_2(2|U|)$, then there exist $U_0, U_1 \subset U$ with cardinalities at least $\frac{1}{100}|U|$ such that*

$$(U_0^{-1}x_0, U_1x_0)_{x_0} \leqslant 1000 \log_2(2|U|)\delta \text{ and } (U_0x_0, U_1^{-1}x_0)_{x_0} \leqslant 1000 \log_2(2|U|)\delta.$$

In addition, for all $u_0 \in U_0$ and $u_1 \in U_1$ we have that

$$d(x_0, u_0x_0) \geqslant 10^{10}\kappa_0 \log_2(2|U|) \text{ and } d(x_0, u_1x_0) \geqslant 10^{10}\kappa_0 \log_2(2|U|).$$

Figure 4.2 gives a representation of how such sets U_0 and U_1 may look when applied to x_0. Every element displaces x_0 by a large distance, while U_0^{-1} and U_1

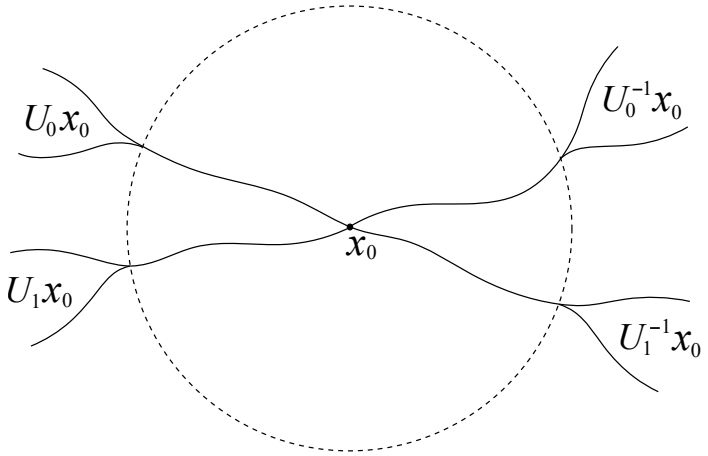


Figure 4.2: The reduction lemma

(respectively U_0 and U_1^{-1}) translate x_0 in opposite directions in the space. Once such sets are found, Delzant and Steenbock estimate $|U^{2n+1}|$ by instead considering $|(U_0U_1)^nU_0|$, and showing that a lower bound on this can be given in terms of $|U_0|^n$. As $|U_0| \geqslant \frac{1}{100}|U|$, a lower bound can then be given in terms of $|U|^n$.

The logarithm terms in the reduction lemma are introduced because the proof makes use of Gromov's Tree Approximation Lemma. In their paper Delzant and Steenbock remark that if the space being acted on is a quasi-tree, then the logarithm terms in Theorem 4.2.3 should disappear (see [31, Remark 1.18]). We are able to check that this is indeed the case. We include the proofs here for completeness, however the method and ideas are the same as those used in [31], with the substitution of Proposition 2.4.1 in place of Gromov's Tree Approximation Lemma.

In everything that follows, we will assume that x_0 has been chosen to minimise $E(U)$ for our given set of isometries U, as in Definition 4.2.1. For the purposes of the reduction lemma the constant κ_0 only needs to satisfy $\kappa_0 \geqslant \Delta$, where Δ is the hyperbolicity or bottleneck constant for the space, however when this lemma is applied to an acylindrical action it will be taken to be the acylindricity constant from Definition 4.1.2.

Remark 4.2.7. To reduce the number of constants, we recall from Remark 2.1.32 that a quasi-tree with bottleneck constant $\Delta \geqslant 0$ is Δ-hyperbolic, so we will use this Δ for both.

We will first consider subsets of U of the form

$$U_{y,z} = \{u \in U : d(x_0, ux_0) \geqslant 10^{10}\kappa_0,\ (x_0, ux_0)_y \leqslant \Delta,\ \text{and}\ (x_0, u^{-1}x_0)_z \leqslant \Delta\},$$

where $y, z \in S(x_0, 1000\Delta)$, the sphere of radius 1000Δ centred at x_0 in X. Note that some of these sets may be empty.

The following result is key to the proof of the reduction lemma.

Lemma 4.2.8 (Minimal energy)**.** [31] *Let U be a finite set of isometries of a Δ-hyperbolic geodesic metric space X, with $\Delta > 0$. Let $y_0 \in S(x_0, 1000\Delta)$, and $Y = B(y_0, 100\Delta) \cap S(x_0, 1000\Delta)$. Then*

$$\left| \bigcup_{y,z \in Y} U_{y,z} \right| \leqslant \frac{2}{3}|U|.$$

The idea of the minimal energy lemma is that if too many of the isometries in U and U^{-1} send x_0 in the same direction, then this would contradict our choice of x_0 as minimising $E(U) = \frac{1}{|U|} \sum_{u \in U} d(x, ux)$. In particular, y_0 would be a better choice.

The statement of the reduction lemma for quasi-trees is as follows.

Lemma 4.2.9 (Reduction lemma)**.** *Let U be a finite set of isometries of a quasi-tree X with bottleneck constant $\Delta > 0$. If at most $\frac{1}{4}$ of the isometries $u \in U$ have displacement $d(x_0, ux_0) \leqslant 10^{10}\kappa_0$, then there exist $U_0, U_1 \subset U$ with cardinalities at least $\frac{1}{100}|U|$ such that*

$$(U_0^{-1}x_0, U_1x_0)_{x_0} \leqslant 1000\Delta \text{ and } (U_0x_0, U_1^{-1}x_0)_{x_0} \leqslant 1000\Delta.$$

In addition, for all $u_0 \in U_0$ and $u_1 \in U_1$ we have that

$$d(x_0, u_0x_0) \geqslant 10^{10}\kappa_0 \text{ and } d(x_0, u_1x_0) \geqslant 10^{10}\kappa_0.$$

To prove the reduction lemma we need two preliminary lemmas, for which we will assume that X and U are as above. As with everything in this subsection, their proofs follow their analogues in [31], with the only difference being the substitution of constants to match the uniform tree approximation lemma.

We first set up our notation. Let $r = 1000\Delta$ and $S = S(x_0, r)$. Consider

$$V = \bigcup_{x \in Ux_0 \cup U^{-1}x_0} [x_0, x].$$

As U is finite, by Proposition 2.4.1 there exists an $\mathbb{R}$-tree (T, d^*) and a map $f : (V, d) \to (T, d^*)$ such that:

(1) For all $x \in Ux_0 \cup U^{-1}x_0$, the restriction of f to the geodesic segment $[x_0, x]$ is an isometry.

(2) For all $y, z \in V$, we have that $d(y,z) - 6\Delta \leqslant d^*(f(y), f(z)) \leqslant d(y,z)$.

In particular, this $\mathbb{R}$-tree T is exactly the end-approximating tree T_X, constructed using the basepoint x_0.

Let $\overline{S}$ be the sphere of radius r centred at $f(x_0)$ in T. As f restricted to any geodesic segment $[x_0, x]$ is an isometry, we have that $\overline{S} = f(S \cap V)$. For any $P, Q \subset \overline{S}$, we let $k = 10^{10}\kappa_0$ and define

$$\overline{U_{P,Q}} = \{u \in U : d(x_0, ux_0) \geqslant k, \text{ and there exists } p \in P, q \in Q \text{ such that } (f(x_0), f(ux_0))_p = 0 \text{ and } (f(x_0), f(u^{-1}x_0))_q = 0\}.$$

This is equivalent to saying that $u \in U$ with displacement $d(x_0, ux_0) \geqslant k$ is in $\overline{U_{P,Q}}$ if and only if $[f(x_0), f(ux_0)]$ intersects $\overline{S}$ at P, and $[f(x_0), f(u^{-1}x_0)]$ intersects $\overline{S}$ at Q.

The proof of Lemma 4.2.10 follows the proof of Lemma 6.9 in [31].

Lemma 4.2.10. *Let $P \subset \overline{S}$ and $Q \subset \overline{S} \backslash P$. Then*

$$(\overline{U_{\diamond,Q}}^{-1} x_0, \overline{U_{P,\star}} x_0)_{x_0} \leqslant r,$$

where $\diamond$ and $\star$ are used as placeholders for any subsets of $\overline{S}$.

Proof. Let $u_0 \in \overline{U_{P,\star}}$, and $u_1 \in \overline{U_{\diamond,Q}}$. This means that $[f(x_0), f(u_0 x_0)]$ intersects $\overline{S}$ at P, and $[f(x_0), f(u_1^{-1}x_0)]$ intersects $\overline{S}$ at Q. As $P \cap Q = \emptyset$, and we are in a tree, this means that

$$(f(u_0 x_0), f(u_1^{-1} x_0))^*_{f(x_0)} \leqslant r.$$

We know from the proof of Lemma 2.2.4 that

$$(u_0 x_0, u_1^{-1} x_0)_{x_0} \leqslant (f(u_0 x_0), f(u_1^{-1} x_0))^*_{f(x_0)},$$

and so

$$(u_0 x_0, u_1^{-1} x_0)_{x_0} \leqslant r.$$

□

Sets of the form $\overline{U_{P,\star}}$ and $\overline{U_{\diamond,Q}}$ therefore have the required Gromov product, and have the required displacement by definition. If we can find such sets that have cardinalities at least $\frac{1}{100}|U|$, then we will be done.

The proof of Lemma 4.2.11 follows the proof of Lemma 6.11 from [31]. Let $P \subset \overline{S}$, and $Q = \overline{S} \backslash P$. Fix $p \in P$. Let $P' = P \backslash \{p\}$, and $Q' = Q \cup \{p\}$. Note that $|\overline{U_{P',P'}}| \leqslant |\overline{U_{P,P}}|$, and $|\overline{U_{Q',Q'}}| \geqslant |\overline{U_{Q,Q}}|$.

Lemma 4.2.11. *Suppose at most $\frac{1}{4}$ of the isometries $u \in U$ have displacement $d(x_0, ux_0) \leqslant k$, and that we have the following:*

(1) $|\overline{U_{P,Q}}| \leqslant \frac{1}{100}|U|$, $|\overline{U_{Q,P}}| \leqslant \frac{1}{100}|U|$, *and* $|\overline{U_{Q,Q}}| \leqslant \frac{1}{100}|U|$.

(2) $|\overline{U_{P',Q'}}| \leqslant \frac{1}{100}|U|$, $|\overline{U_{Q',P'}}| \leqslant \frac{1}{100}|U|$, *and* $|\overline{U_{Q',Q'}}| > \frac{1}{100}|U|$.

Then $|\overline{U_{P',P'}}| > \frac{1}{100}|U|$.

Proof. Suppose that $|\overline{U_{P',P'}}| \leqslant \frac{1}{100}|U|$. We want to get a lower bound on $|\overline{U_{p,p}}|$, and use this to get a contradiction with the minimal energy lemma.

First note that as $Q = \overline{S}\backslash P$ and T is a tree, the sets $\overline{U_{P,P}}$, $\overline{U_{P,Q}}$, $\overline{U_{Q,P}}$, and $\overline{U_{Q,Q}}$ are disjoint and

$$\{u \in U : d(x_0, ux_0) \geqslant k\} = \overline{U_{P,P}} \sqcup \overline{U_{P,Q}} \sqcup \overline{U_{Q,P}} \sqcup \overline{U_{Q,Q}}.$$

We know that $|\{u \in U : d(x_0, ux_0) \geqslant k\}| \geqslant \frac{75}{100}$, so $|\overline{U_{P,P}}| \geqslant \frac{72}{100}$, and therefore $|\overline{U_{P,P}}\backslash\overline{U_{P',P'}}| \geqslant \frac{71}{100}$. We also have that

$$\begin{aligned} |\overline{U_{P,P}}\backslash\overline{U_{P',P'}}| &= |\overline{U_{P',p}}| + |\overline{U_{p,P'}}| + |\overline{U_{p,p}}| \\ &\leqslant |\overline{U_{P',Q'}}| + |\overline{U_{Q',P'}}| + |\overline{U_{p,p}}| \\ &\leqslant \frac{2}{100}|U| + |\overline{U_{p,p}}|. \end{aligned}$$

Hence $|\overline{U_{p,p}}| \geqslant \frac{69}{100}$.

To get a contradiction with the minimal energy lemma, we want to show that $\overline{U_{p,p}}$ is a subset of $\bigcup_{y,z \in B(y_0, 6\Delta) \cap S} U_{y,z}$ for some $y_0 \in S$.

Consider the set $f^{-1}(p)$. As $p \in \overline{S}$, and f maps geodesics from x_0 isometrically, we have that $f^{-1}(p) \subset S$. Fix $y_0 \in f^{-1}(p)$. For any $z \in f^{-1}(p)$, we get that $d(y_0, z) \leqslant d^*(f(y_0), f(z)) + 6\Delta = 6\Delta$. Hence $f^{-1}(p) \subset B(y_0, 6\Delta) \cap S$.

It therefore suffices to show that $\overline{U_{p,p}}$ is a subset of $\bigcup_{y,z \in f^{-1}(p)} U_{y,z}$. Let $u \in \overline{U_{p,p}}$. The function f maps $[x_0, ux_0]$ isometrically onto $[f(x_0), f(ux_0)]$, and p lies on $[f(x_0), f(ux_0)]$, so there exists $y \in f^{-1}(p)$ such that $y \in [x_0, ux_0]$. This means that $(x_0, ux_0)_y = 0 < \Delta$. Similarly, we can pick $z \in f^{-1}(p)$ such that $(x_0, u^{-1}x_0)_z = 0 < \Delta$. Hence $u \in U_{y,z}$, so

$$\frac{69}{100} \leqslant |\overline{U_{p,p}}| \leqslant \left| \bigcup_{y,z \in f^{-1}(p)} U_{y,z} \right| \leqslant \left| \bigcup_{y,z \in B(y_0, 6\Delta) \cap S} U_{y,z} \right|.$$

This contradicts the minimal energy lemma, so $|\overline{U_{P',P'}}| > \frac{1}{100}|U|$. □

We will now use these results to prove the reduction lemma, following the proof of Lemma 6.4 in [31].

Proof of reduction lemma. We want to find two sets U_0, U_1 of the form $\overline{U_{P,\star}}$ and $\overline{U_{\diamond,Q}}$, with $P \subset \overline{S}$ and $Q \subset \overline{S} \backslash P$, such that $|\overline{U_{P,\star}}| > \frac{1}{100}$ and $|\overline{U_{\diamond,Q}}| > \frac{1}{100}$.

Let $P^{(0)} = \overline{S}$ and $Q^{(0)} = \emptyset$. Note that $\overline{S}$ is finite, so we can write $P^{(0)}$ as $P^{(0)} = \{p_1, \dots, p_N\}$ for some $N \in \mathbb{N}$. We now recursively define $P^{(n)} = P^{(n-1)} \backslash \{p_n\}$, and $Q^{(n)} = Q^{(n-1)} \cup \{p_n\}$.

If for any n we have that $|\overline{U_{P^{(n)},Q^{(n)}}}| > \frac{1}{100}$ or $|\overline{U_{Q^{(n)},P^{(n)}}}| > \frac{1}{100}$, then we can set $U_0 = U_1$ to be this set, and we are done. Suppose otherwise, so $|\overline{U_{P^{(n)},Q^{(n)}}}| \leqslant \frac{1}{100}$ and $|\overline{U_{Q^{(n)},P^{(n)}}}| \leqslant \frac{1}{100}$ for all n. We can see that $|\overline{U_{Q^{(n)},Q^{(n)}}}|$ is an increasing sequence, with $|\overline{U_{Q^{(0)},Q^{(0)}}}| = 0$ and

$$|\overline{U_{Q^{(N)},Q^{(N)}}}| = |\overline{U_{\overline{S},\overline{S}}}| = |\{u \in U : d(x_0, ux_0) \geqslant k\}| \geqslant \frac{75}{100}.$$

Hence there exists n such that $|\overline{U_{Q^{(n-1)},Q^{(n-1)}}}| \leqslant \frac{1}{100}$ and $|\overline{U_{Q^{(n)},Q^{(n)}}}| > \frac{1}{100}$. By Lemma 4.2.11 we also get that $|\overline{U_{P^{(n)},P^{(n)}}}| > \frac{1}{100}$, so we can set $U_0 = \overline{U_{P^{(n)},P^{(n)}}}$ and $U_1 = \overline{U_{Q^{(n)},Q^{(n)}}}$. □

The subsets used in [31] actually take a slightly stronger form, where the displacement of x_0 by one subset is great than that of the other, however as in [31] such sets are easily obtainable as a corollary of the reduction lemma.

Corollary 4.2.12. *Let U be a finite set of isometries of a quasi-tree X with bottleneck constant $\Delta > 0$. If at most $\frac{1}{4}$ of the isometries $u \in U$ have displacement $d(x_0, ux_0) \leqslant 10^{10}\kappa_0$, then there exist $U_0, U_1 \subset U$ with cardinalities at least $\frac{1}{200}|U|$ such that*

$$(U_0^{-1}x_0, U_1x_0)_{x_0} \leqslant 1000\Delta \text{ and } (U_0x_0, U_1^{-1}x_0)_{x_0} \leqslant 1000\Delta.$$

In addition, for all $u_0 \in U_0$ and $u_1 \in U_1$ we have that

$$10^{10}\kappa_0 \leqslant d(x_0, u_0x_0) \leqslant d(x_0, u_1x_0).$$

Proof. Let U_0' and U_1' be the sets chosen in the reduction lemma. Let m_0 be the median of $\{d(x_0, ux_0) : u \in U_0'\}$, and m_1 be the median of $\{d(x_0, ux_0) : u \in U_1'\}$.

If $m_0 \leqslant m_1$, then let $U_0 = \{u \in U_0' : d(x_0, ux_0) \leqslant m_0\}$ and $U_1 = \{u \in U_1' : d(x_0, ux_0) \geqslant m_1\}$. If $m_1 < m_0$ then let $U_0 = \{u \in U_1' : d(x_0, ux_0) \leqslant m_1\}$ and $U_1 = \{u \in U_0' : d(x_0, ux_0) \geqslant m_0\}$. In both cases note that U_0 and U_1 have cardinalities at least half of that of the original sets. □

4.2.2 Growth of sets with large displacement

When combined with the rest of the work in [31], Corollary 4.2.12 gives us the desired result for acylindrical actions on quasi-trees.

Theorem 4.2.13. *Let G be a group acting (κ_0, N_0)-acylindrically on a quasi-tree X with bottleneck constant $\Delta > 0$. Assume $\kappa_0 \geqslant \Delta$ and $N_0 \geqslant 1$. There exist constants $\alpha = \alpha(\Delta, \kappa_0, N_0) > 0$ and $K = K(\kappa_0) > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K$.

(3) $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

In particular, we can take $\alpha = \frac{\Delta^2}{10^{52} N_0^6 \kappa_0^2}$ and $K = 10^{14}\kappa_0$.

Proof. This is a combination of Corollary 4.2.12 from this thesis along with Corollary 5.6 and Proposition 6.18 from [31], using $d = 1$, $b = 10$, and $\delta = \Delta$. □

Remark 4.2.14. This is a generalisation of Delzant and Steenbock's result for groups acting acylindrically on trees, stated in this thesis as Theorem 3.4.14.

We now recall Theorem 4.1.10, which stated that every acylindrically hyperbolic group admits a non-elementary acylindrical action on a quasi-tree. This allows us to apply Theorem 4.2.13 to the entire class of acylindrically hyperbolic groups.

Corollary 4.2.15. *Let G be an acylindrically hyperbolic group. Then there exists a quasi-tree with bottleneck constant $\Delta > 0$ on which G has a non-elementary (κ_0, N_0)-acylindrical action, with and $\kappa_0 \geqslant \Delta$ and $N_0 \geqslant 1$, and constants $\alpha = \alpha(\Delta, \kappa_0, N_0) > 0$ and $K = K(\kappa_0) > 0$ such that for every finite $U \subset G$ at least one of the following must hold:*

(1) $\langle U \rangle$ is virtually $\mathbb{Z}$.

(2) $\lambda_0(U) < K$.

(3) $|U^n| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor}$ for every $n \in \mathbb{N}$.

Remark 4.2.16. We emphasise here that this is not a direct improvement of Theorem 4.2.3, as in both statements the displacement condition is dependent on the action under consideration. In particular, for a certain acylindrically hyperbolic group G, and finite subset U of G, we may have that the displacement of U is large under the acylindrical action of G on some general hyperbolic space, but small under the acylindrical action of G on a quasi-tree.

4.2.3 Loxodromic elements and displacement

To use Theorem 4.2.13 to say anything about the growth of finite subsets of a specific acylindrically hyperbolic group, we must first be able to say something about which finite subsets will have large displacement under some action on a quasi-tree.

One difficulty with this is that the displacement λ_0 is dependent on the choice of basepoint x_0, which is dependent on the finite set U. It is easier instead to try to find sets U such that $\max_{u\in U} d(x, ux)$ is large for every point x in the space being acted on. As is often the case with exponential growth questions, this can be done by finding loxodromic elements.

Proposition 4.2.17. [15] *Let G be a group acting acylindrically on a δ-hyperbolic space. There exists $\nu > 0$, dependent only on δ and the acylindricity constants, such that if $g \in G$ is loxodromic then $\tau(g) \geqslant \nu$.*

The following application of this to Theorem 4.2.13 was pointed out by Thomas Delzant.

Proposition 4.2.18. *Let G be a group acting acylindrically on a quasi-tree. There exist $\alpha, \beta > 0$ such that for every finite $U \subset G$ such that U contains a loxodromic element and $\langle U\rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.*

Proof. Let X be the quasi-tree on which G acts acylindrically. Let $u \in U$ be our given loxodromic element. By Proposition 4.2.17 we know that there exists $\nu > 0$, dependent only on the bottleneck constant of X and the acylindricity constants, such that $\tau(u) \geqslant \nu$. Let $K > 0$ be as in Theorem 4.2.13, and let $m = \lceil \frac{K}{\nu} \rceil$. It is clear from the definition of the stable translation length that $\tau(u^m) = m\tau(u) \geqslant K$.

Let $x \in X$ be arbitrary, and note that for $n \in \mathbb{N}$ we have that $d(x, u^{nm}x) \leqslant nd(x, u^m x)$, so $\tau(u^m) \leqslant d(x, u^m x)$. Therefore, as $u^m \in U^m$, we conclude that $\lambda_0(U^m) \geqslant K$. We assumed that $\langle U\rangle$ is not virtually $\mathbb{Z}$, so we can apply Theorem 4.2.13 to say that

$$|U^{nm}| \geqslant (\alpha|U^m|)^{\lfloor \frac{n+1}{2} \rfloor} \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor},$$

for all $n \in \mathbb{N}$, where $\alpha > 0$ is dependent only on the bottleneck constant of X and the acylindricity constants of the action.

Suppose $i \in \mathbb{N}$ is such that $i \geqslant m$. Then there exists $n \in \mathbb{N}$ such that $nm \leqslant i < (n+1)m$. Therefore

$$|U^i| \geqslant |U^{nm}| \geqslant (\alpha|U|)^{\lfloor \frac{n+1}{2} \rfloor} \geqslant (\alpha|U|)^{\frac{i}{4m}}.$$

Now suppose $i < m$. If $\alpha|U| < 1$ then it is trivial that $|U^i| \geqslant (\alpha|U|)^{\frac{i}{4m}}$, as $|U^i| \geqslant 1$. If we instead have that $\alpha|U| \geqslant 1$, then Theorem 4.2.13 tells us that $\alpha < 1$, so $|U^i| \geqslant |U| \geqslant \alpha|U| \geqslant (\alpha|U|)^{\frac{i}{4m}}$, as $\frac{i}{4m} < 1$. □

Remark 4.2.19. The constants α and β are dependent on the acylindricity constants of the action, and the bottleneck constant of the space being acted on, so two different acylindrical actions on quasi-trees by the same group will give different constants.

Corollary 4.2.20. *Let G be a group acting acylindrically on a quasi-tree. There exist $\alpha, \beta > 0$ such that for every finite $U \subset G$ such that U^k contains a loxodromic element for $k \in \mathbb{N}$, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\frac{\beta n}{k}}$ for every $n \in \mathbb{N}$.*

Proof. The set U^k satisfies the conditions of Proposition 4.2.18, so we have that

$$|U^n| \geqslant |U^{nk}|^{\frac{1}{k}} \geqslant (\alpha|U^k|)^{\frac{\beta n}{k}} \geqslant (\alpha|U|)^{\frac{\beta n}{k}}.$$

□

The growth we get for a certain subset of our group is therefore dependent on k, however if we can find an upper bound on k then we will have a lower bound for the growth. In other words, given a class of finite subsets of G, if we can always generate a loxodromic element on quasi-tree within a bounded product of each set, then every U in this class will satisfy $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$ for some uniform constants $\alpha, \beta > 0$. We note here that these loxodromic elements do not have to be obtained from a single acylindrical action on a quasi-tree, we only require that the quasi-trees and associated acylindrical actions in question admit an upper bound on the bottleneck constants and acylindricity constants.

Corollary 4.2.21. *Let G be a group. Suppose there exist $N \in \mathbb{N}, \Delta > 0, \kappa_0 \geqslant \Delta, N_0 \geqslant 1$ such that for every finite $U \subset G$, where $\langle U \rangle$ is not virtually $\mathbb{Z}$, there exists a quasi-tree X with bottleneck constant Δ such that $\langle U \rangle$ acts (κ_0, N_0)-acylindrically on X, and for some $n \leqslant N$ the set U^n contains a loxodromic element for this action. Then there exist $\alpha, \beta > 0$ such that for every finite $U \subset G$, where $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.*

Remark 4.2.22. In particular, the conditions of this corollary will be satisfied if the loxodromic elements can be found from a finite collection of quasi-trees and acylindrical actions, as then the associated constants will automatically be bounded.

The statement of Corollary 4.2.21 concerns every finite $U \subset G$ such that $\langle U \rangle$ is not virtually $\mathbb{Z}$, however it could also have been phrased to instead only consider those U that also satisfy some additional property. This is important, as without any additional restrictions this class of subsets would generally be far too wide to prove such a strong statement about, as noted in Section 3.6.

4.3 Growth from actions on hyperbolic spaces

In the previous section we showed that being able to quickly generate loxodromic elements in an acylindrical action on a quasi-tree is enough to prove uniform product set growth. Theorem 4.1.10 tells us that every acylindrically hyperbolic group admits an acylindrical action on a quasi-tree, however this construction does not necessarily give us an easy way to find loxodromic elements. There are groups with more natural actions on quasi-trees where this is possible to do, as we will see in Chapter 5, however for many acylindrically hyperbolic groups their most natural acylindrical action is on a hyperbolic space that is not a quasi-tree.

In such a case, we can instead use a result of Fujiwara, who showed that the statements we proved in Subsection 4.2.3 also hold for finite symmetric subsets of groups that act acylindrically on hyperbolic spaces [37], using the idea of generating free subgroups with rank proportional to the size of the initial set. Note that this does not replace the generality of Theorem 4.2.13, as here we only consider symmetric subsets, and do not allow for other possible ways of finding sets of large displacement. On the other hand, it does give us a practical way of answering Question 5 for a wider variety of acylindrically hyperbolic groups.

Proposition 4.3.1. [37] *Let G be a group acting acylindrically on a hyperbolic space. There exist $a, b \in \mathbb{N}$ such that for every finite symmetric $U \subset G$ such that U^k contains a loxodromic element for $k \in \mathbb{N}$, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that there exists $F \subset U^{bk}$ such that $|F| \geqslant \frac{|U|}{a}$ and F generates a free subgroup of rank $|F|$.*

This allowed Fujiwara to obtain an alternative version of Corollary 4.2.20, for finite symmetric subsets of groups with acylindrical actions on hyperbolic spaces.

Corollary 4.3.2. [37] *Let G be a group acting acylindrically on a hyperbolic space. There exist $\alpha, \beta > 0$ such that for every finite symmetric $U \subset G$ such that U^k contains a loxodromic element for $k \in \mathbb{N}$, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha |U|)^{\frac{\beta n}{k}}$ for every $n \in \mathbb{N}$.*

Proof. By Proposition 4.3.1, there exist $a, b \in \mathbb{N}$ that are not dependent on U such that we can find $F \subset U^{bk}$, where $|F| \geqslant \frac{|U|}{a}$ and F generates a free subgroup of rank $|F|$. This gives us that

$$|U^n| \geqslant |U^{nbk}|^{\frac{1}{bk}} \geqslant |F^n|^{\frac{1}{bk}} = |F|^{\frac{n}{bk}} \geqslant \left(\frac{1}{a}|U|\right)^{\frac{n}{bk}}.$$

Letting $\alpha = \frac{1}{a}$ and $\beta = \frac{1}{b}$ completes the proof. □

As in the quasi-tree case with Corollary 4.2.21, this means that if we have a collection of subgroups that each act acylindrically on one of a collection of hyperbolic spaces, such that all the constants involved are bounded, and for each symmetric generating set of a subgroup we can quickly generate a loxodromic element in one of these actions, then this will give us our desired uniform product set growth.

5

Applications to subgroups of mapping class groups

But Pure Mathematicians are noble and good,
And wonderful people, though misunderstood!
– Kenneth Falconer, A Poem of Academic Prejudice

The aim of this chapter is to apply the results of Chapter 4 to answer Question 5 for mapping class groups. That is, we would like to find a dichotomy for the finitely generated subgroups of our mapping class group, where either $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every symmetric generating set of our subgroup, with $\alpha, \beta > 0$ dependent only on the mapping class group in question, or our subgroup cannot satisfy this property for any $\alpha, \beta > 0$.

Answering this question for mapping class groups also answers it for any group that embeds as a subgroup of a mapping class group, which includes right-angled Artin groups. We will however deal with the right-angled Artin group case separately here. The reason for doing this is that right-angled Artin groups have a natural acylindrical action on an associated quasi-tree called the extension graph, and so in proving something about the growth of these groups we have to prove that we can quickly generate loxodromic elements on these extension graphs. This is analogous to an already known result for mapping class groups [63].

The other reason for treating right-angled Artin group separately is that, as they naturally act acylindrically on quasi-trees, the only reason we need to restrict ourselves to symmetric subsets is because of the method we use to find loxodromic elements. If a non-symmetric way of doing this were to be found, then a non-symmetric version of uniform product set growth would automatically follow from the results in this chapter.

In Section 5.1 we will recall some necessary background information about right-angled Artin groups and mapping class groups. In Section 5.2 we will prove our result regarding the uniformly quick generation of loxodromic elements in right-angled Artin groups, and in Section 5.3 we use this to obtain a product set growth result. We then finish by extending this result to mapping class groups in Section 5.4.

5.1 Right-angled Artin groups and mapping class groups

We begin by gathering some basic definitions and results relating to right-angled Artin groups and mapping class groups.

Definition 5.1.1. Let $\Gamma = (V, E)$ be a graph, where V is its vertex set and E is its edge set. We say that Γ is *simple* if it does not contain any double edges or loops.

Definition 5.1.2. For a finite simple graph $\Gamma = (V, E)$, its *right-angled Artin group* is

$$A(\Gamma) = \langle\ V \mid [v, w] = 1 \text{ if } \{v, w\} \in E\ \rangle.$$

Example 5.1.3. The following are examples of right-angled Artin groups:

(1) If Γ does not contain any edges, then $A(\Gamma)$ will be the free group of rank $|V|$.

(2) If Γ is a complete graph, which is a graph where every pair of vertices are joined by an edge, then $A(\Gamma)$ will be the free abelian group of rank $|V|$.

(3) If Γ is the union of two connected components Γ_1 and Γ_2, then $A(\Gamma)$ will be the free product $A(\Gamma_1) * A(\Gamma_2)$.

Remark 5.1.4. Right-angled Artin groups are always torsion-free.

Every right-angled Artin group $A(\Gamma)$ has another associated graph called the extension graph, first defined in [56], where the vertices are conjugates of the vertices in Γ, and two vertices have an edge between them if and only if they commute as elements of $A(\Gamma)$.

Definition 5.1.5. Let $\Gamma = (V, E)$ be a finite simple graph, and let $A(\Gamma)$ be its right-angled Artin group. The *extension graph* of Γ is $\Gamma^e = (V^e, E^e)$, where $V^e = \{gvg^{-1} : v \in V, g \in A(\Gamma)\}$ and $E^e = \{\{g, h\} : g, h \in V^e \text{ and } [g, h] = 1\}$.

So long as the defining graph Γ is connected, we get the following two results.

Lemma 5.1.6. [56] *Let Γ be a finite simple connected graph. The extension graph Γ^e is a quasi-tree.*

Theorem 5.1.7. [57] *Let Γ be a finite simple connected graph. The action of $A(\Gamma)$ on Γ^e is acylindrical.*

In Section 5.2 we will be interested in generating loxodromic elements in the action of $A(\Gamma)$ on Γ^e. This is made easier by the fact that there exists a characterisation of these elements, given in [57]. To state this characterisation, we first need to give some terminology.

Remark 5.1.8. All graphs in the context of right-angled Artin groups will be assumed to be simple, and with the exception of the extension graph they will also be assumed to be finite. Unless stated otherwise, we will denote the vertex set of a graph Γ by V, and the edge set by E.

Definition 5.1.9. The *complement* of a graph Γ is $\Gamma^c = (V, E^c)$, where $E^c = \{\{v, w\} : v, w \in V \text{ and } v \neq w\} \backslash E$.

Definition 5.1.10. The *join* of two graphs Γ_1 and Γ_2 is $\Gamma_1 * \Gamma_2 = (\Gamma_1^c \sqcup \Gamma_2^c)^c$.

Alternatively, we can see that if $\Gamma_1 = (V_1, E_1)$ and $\Gamma_2 = (V_2, E_2)$ then $\Gamma_1 * \Gamma_2 = (V_1 \cup V_2, E)$, where $E = E_1 \cup E_2 \cup \{\{s_1, s_2\} : s_1 \in V_1 \text{ and } s_2 \in V_2\}$.

Definition 5.1.11. A graph Γ is said to *split as a non-trivial join* if $\Gamma = \Gamma_1 * \Gamma_2$, where Γ_1 and Γ_2 are non-empty graphs. A *subjoin* of a graph is a subgraph that splits as a non-trivial join.

Example 5.1.12. If $\Gamma = \Gamma_1 * \Gamma_2$, then $A(\Gamma) = A(\Gamma_1) \times A(\Gamma_2)$.

Let $\Gamma = (V, E)$ be a graph, and let $V = \{s_1, \ldots, s_k\}$. Recall that $A(\Gamma)$ is generated by V, so $g \in A(\Gamma)$ can be represented by a word w written in the alphabet $V \cup V^{-1}$.

Definition 5.1.13. A word w representing $g \in A(\Gamma)$ in the alphabet $V \cup V^{-1}$ is *reduced* if its length is the same as the word length of g. In other words, it is a minimal length representative of g. We will often identify g with a reduced word representing it, and say that g is reduced.

The word w is *cyclically reduced* if every cyclic permutation of w is also reduced. Equivalently, it is a minimal length representative of the conjugation class of g.

Remark 5.1.14. This is very similar to the notions of reduced and cyclically reduced words in a free group. The difference is that in the free group every element has a unique reduced word representing it, whereas in a right-angled Artin group each reduced representative is only unique up to being able to swap the order of vertices when they commute. See Section 2 of [3] for a more detailed explanation of this.

Remark 5.1.15. Every reduced word w can be written in the form $w = uw'u^{-1}$, where w' is a cyclically reduced word, and u is a possibly empty word. This w' is a minimal length representative of the conjugation class of w.

Definition 5.1.16. The *support* of a word w is the set of vertices $s \in V$ such that s or s^{-1} is a letter of w. The support of $g \in A(\Gamma)$ is the support of a reduced word that represents g.

Remark 5.1.17. By Remark 5.1.14, the support of $g \in A(\Gamma)$ does not depend on the reduced word chosen.

The characterisation of elliptic and loxodromic elements in the action of $A(\Gamma)$ on Γ^e is then as follows.

Theorem 5.1.18. [57] *Let Γ be a connected graph that is not an isolated vertex, and suppose $1 \neq g \in A(\Gamma)$ is cyclically reduced. Then g is elliptic in the action of $A(\Gamma)$ on Γ^e if and only if the support of g is contained in a subjoin of Γ.*

Although this theorem only mentions cyclically reduced elements, we know from Remark 5.1.15 that every $g \in A(\Gamma)$ is conjugate to a cyclically reduced element, and both elliptic and loxodromic elements are preserved under conjugation.

We now move on to giving some basic definitions and facts about mapping class groups.

Definition 5.1.19. Let S be an oriented surface with finite genus and finitely many boundary components, punctures, and connected components. The *mapping class group* of S, denoted $MCG(S)$, is the group of orientation preserving isotopy classes of homeomorphisms of S that restrict to the identity on the boundary ∂S, where the isotopies fix components of the boundary pointwise. These isotopy classes are called *mapping classes*.

Convention. All surfaces in this chapter will have finite genus and finitely many boundary components, punctures, and connected components, which is known as being of *finite type*. We will also assume that every surface in this chapter is oriented.

Example 5.1.20. The following are examples of mapping class groups:

(1) If S is a sphere, then $MCG(S)$ will be the trivial group.

(2) If S is an annulus, then $MCG(S)$ will be isomorphic to $\mathbb{Z}$.

(3) If S is a torus, then $MCG(S)$ will be isomorphic to $SL_2(\mathbb{Z})$.

Definition 5.1.21. The *sporadic* surfaces are the sphere with up to four punctures, and the torus with up to one puncture.

Remark 5.1.22. The mapping class group of a sporadic surface is hyperbolic [35, 73].

In general this will not be true, however as in the right-angled Artin group case, we can associate a hyperbolic space to our mapping class group.

Definition 5.1.23. A *simple closed curve* on a surface S is the image of a continuous injective function $f : S^1 \to S$.

Definition 5.1.24. A simple closed curve on a surface S is *essential* if it is not homotopic to a point, a puncture, or to any component of the boundary ∂S.

Definition 5.1.25. Let S be a connected surface. Its *curve complex* $C(S)$ is the simplicial complex where the vertices are the isotopy classes of essential simple closed curves in S, and a set of these isotopy classes span a simplex if and only if they have representatives that are pairwise disjoint.

Theorem 5.1.26. [67] *Let S be a non-sporadic connected surface. The curve complex $C(S)$ is a hyperbolic space.*

Remark 5.1.27. The hyperbolicity constant of the curve complex is independent of S [4, 24]. In particular, we can take $\delta = 17$, and this will work for every connected surface [47].

As with the extension graph for a right-angled Artin group, it turns out that the action of a mapping class group on its associated curve complex is acylindrical in most cases.

Theorem 5.1.28. [15] *Let S be a non-sporadic connected surface without boundary. The mapping class group $MCG(S)$ acts acylindrically on the curve complex $C(S)$.*

Remark 5.1.29. A mapping class group of a surface with non-empty boundary would have an infinite finitely generated centre [35], and so would not have uniform product set growth by Proposition 3.6.2. It would also not be acylindrically hyperbolic [73], and for this reason in many places the mapping class group is defined either for surfaces without boundary, or the boundary is fixed setwise rather than pointwise. In the latter case the boundary components can effectively be seen as punctures that are not allowed to permute, and this version of the mapping class group is a finite index subgroup of the version where the punctures are allowed to permute. Sometimes the mapping class group is defined in such a way as to allow the boundary components themselves to permute, as in [51] for example, in which case this is exactly the same group as considering the boundary components as punctures. Although in our definition we fix our boundary pointwise, we will later see that we can easily reduce to the case where we only consider punctures.

Again, as with the extension graph, there exists a characterisation of the loxodromic elements in this action.

Definition 5.1.30. A mapping class $f \in MCG(S)$ is a *pseudo-Anosov* if there is no isotopy class of an essential simple closed curve in S that is fixed by a power of f (up to isotopy). A mapping class f is said to be *pseudo-Anosov on a connected subsurface* S' if there exists a representative of f such that its restriction to S' gives a pseudo-Anosov mapping class in $MCG(S')$.

Proposition 5.1.31. [67] *Let S be a non-sporadic surface without boundary. A mapping class $f \in MCG(S)$ is a loxodromic element in the acylindrical action on the curve complex $C(S)$ if and only if it is a pseudo-Anosov.*

When we consider mapping class groups, we will therefore not talk about loxodromic elements, but refer to pseudo-Anosovs instead. In other words, to apply our results from Chapter 4 to mapping class groups, we will need the ability to quickly generate pseudo-Anosovs, which is provided by the following theorem.

Theorem 5.1.32. [63] *Let S be a non-sporadic connected surface without boundary. There exists a constant $N = N(S) \in \mathbb{N}$ such that for any finite symmetric $U \subset MCG(S)$ where $\langle U \rangle$ contains a pseudo-Anosov, there exists $n \leqslant N$ such that U^n contains a pseudo-Anosov.*

The problem with using this strategy to show that subgroups of mapping class groups have uniform product set growth is that many interesting subgroups may not

contain a pseudo-Anosov on the whole surface. The standard way to deal with this situation is to cut our surface in some canonical way, and consider the groups induced by the restrictions of our mapping classes to the remaining subsurfaces.

In doing this cutting along curves we may lose a few elements of our subgroup, specifically those that act as twists along those curves.

Definition 5.1.33. Let α be a simple closed curve in a surface S, let $A = S^1 \times [0,1]$ be an annulus, and let N_α be a neighbourhood of α such that there is a homeomorphism $\phi : A \to N_\alpha$. Let $\psi : A \to A$ be the map defined by $\psi(e^{2\pi i x}, y) = (e^{2\pi i(x+y)}, y)$. We can then define a homeomorphism on S by

$$\tau_\alpha(x) = \begin{cases} \phi \circ \psi \circ \phi^{-1}(x) & x \in N_\alpha \\ x & x \notin N_\alpha. \end{cases}$$

The mapping class of τ_α is called the *Dehn twist about* α, denoted T_α.

Remark 5.1.34. If α and β are isotopic simple closed curves in S then $T_\alpha = T_\beta$, assuming that the homeomorphism $N_\alpha \to A \to N_\beta$ is orientation preserving [35].

Proposition 5.1.35 (Cutting homomorphism)**.** [35, Proposition 3.20] *Let S be a surface, and let $\{\gamma_1, \ldots, \gamma_k\}$ be a set of disjoint and isotopically distinct essential simple closed curves in S. We denote by $MCG(S, \{\gamma_1, \ldots, \gamma_n\})$ the subgroup of $MCG(S)$ containing the mapping classes that fix this set of curves. Then there exists a natural homomorphism $\varphi : MCG(S, \{\gamma_1, \ldots, \gamma_n\}) \to MCG(S \setminus \bigcup_{i=1}^k \gamma_i)$, with kernel $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle$, where T_{γ_i} is a Dehn twist around γ_i.*

When we cut along these curves we end up with a collection of possibly smaller surfaces, with punctures where the curves used to be. Any Dehn twist about a puncture is isotopic to the identity, which is why these are exactly the mapping classes which appear in the kernel of the cutting homomorphism.

The important facts for us to know here about Dehn twists are that they are infinite order elements, and are central in any group that fixes γ.

Lemma 5.1.36. [35, Fact 3.7] *Let γ be a simple closed curve in a surface S. Let T_γ be a Dehn twist around γ, then for every $g \in MCG(S)$ we have that $gT_\gamma g^{-1} = T_{g(\gamma)}$.*

Corollary 5.1.37. [35, Fact 3.8] *Let γ be a simple closed curve in a surface S. The Dehn twist T_γ commutes with every mapping class in $MCG(S, \{\gamma\})$.*

In particular, Dehn twists around disjoint and isotopically distinct simple closed curves commute with each other, so in Proposition 5.1.35 we have that $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cong \mathbb{Z}^k$.

Recall from Section 3.6.1 that if a group has an infinite order element in its centre, then it does not have uniform product set growth. This means that if we work with subgroups of $MCG(S)$ that fix individual curves (rather than possibly permute a set of them), then ruling out the subgroups that have such infinite order central elements will mean that the restriction of φ in Proposition 5.1.35 will give us an injective homomorphism into $MCG(S \setminus \bigcup_{i=1}^{k} \gamma_i)$. If we further assume that our subgroup does not permute the connected components of $S \setminus \bigcup_{i=1}^{k} \gamma_i$, then this will be a map into a direct product, which Section 3.6 tells us how to deal with.

For these reasons, it makes sense for us to temporarily restrict ourselves to thinking about the pure subgroups of mapping class groups.

Definition 5.1.38. A mapping class $f \in MCG(S)$ is said to be *pure* if there exists a set $\{\gamma_1, \ldots, \gamma_k\}$ of disjoint and isotopically distinct essential simple closed curves in S such that f fixes each γ_i, does not permute the connected components of $S \setminus \bigcup_{i=1}^{k} \gamma_i$, and there exists a representative of f such that the restriction of this representative to any connected component of $S \setminus \bigcup_{i=1}^{k} \gamma_i$ is either pseudo-Anosov or the identity. A subgroup of $MCG(S)$ is pure if every element in that subgroup is a pure mapping class.

Not only are pure subgroups easier to work with, it turns out that every subgroup of mapping class group has a pure subgroup of finite index. A common approach to studying subgroups is therefore to first work with the pure subgroups, and then extend what we find to their supergroups. This is the approach we will take when trying to prove uniform product set growth for certain subgroups of mapping class groups, with the extension to supergroups allowed by Proposition 3.6.25.

Theorem 5.1.39. [51] *Every mapping class group has a pure normal subgroup of finite index.*

Given a subgroup of a mapping class group, we can take a pure finite index subgroup, cut along a set of curves that are fixed by this pure subgroup, and consider the restriction to subsurfaces on which our group is not the identity. There may however be many curves that are fixed, some of which may intersect each other. Fortunately there is a canonical way of choosing a multicurve that allows us to cut our surface such that the restriction of our pure subgroup to every component is either the identity, or contains a pseudo-Anosov.

Definition 5.1.40. For a pure subgroup $G \leqslant MCG(S)$, the *canonical reduction multicurve* σ is the (possibly empty) union of all isotopically distinct essential simple closed curves γ in S such that G fixes γ, but if any other essential simple closed curve ξ intersects γ then G does not fix ξ.

Theorem 5.1.41. [51] *Let $G \leqslant MCG(S)$ be a pure subgroup of a mapping class group. Let $\{\gamma_1, \ldots, \gamma_k\}$ be the set of curves in the canonical reduction multicurve for G. Let $\mathcal{S}$ be the set of connected components of $S \setminus \bigcup_{i=1}^{k} \gamma_i$. Then the homomorphism in Proposition 5.1.35 restricts to a homomorphism $\varphi : G \to \Pi_{\Sigma \in \mathcal{S}} MCG(\Sigma)$, with kernel $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap G$. Moreover, the projection of G to any factor $MCG(\Sigma)$ is either trivial, or contains a pseudo-Anosov on Σ.*

Remark 5.1.42. It follows from Corollary 5.1.37 that $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap G$ is central in G, as $g(\gamma_i) = \gamma_i$ for every $g \in G$. We also note here that if $G' \leqslant MCG(S)$ has G as a finite index subgroup, then although G' does not fix each γ_i, it does fix the set $\{\gamma_1, \ldots, \gamma_k\}$ (see [51]). Hence by Lemma 5.1.36 we have that $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap G'$ is normal in G'.

To apply our results from Section 3.6.3 to the direct product in Theorem 5.1.41, we need a bound on the number of factors. The following result allows us to do this.

Theorem 5.1.43. [11] *Let S be a surface of genus g with p punctures, and c connected components. A free abelian subgroup of $MCG(G)$ has rank at most $3g + p - 3c$.*

Mangahas' result regarding the quick generation of pseudo-Anosovs (see Theorem 5.1.32) was in fact a particular case of a more general result that she proved. The rough idea is that if we take the union of all subsurfaces on which some element of our group acts as a pseudo-Anosov, then we can find an element that acts as a pseudo-Anosov on this entire union within some bounded product of our generating set. The language defined here will be useful in our proofs about right-angled Artin groups.

Definition 5.1.44. The *active subsurface* $\mathcal{A}(G)$ of a pure subgroup $G \leqslant MCG(S)$, with canonical reduction multicurve σ, is the union of the connected components of $S \setminus \sigma$ such that some mapping class in G is pseudo-Anosov on that connected component, plus the annuli that are the neighbourhood of any $\gamma \in \sigma$ that are not in the boundary of an already selected component. The active subsurface of a pure mapping class f is the active subsurface of $\langle f \rangle$.

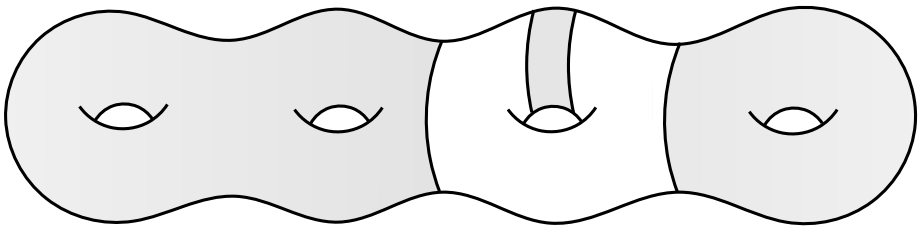


Figure 5.1: An example of an active subsurface

Remark 5.1.45. We note here that if $G \leqslant MCG(S)$ is a pure subgroup generated by a finite set U, then the union of the active subsurfaces of generators in U is a subset of the active subsurface of G. Moreover, if the active subsurface of G is disconnected then the union of the active subsurfaces of generators will also be disconnected, as if every generator fixes a subsurface then so does G.

The fact that every subgroup of a mapping class group has a pure finite index subgroup allows the extension of the definition of the active subsurface to all subgroups of the mapping class group.

Definition 5.1.46. The *active subsurface* $\mathcal{A}(G)$ of an arbitrary subgroup $G \leqslant MCG(S)$, with pure finite index subgroup H, is given by $\mathcal{A}(G) = \mathcal{A}(H)$.

This subsurface does not depend on the choice of finite index subgroup [51]. With these definitions we can now state Mangahas' full result.

Theorem 5.1.47. [63] *Let S be a non-sporadic connected surface without boundary. There exists a constant $N = N(S) \in \mathbb{N}$ such that for any finite symmetric $U \subset MCG(S)$ there exists $n \leqslant N$ and $f \in U^n$ such that f has the same active subsurface as $\langle U \rangle$.*

Many of the tools that we have introduced here require that the surface in question has no boundary. As noted in Remark 5.1.29, if the surface had non-empty boundary then the mapping class group would have infinite centre, namely the subgroup generated by a Dehn twist around a curve parallel to the boundary. This means that it would not be acylindrically hyperbolic, so many of our tools would not apply, however Proposition 3.6.2 tells us that it would also not have uniform product set growth.

To deal with subgroups of such groups, we will employ the capping homomorphism, which effectively says that either our subgroup has infinite centre, or we can view it as a subgroup of a mapping class group of a surface without boundary.

Proposition 5.1.48 (Capping homomorphism)**.** [35, Theorem 3.18] *Let S be a surface with boundary components $\{\xi_1, \dots, \xi_l\}$, and let S' be the surface obtained from S by capping each boundary component by a once punctured disc. Then there exists a natural homomorphism $\psi : MCG(S) \to MCG(S')$, with kernel $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle$, where T_{ξ_i} is a Dehn twist around ξ_i.*

Remark 5.1.49. As every element of $MCG(S)$ fixes the boundary components by definition, we have by Corollary 5.1.37 that $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle$ is central in $MCG(S)$.

5.2 Short loxodromics in right-angled Artin groups

Recall that the aim of this chapter is to say something about product set growth in right-angled Artin groups and mapping class groups. In the right-angled Artin groups case we would like to use Corollary 4.2.20 to prove uniform product set growth for as many subgroups as possible. To do so we need to show that we can quickly generate loxodromic elements in the action of our group on its extension graph. We already have a characterisation of these loxodromic elements, as given by Theorem 5.1.18, and we also have a result about the quick generation of pseudo-Anosovs in subgroups of mapping class groups from Theorem 5.1.47.

In this section we will use these facts to get an analogous result to Theorem 5.1.47 for right-angled Artin groups, using the existence of a natural embedding of right-angled Artin groups into mapping class groups [25]. The idea is that we can embed our right-angled Artin group into a mapping class group, then use Theorem 5.1.47 to quickly generate a mapping class within this embedding that has the same active subsurface as the right-Angled Artin group. Under certain circumstances this mapping class will correspond to a loxodromic element in the original right-angled Artin group.

There are several ways to see right-angled Artin groups as subgroups of mapping class groups, see for example Section 7.3 of [59]. The construction we use here is taken from Section 2.4 of [25], and for a more detailed picture of it we refer to there.

Given a finite simple graph Γ, with vertex set $V = \{s_1, \dots, s_k\}$, there exists a surface S and a collection of non-annular connected subsurfaces $\mathbb{X} = \{X_1, \dots, X_k\}$ such that:

(1) $X_i \cap X_j = \emptyset$ if and only if $\{s_i, s_j\}$ is an edge in Γ.

(2) $X_i \cap X_j \neq \emptyset$ if and only if X_i and X_j cannot be isotoped to be disjoint.

(3) If $X_i \cap X_j \neq \emptyset$ then this intersection is homeomorphic to a disc.

(4) X_i and X_j cannot be isotoped such that $X_i \subset X_j$ when $i \neq j$.

(5) Each X_i is a twice punctured torus.

A picture of such a surface can be found in Figure 4 of [25]. For each X_i we then pick an f_i in $MCG(S)$ that is pseudo-Anosov on X_i, and the identity elsewhere.

Theorem 5.2.1. [25] *Given the collection $\mathbb{F} = \{f_1, \dots, f_k\}$, there exists $M \in \mathbb{N}$ such that the map $\phi : A(\Gamma) \to MCG(S)$ defined by $\phi(s_i) = f_i^M$ is an injective homomorphism.*

This statement allows us to see $A(\Gamma)$ as a subgroup of $MCG(S)$. In particular, every element of $A(\Gamma)$ will be sent to a mapping class that is pseudo-Anosov on some collection of connected subsurfaces of S, and is the identity elsewhere. These subsurfaces can be constructed explicitly. The following notation is taken directly from [25].

Notation 5.2.2. Let $\{X_1, \dots, X_r\} \subset \mathbb{X}$. We denote by $\mathrm{Fill}(X_1, \dots, X_r)$ the minimal essential subsurface of S that contains $X_1 \cup \dots \cup X_r$. That is, if $X_1 \cup \dots \cup X_r$ has any discs in its complement then we add these discs to get $\mathrm{Fill}(X_1, \dots, X_r)$.

Remark 5.2.3. Note that $\mathrm{Fill}(X_1, \dots, X_r)$ is connected if and only if $X_1 \cup \dots \cup X_r$ is connected, and that $\mathrm{Fill}(X_i) \cap \mathrm{Fill}(X_j) = \emptyset$ if and only if $X_i \cap X_j = \emptyset$.

Notation 5.2.4. Suppose that $1 \neq g \in A(\Gamma)$ is cyclically reduced, and that r is the minimal number such that g is a word in the first r generators $s_1, \dots, s_r$, changing the indices if necessary. Then we say that $\mathrm{Fill}(g) = \mathrm{Fill}(X_1, \dots, X_r)$. If $g = hg'h^{-1}$ for some $h, g' \in A(\Gamma)$ then $\mathrm{Fill}(g) = \phi(h)\mathrm{Fill}(g')$. If $g = 1$, we say that $\mathrm{Fill}(g) = \emptyset$.

Theorem 5.2.5. [25] *For any $g \in A(\Gamma)$ we have that $\phi(g)$ is pseudo-Anosov on each connected component of Fill(g), and is the identity elsewhere.*

Remark 5.2.6. In other words, Theorem 5.2.5 tells us that $A(\Gamma)$ can be seen as a pure subgroup of $MCG(S)$, and for any $g \in A(\Gamma)$ the active subsurface of $\phi(g)$ is exactly $\mathrm{Fill}(g)$.

We can relate this construction to the characterisation of loxodromic elements given in Theorem 5.1.18.

Lemma 5.2.7. *Suppose Γ is a connected graph that is not an isolated vertex, and let $1 \neq g \in A(\Gamma)$ be cyclically reduced. Then g is a loxodromic on Γ^e if and only if the subsurface $\mathrm{Fill}(g)$ is connected, and $\mathrm{Fill}(g) \cap \mathrm{Fill}(s) \neq \emptyset$ for all $s \in V$.*

Proof. Let $\{s_1, \dots, s_r\}$ be the vertices in the support of g. By Theorem 5.1.18 we have that g is elliptic on Γ_e if and only if $\{s_1, \dots, s_r\}$ is contained in a subjoin of Γ. This is the case if and only if $\{s_1, \dots, s_r\}$ induces a subjoin, or there exists $s \in V$ such that $\{s_1, \dots, s_r\}$ is contained in the 1-neighbourhood of s.

Now $\{s_1, \dots, s_r\}$ induces a subjoin if and only if we can write it as a disjoint partition $A \sqcup B$ such that A and B are non-empty and for every $a \in A$, $b \in B$, we have that $\{a, b\}$ is an edge in Γ. By construction of ϕ, this is true if and only if the corresponding partition $A' \sqcup B'$ of $\{X_1, \dots, X_r\}$ is such that for every $X \in A'$, $Y \in B'$, we have that $X \cap Y = \emptyset$. This correspondence gives us that $\{s_1, \dots, s_r\}$ induces a subjoin if and only if $X_1 \cup \dots \cup X_r$ is disconnected, which is the case if and only if $\mathrm{Fill}(g) = \mathrm{Fill}(X_1, \dots, X_r)$ is disconnected.

For the other case, there exists $s \in V$ such that $\{s_1, \dots, s_r\}$ lies in the 1-neighbourhood of s if and only if there exists $s \in V$ such that $\mathrm{Fill}(s_i) \cap \mathrm{Fill}(s) = \emptyset$ for every s_i. This is equivalent to saying that $\mathrm{Fill}(g) \cap \mathrm{Fill}(s) = \emptyset$. We therefore conclude that g is elliptic if and only if $\mathrm{Fill}(g)$ is disconnected, or there exists $s \in V$ such that $\mathrm{Fill}(g) \cap \mathrm{Fill}(s) = \emptyset$. □

Our strategy now is that given a finite symmetric $U \subset A(\Gamma)$, we can consider $\phi(U) \subset MCG(S)$, and use Theorem 5.1.47 to find $g \in U^n$ such that $\phi(g)$ has the same active subsurface as $\phi(\langle U \rangle)$. By Lemma 5.2.7, if we want this g to be a loxodromic on the extension graph Γ^e then we need the active subsurface of $\phi(\langle U \rangle)$ to be connected. We will give a sufficient condition for this to be the case.

Lemma 5.2.8. *Let $U \subset A(\Gamma)$ be finite. If the active subsurface of $\phi(\langle U \rangle)$ is disconnected, then $\langle U \rangle \leqslant A(\Gamma)$ is contained non-trivially in a direct product.*

Proof. Let $U = \{u_1, \dots, u_n\}$. Suppose the active subsurface of $\phi(\langle U \rangle)$ is disconnected. As noted in Remark 5.1.45, this means that $\mathrm{Fill}(u_1) \cup \dots \cup \mathrm{Fill}(u_n)$ is disconnected. Write $u_i \in U$ as $u_i = g_i s_{i,1} \cdots s_{i,r_i} g_i^{-1}$ with the $s_{i,j}$ being vertices of Γ and $s_{i,1} \cdots s_{i,r_i}$ being cyclically reduced.

We then have that

$$\begin{aligned}
\mathrm{Fill}(u_i) &= \phi(g_i)\mathrm{Fill}(s_{i,1} \cdots s_{i,r_i}) \\
&= \phi(g_i)\mathrm{Fill}(X_{i,1}, \dots, X_{i,r_i}) \\
&\supset \phi(g_i)(\mathrm{Fill}(X_{i,1}) \cup \dots \cup \mathrm{Fill}(X_{i,r_i})) \\
&= \mathrm{Fill}(g_i s_{i,1} g_i^{-1}) \cup \dots \cup \mathrm{Fill}(g_i s_{1,r_i} g_i^{-1})
\end{aligned}$$

Let S' be a connected component of $\mathrm{Fill}(u_i)$. Then for some j we must have that $\mathrm{Fill}(g_i s_{i,j} g_i^{-1}) \subset S'$, as otherwise every $\phi(g_i s_{i,j} g_i^{-1})$ would be the identity on S', so $\phi(u_i)$ would be the identity on S'. This would contradict the fact that, by Theorem 5.2.5, $\phi(u_i)$ is pseudo-Anosov on S'. We can therefore see that the union $\bigcup_{i=1}^{n} \bigcup_{j=1}^{r_i} \mathrm{Fill}(g_i s_{i,j} g_i^{-1})$ is disconnected.

As each $\mathrm{Fill}(g_i s_{i,j} g_i^{-1})$ is connected, we partition these subsurfaces into $A \sqcup B$, where $A \neq \emptyset$, $B \neq \emptyset$, and all subsurfaces in A are disjoint from those in B. We denote the corresponding partition of the elements $g_i s_{i,j} g_i^{-1}$ by $A' \sqcup B'$, and note that for $a \in A'$, $b \in B'$, we have that $[a, b] = 1$.

We therefore have that

$$\langle U \rangle \leqslant \langle A' \sqcup B' \rangle = \langle A' \rangle \times \langle B' \rangle.$$

We conclude by noting that we cannot have $\langle U \rangle \leqslant \langle A' \rangle$, as then $\phi(\langle U \rangle)$ would act trivially on the subsurfaces in B, which would contradict the fact that these are contained in the active subsurface of $\phi(\langle U \rangle)$. Similarly, we cannot have $\langle U \rangle \leqslant \langle B' \rangle$, so $\langle U \rangle$ is contained non-trivially in a direct product. □

We now need to address the second condition in Lemma 5.2.7, which says that for a cyclically reduced element to be loxodromic in our right-angled Artin group then we need its active subsurface to intersect every subsurface associated to a vertex. To give a sufficient condition for this to be the case, we need the following two lemmas. The method used in the proof of Lemma 5.2.11 was suggested by Ric Wade.

Definition 5.2.9. For a vertex v in a graph Γ, the *link* of v, denoted $\mathrm{link}(v)$, is the set of vertices adjacent to v in Γ.

Notation 5.2.10. If $V' \subset V$, we denote the subgraph of Γ induced by V' as $\Gamma(V')$.

Lemma 5.2.11. *Let Γ be a finite graph, and let $U \subset A(\Gamma)$ be finite, where $U \neq \{1\}$. Suppose $V' \subset V$ is minimal under inclusion such that U is conjugate into $A(\Gamma(V'))$. Then for every $s \in V'$ there exists $g \in U^2$ such that when we write g in the form $g = hg'h^{-1}$, where $h, g' \in A(\Gamma)$ and g' is cyclically reduced, we have that s is in the support of g'.*

Proof. Suppose that $V' \subset V$ is such a minimal set. The conclusion of this lemma is true of U if and only if it is true of any conjugate of U, so we will assume that $U \subset A(\Gamma(V'))$.

For every $g \in \langle U \rangle$, we can write g in the form $g = hg'h^{-1}$, where $h, g' \in A(\Gamma)$ and g' is cyclically reduced. Note that this g' is unique up to cyclic permutation and

reordering letters if they commute. Let $s \in V$. Suppose that for some $g \in U$ we have that s in the support of g'. Then s is in the support of $(g')^2 = (g^2)'$, so we are done. We therefore suppose that for every $g \in U$ we have that s is not in the support of g'.

Let $V'' = V' \backslash \{s\}$. Consider $H = A(\Gamma(V'' \cap \text{link}(s))) \leqslant A(\Gamma(V''))$, and let φ be the identity map on H. From the definition of a HNN extension we can see that this homomorphism gives us that $A(\Gamma(V')) = A(\Gamma(V'')) *_H$. That is, if $A(\Gamma(V'')) = \langle V'' | R \rangle$ then

$$A(\Gamma(V')) = \langle V'' \cup \{s\} \mid R \cup \{svs^{-1} = v \mid v \in V'' \cap \text{link}(s)\} \rangle.$$

Consider the action of $\langle U \rangle \leqslant A(\Gamma(V'))$ on the Bass-Serre tree associated with this HNN extension. As every $g \in U$ is conjugate into $A(\Gamma(V''))$, every g has a fixed point in this action. If the same was true of every $g \in U^2$, then by [78, p. 64] we would have that $\langle U \rangle$ has a fixed point in this action, and so U is conjugate into $A(\Gamma(V''))$. This contradicts the minimality of V', therefore for some $g \in U^2$ we have that s is in the support of g'. □

The other lemma we need is the following.

Lemma 5.2.12. *For any $s \in V$ and $g \in A(\Gamma)$, we have that* $\text{Fill}(s)$ *and* $\phi(g)\text{Fill}(s)$ *cannot be isotoped to be disjoint.*

Proof. Let $s \in V$, and consider $g \in A(\Gamma)$ written as a reduced word $g = g_1 \cdots g_n$, where each g_i is a vertex or its inverse. We will prove our statement by induction on how many of the g_i's are not equal to s or s^{-1}, so let $\mu(g) = |\{g_i : g_i \neq s^{\pm 1}\}|$.

We will show that for every $g \in A(\Gamma)$ there exists a genus one subsurface $T^{(g)}$ of $\text{Fill}(s)$ such that $\phi(g)T^{(g)} \subset \text{Fill}(s)$. From this we will be able to conclude that $T^{(g)}$ and $\phi(g)T^{(g)}$ cannot be isotoped to be disjoint, as both subsurfaces have genus one, as does $\text{Fill}(s)$. This will give us that $\text{Fill}(s)$ and $\phi(g)\text{Fill}(s)$ cannot be isotoped to be disjoint.

We begin with the base case, $\mu(g) = 0$. Then as each $g_i = s^{\pm 1}$ we have that $\phi(g_i)\text{Fill}(s) = \text{Fill}(s)$, so $\phi(g)\text{Fill}(s) = \text{Fill}(s)$. We therefore have that $T^{(g)} = \text{Fill}(s)$ is our required subsurface.

Suppose that the statement is true for every $g \in A(\Gamma)$ such that $\mu(g) = k$. Now consider $g \in A(\Gamma)$ such that $\mu(g) = k + 1$. Let g_i be the first letter in $g_1 \cdots g_n$ such that $g_i \neq s^{\pm 1}$. Then $\mu(g_{i+1} \cdots g_n) = k$, so there exists a genus one subsurface $T^{(g_{i+1} \cdots g_n)}$ of $\text{Fill}(s)$ such that $\phi(g_{i+1} \cdots g_n)T^{(g_{i+1} \cdots g_n)} \subset \text{Fill}(s)$.

Let $Y = \text{Fill}(s) \backslash \bigcup_{s' \in V \backslash \{s\}} \text{Fill}(s')$, so the part of $\text{Fill}(s)$ that is fixed under the action of any other $\phi(s')$. Consider $Y \cap \phi(g_{i+1} \cdots g_n)T^{(g_{i+1} \cdots g_n)}$, and note that some

connected component must have genus one, as each intersection that we remove is a disc. Let T be this component. We then have that $\phi(g_i)T = T$, and hence $\phi(g_1 \cdots g_i)T \subset \text{Fill}(s)$, as for each $j < i$ we have that $g_j = s^{\pm 1}$.

Let $T^{(g)} = \phi(g_{i+1} \cdots g_n)^{-1}T$. As $T \subset \phi(g_{i+1} \cdots g_n)T^{(g_{i+1} \cdots g_n)}$ we get that $T^{(g)} \subset T^{(g_{i+1} \cdots g_n)} \subset \text{Fill}(s)$. We also have that $\phi(g)T^{(g)} = \phi(g_1 \cdots g_i)T \subset \text{Fill}(s)$, so this is our required genus one subsurface. This concludes our induction, and as a consequence the statement is proved. □

We can now combine these results to show that under certain circumstances we can quickly generate loxodromic elements in $A(\Gamma)$.

Proposition 5.2.13. *Suppose Γ is a connected graph that is not an isolated vertex. There exists $N = N(\Gamma) \in \mathbb{N}$ such that for every finite symmetric $U \subset A(\Gamma)$, where $\langle U \rangle$ is not contained non-trivially in a direct product and U is not conjugate into $A(\Gamma')$ for any induced subgraph $\Gamma' \subset \Gamma$, there exists $n \leqslant N$ such that U^n contains a loxodromic element on Γ^e.*

Proof. Let $S = S(\Gamma)$ be the surface constructed such that $\phi : A(\Gamma) \to MCG(S)$ is an injective homomorphism, in the sense of Theorem 5.2.1. As $\langle U \rangle$ is not contained non-trivially in a direct product, Lemma 5.2.8 tells us that the active subsurface of $\phi(\langle U \rangle)$ is connected.

Note that the construction of these surfaces in [25] means that S cannot be sporadic, so by Theorem 5.1.47 there exists a constant $N \in \mathbb{N}$, dependent only on S, such that for some $n \leqslant N$ the set $\phi(U)^n = \phi(U^n)$ contains an element $\phi(g)$ with the same active subsurface as $\phi(\langle U \rangle)$. By Remark 5.2.6 this subsurface is exactly $\text{Fill}(g)$. We want to show that $g \in U^n$ is a loxodromic element on Γ^e using Lemma 5.2.7.

Let $g = hg'h^{-1}$ where g' is cyclically reduced, and note that g is loxodromic if and only if g' is loxodromic. Recall that $\text{Fill}(g) = \phi(h)\text{Fill}(g')$. We have already noted that $\text{Fill}(g)$ is connected, so we therefore have that $\text{Fill}(g')$ is connected.

We now want to show that $\text{Fill}(g')$ has non-empty intersection with $\text{Fill}(s)$ for every $s \in V$. Let $s \in V$, and note that by Lemma 5.2.11 there exists $f = af'a^{-1} \in \langle U \rangle$ such that f' is cyclically reduced and s is in the support of f'. Therefore as $\text{Fill}(f) \subset \text{Fill}(g)$ we have that

$$\phi(a)\text{Fill}(s) \subset \phi(a)\text{Fill}(f') \subset \phi(h)\text{Fill}(g'),$$

so $\phi(h^{-1}a)\text{Fill}(s) \subset \text{Fill}(g')$.

Recall from Lemma 5.2.12 that $\text{Fill}(s)$ and $\phi(h^{-1}a)\text{Fill}(s)$ cannot be isotoped to be disjoint, so $\text{Fill}(s) \cap \text{Fill}(g') \neq \emptyset$. As Γ is not an isolated vertex, we can therefore

conclude that g' is a loxodromic element on Γ^e by Lemma 5.2.7, so $g \in U^n$ is also a loxodromic element. □

We can rephrase this result to give it in a similar format to Theorem 5.1.47.

Proposition 5.2.14. *Suppose Γ is a finite graph. For $U \subset A(\Gamma)$, let $V_U \subset V$ be minimal under inclusion such that U is conjugate into $A(\Gamma(V_U))$. There exists $N = N(\Gamma) \in \mathbb{N}$ such that for every finite symmetric $U \subset A(\Gamma)$, where $\Gamma(V_U)$ is connected and $\langle U \rangle$ is neither cyclic nor contained non-trivially in a direct product, there exists $n \leqslant N$ such that U^n contains a loxodromic element on $\Gamma(V_U)^e$.*

Proof. It was shown in Section 3 of [3] that for any $U \subset A(\Gamma)$ there exists such a $V_U \subset V$. Let $U \subset A(\Gamma)$ be finite and symmetric, and suppose that $\Gamma(V_U)$ is connected and $\langle U \rangle$ is neither cyclic nor contained non-trivially in a direct product. Let $g \in A(\Gamma)$ be such that $gUg^{-1} \subset A(\Gamma(V_U))$. As $\langle gUg^{-1} \rangle$ is not cyclic, we have that $\Gamma(V_U)$ is not an isolated vertex, and we also have that $\langle gUg^{-1} \rangle$ is not contained non-trivially in a direct product.

Hence we can apply Proposition 5.2.13 to show that there exists $N(\Gamma(V_U)) \in \mathbb{N}$ such that there exists $n \leqslant N(\Gamma(V_U))$ such that gU^ng^{-1} contains a loxodromic element on $\Gamma(V_U)^e$. The action of $\langle gUg^{-1} \rangle$ on $\Gamma(V_U)^e$ translates to an action of $\langle U \rangle$ on $\Gamma(V_U)^e$, so we can view this as U^n containing a loxodromic element.

Let $\mathcal{U}$ be the set of all possible finite symmetric $U \subset A(\Gamma)$ such that $\Gamma(V_U)$ is connected and $\langle U \rangle$ is neither cyclic nor contained non-trivially in a direct product. As Γ is a finite graph, there are only finitely many possibilities for V_U. We therefore let $N = \max\{N(\Gamma(V_U)) : U \in \mathcal{U}\}$, and the result follows. □

As mentioned in the introduction, the ability to quickly generate these loxodromic elements has an application to $\xi(G) = \{\omega(G, S) \ : \ S \text{ is a finite generating set of } G\}$, the set of exponential growth rates of G with respect to its finite generating sets. In particular, the fact that loxodromic elements can be used to link the growth of a set with its size is a key step in the proof of the following theorem.

Theorem 5.2.15. [37] *Let G be a group with an acylindrical and non-elementary action on a hyperbolic graph X. Suppose that there exists a constant M such that for any finite symmetric generating set S of G, the product set S^M contains a loxodromic element on X. Assume that G is equationally Noetherian. Then $\xi(G)$ is a well-ordered set.*

We can combine this with Proposition 5.2.14 to obtain the following.

Theorem 5.2.16. *Let Γ be a finite graph. Suppose that $G \leqslant A(\Gamma)$ is finitely generated, and is not contained non-trivially in a direct product where the projection of G to more than one factor has exponential growth. Then $\xi(G)$ is a well-ordered set.*

Proof. We begin by noting that there is a bound on the number of factors of any direct product subgroup of a right-angled Artin group, as a consequence of Theorem 5.1.43 and the fact that right-angled Artin groups are subgroups of mapping class groups [25].

If G does not have exponential growth then $\xi(G) = \{1\}$, which is trivially well-ordered. We therefore assume that G is not cyclic, and by Lemma 3.1.15 we also assume that G is not contained non-trivially in a direct product where none of the factors have exponential growth.

Now suppose G is contained non-trivially in a direct product, and the projection of G to exactly one factor has exponential growth. Then by Lemma 3.1.15 we have that the set of growth rates of G is the same as the set of growth rates of the projection of G to this factor. We can therefore assume that G is not contained non-trivially in a direct product.

Note that G is a subgroup of a right-angled Artin group, so is linear [48], and therefore is equationally Noetherian [7]. Let $V_G \subset V$ be minimal under inclusion such that G is conjugate into $A(\Gamma(V_G))$. If $\Gamma(V_G)$ is connected, then we conclude using Proposition 5.2.14 and Theorem 5.2.15. If $\Gamma(V_G)$ is not connected, we instead use its action on the associated Bass-Serre tree (see Example 1.7 in [37]), and again conclude using Theorem 5.2.15. □

5.3 Product set growth in right-angled Artin groups

We can now use the results of the previous section to answer Question 5 for right-angled Artin groups. For a given finite symmetric $U \subset A(\Gamma)$, we will use either Proposition 5.2.14 or the action of a free product on its Bass-Serre tree to quickly generate a loxodromic element in the action of $\langle U \rangle$ on one of a finite collection of quasi-trees, which will allow us to apply Corollary 4.2.21.

This will not give us our final result, as we will then be able to break down the direct product case created by Lemma 5.2.8 using some of our observations from Section 3.6, however it does get us very close to the best possible result.

Remark 5.3.1. In the following we assume that $\langle U \rangle$ is not $\mathbb{Z}$, rather than not virtually $\mathbb{Z}$. This is possible as $A(\Gamma)$ is torsion-free, so as any torsion-free virtually free group is free [82] we have that any virtually $\mathbb{Z}$ subgroup of $A(\Gamma)$ is in fact isomorphic to $\mathbb{Z}$.

Proposition 5.3.2. *Let Γ be a finite graph. There exist constants $\alpha, \beta > 0$ such that for every finite symmetric $U \subset A(\Gamma)$ at least one of the following must hold:*

(1) $\langle U \rangle$ is isomorphic to $\mathbb{Z}$

(2) $\langle U \rangle$ is contained non-trivially in a direct product.

(3) $|U^n| \geqslant (\alpha |U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Proof. Note first that if $\langle U \rangle$ is finite, then it is the trivial group, and so will satisfy the inequality $|U^n| \geqslant (\alpha |U|)^{\beta n}$ as long as $\alpha \leqslant 1$. We will therefore assume that $\langle U \rangle$ is infinite.

We want to show that there exists $N \in \mathbb{N}$ such that for every finite symmetric $U \subset A(\Gamma)$, where $\langle U \rangle$ is not contained non-trivially in a direct product and is not isomorphic to $\mathbb{Z}$, there exists $n \leqslant N$ such that U^n contains a loxodromic element for an acylindrical action of $\langle U \rangle$ on one of a finite collection of quasi-trees.

Let U be such a set, and let $V_U \subset V$ be minimal under inclusion such that U is conjugate into $A(\Gamma(V_U))$. We will technically be proving our result for some conjugate $gUg^{-1} \subset A(\Gamma(V_U))$, however as noted previously we can see any action of $\langle gUg^{-1} \rangle$ as being an action of $\langle U \rangle$. It was also noted in Chapter 3 that product set growth is preserved by conjugation, so proving our result for some gUg^{-1} will be sufficient. There are now only two cases for us to consider.

The first case is that $\Gamma(V_U)$ is connected, which means that we can apply Proposition 5.2.14. Let $N' = N'(\Gamma) \in \mathbb{N}$ be the constant that we get from this proposition. Note that the fact that there are only finitely many such $V_U \subset V$ means that there are also only finitely many quasi-trees $\Gamma(V_U)^e$, and finitely many associated acylindrical actions.

The second case is that $\Gamma(V_U)$ is disconnected. Write $\Gamma(V_U) = \Gamma_1 \sqcup \Gamma_2$, and note that Γ_1 and Γ_2 are both induced subgraphs of Γ. Then $A(\Gamma(V_U)) = A(\Gamma_1) * A(\Gamma_2)$, and we recall from Lemma 4.1.7 that any free product acts acylindrically on its associated Bass-Serre tree.

If $\langle U \rangle$ fixes a vertex in the tree, then $\langle U \rangle$ is contained in the stabiliser of that vertex. This implies U is conjugate into $A(\Gamma_1)$ or $A(\Gamma_2)$, which contradicts the minimality of V_U. Therefore $\langle U \rangle$ does not fix a vertex, so U^2 must contain a loxodromic

element in this action, as this is an action on a tree [78, p. 64]. Note that as Γ is a finite graph, there are only finitely many possible choices for $A(\Gamma_1) * A(\Gamma_2) \leqslant A(\Gamma)$, so only finitely many acylindrical actions to consider on the associated Bass-Serre trees.

We now combine the cases by letting $U \subset A(\Gamma)$ be any finite symmetric subset such that $\langle U \rangle$ is not contained non-trivially in a direct product and is not isomorphic to $\mathbb{Z}$, and letting $N = \max\{2, N'\}$. We can conclude that there exists $n \leqslant N$ such that U^n contains a loxodromic element for some acylindrical action on one of a finite collection of quasi-trees. Therefore by Corollary 4.2.21 and Remark 4.2.22 there exist constants $\alpha, \beta > 0$ such that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for all $n \in \mathbb{N}$. □

As noted in Section 3.6, it is unsurprising that the presence of direct products is a problem when trying to link the size of U^n with the size of U. However we also showed in that section that we can still sometimes say something in the case of direct products, so long as all of the factors exhibit uniform product set growth. We can use this observation to give an improved version of Proposition 5.3.2.

One potential difficulty with proving uniform product set growth using this method is that the constants that we obtain depend on the number of factors in the direct product. Fortunately for us, it is a well known fact that the maximal rank of a free abelian subgroup of a right-angled Artin group $A(\Gamma)$ is equal to the number of vertices in a maximal complete subgraph of Γ, see for example [23], [58], or [60]. As mentioned in the proof of Theorem 5.2.16, the maximal rank being bounded is also a trivial consequence of Theorem 5.1.43 and the fact that right-angled Artin groups are subgroups of mapping class groups. On the other hand, it is useful to have a specific bound linked directly to the properties of Γ.

For the sake of completeness we will give a short proof of this fact here, for which we need the Centraliser Theorem.

Theorem 5.3.3 (Centraliser Theorem)**.** [79] *Let Γ be a finite graph, let $g \in A(\Gamma)$, and write $g = hg'h^{-1}$ where g' is cyclically reduced. Let V' be the support of g', and let $\Gamma_1, \ldots, \Gamma_k$ be the connected components of $\Gamma(V')^c$. Then we can write $g' = g_1^{n_1} \cdots g_k^{n_k}$, where the support of g_i is contained in Γ_i, each $n_i \in \mathbb{N}$ is maximal, and the centraliser of g is*

$$\mathrm{C}(g) = h(\langle g_1 \rangle \times \cdots \times \langle g_k \rangle \times \langle W \rangle)h^{-1},$$

where W is the set of vertices in $V \backslash V'$ that are adjacent to every $v \in V'$.

Lemma 5.3.4. *Let Γ be a finite graph, and let $N \in \mathbb{N}$ be the number of vertices in a maximal complete subgraph of Γ. If $\mathbb{Z}^n \leqslant A(\Gamma)$ then $n \leqslant N$.*

Proof. We proceed by induction on the number of vertices in Γ. If $|V| = 1$ then $A(\Gamma) \cong \mathbb{Z}$, so we are done. Now take arbitrary Γ, and suppose that this statement is true for all graphs on fewer vertices.

Suppose $\mathbb{Z}^n \cong \langle a_1 \rangle \times \cdots \times \langle a_n \rangle \leqslant A(\Gamma)$, and let $a_1 = hg'h^{-1}$ where g' is cyclically reduced. Let V' be the support of g', and denote the connected components of $\Gamma(V')^c$ by $\Gamma_1, \ldots, \Gamma_k$. We then let $g_1, \ldots, g_k$ and W be as in the Centraliser Theorem. We can see that

$$\mathbb{Z}^n \cong \langle a_1 \rangle \times \cdots \times \langle a_n \rangle \leqslant \mathrm{C}(a_1) = h(\langle g_1 \rangle \times \cdots \times \langle g_k \rangle \times \langle W \rangle)h^{-1} \cong \mathbb{Z}^k \times A(\Gamma(W)).$$

As $|W| < |V|$, we have by our induction hypothesis that the maximum rank of any free abelian subgroup of $A(\Gamma(W))$ is the number of vertices in a maximal complete subgraph of $\Gamma(W)$, which we will say is m. We therefore get that $n \leqslant k + m$.

Let $v_1, \ldots, v_k$ be vertices in $\Gamma_1, \ldots, \Gamma_k$ respectively, then for $i \neq j$ we get that v_i and v_j are adjacent in $\Gamma(V')$. Hence $v_1, \ldots, v_k$ induce a complete subgraph in $\Gamma(V')$, and so induce a complete subgraph in Γ. Let $w_1, \ldots, w_m \in W$ be such that they induce a complete subgraph in $\Gamma(W)$, and so also in Γ. Each v_i is adjacent to each w_j, so $v_1, \ldots, v_k, w_1, \ldots, w_m$ induce a complete subgraph in Γ. Therefore $k + m \leqslant N$, where N is the number of vertices in a maximal complete subgraph of Γ, and hence $n \leqslant N$. □

The following lemma is a natural consequence of this fact.

Lemma 5.3.5. *Let Γ be a finite graph, and let $N \in \mathbb{N}$ be the number of vertices in the maximal complete subgraph of Γ. If $G_1 \times \cdots \times G_n \leqslant A(\Gamma)$ with each G_i non-trivial, then $n \leqslant N$.*

Proof. As $A(\Gamma)$ is torsion free, if we take arbitrary $1 \neq g_i \in G_i$ we get that $\mathbb{Z}^n \cong \langle g_1 \rangle \times \cdots \times \langle g_n \rangle \leqslant G_1 \times \cdots \times G_n \leqslant A(\Gamma)$, so we must have that $n \leqslant N$. □

We can now combine this with Corollary 3.6.14 to improve Proposition 5.3.2.

Theorem 5.3.6. *Let Γ be a finite graph. There exist constants $\alpha, \beta > 0$ such that for every finite symmetric $U \subset A(\Gamma)$ at least one of the following must hold:*

(1) The centre of $\langle U \rangle$ is non-trivial.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Proof. Let $U \subset A(\Gamma)$ be finite and symmetric, and suppose that $\langle U \rangle$ has trivial centre. In particular, $\langle U \rangle$ cannot be $\mathbb{Z}$. Note that, by the same reasoning as in Corollary 3.6.3, if $\langle U \rangle$ is a subgroup of a direct product of the form $H \times \mathbb{Z}$, then the projection of $\langle U \rangle$ to H is injective, as otherwise $\langle U \rangle$ would have infinite centre.

If $\langle U \rangle$ is not contained non-trivially in a direct product, then we get that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, with the α and β coming from Proposition 5.3.2. Now suppose that $\langle U \rangle$ is contained non-trivially in a direct product $G_1 \times \cdots \times G_m \leqslant A(\Gamma)$, and let U_i be the projection of U to G_i.

Suppose that some $\langle U_i \rangle \leqslant G_i$ is contained non-trivially in a direct product $G_{i,1} \times G_{i,2} \leqslant G_i$. Then $\langle U \rangle$ is contained in $G_1 \times \cdots \times G_{i-1} \times G_{i,1} \times G_{i,2} \times G_{i+1} \times \cdots \times G_m$. Lemma 5.3.5 tells us that there exists $N \in \mathbb{N}$ that bounds the number of factors in a direct product in $A(\Gamma)$, so we can assume that each $\langle U_i \rangle$ is not contained non-trivially in a direct product, and that $m \leqslant N$.

Suppose that some $\langle U_i \rangle$ is $\mathbb{Z}$. Then, by our assumption, we must have that the projection of $\langle U \rangle$ to $G_1 \times \cdots \times G_{i-1} \times G_{i+1} \times \cdots \times G_m$ is injective, so $\langle U \rangle$ can instead be viewed as a subgroup of this direct product. We can therefore assume that no $\langle U_i \rangle$ is $\mathbb{Z}$, without loss of generality. We also note that each U_i is finite and symmetric. Hence Proposition 5.3.2 tells us that $|U_i^n| \geqslant (\alpha|U_i|)^{\beta n}$ for every $n \in \mathbb{N}$ and $i \in \{1, \ldots, m\}$.

We can now apply Corollary 3.6.14 to get that

$$|U^n| \geqslant (\alpha|U|^{\frac{1}{m}})^{\beta n} \geqslant (\alpha|U|^{\frac{1}{N}})^{\beta n} = (\alpha^N|U|)^{\frac{\beta n}{N}}.$$

Let $\alpha' = \min\{\alpha, \alpha^N\}$ and $\beta' = \frac{\beta}{N}$. Combining the two cases, we have shown that for every finite symmetric $U \subset A(\Gamma)$, where $\langle U \rangle$ has trivial centre, we have that $|U^n| \geqslant (\alpha'|U|)^{\beta' n}$ for all $n \in \mathbb{N}$. □

We note here that, given that Proposition 3.6.2 tells us that any group with an infinite order element in its centre cannot have uniform product set growth, and the fact that right-angled Artin groups are torsion free, this is the best result we can hope for in this case. We have therefore answered Question 5 for symmetric subsets of right-angled Artin groups. The question as to whether the same result holds for non-symmetric sets remains open.

A final remark on right-angled Artin groups is that in the proof of Theorem 5.3.6 we use that any subgroup with trivial centre has the property that if it is a subgroup of a direct product of the form $H \times \mathbb{Z}$, then its projection to H is injective. We could have used this as initial our assumption on our subgroups in Theorem 5.3.6 to get the same result, so as Theorem 5.3.6 is a dichotomy we re-obtain the following fact, which can also be deduced from the Centraliser Theorem.

Corollary 5.3.7. *Let Γ be a finite graph, and let $G \leqslant A(\Gamma)$. Then G has non-trivial centre if and only if it is contained in a group of the form $H \times \mathbb{Z}$, where the projection of G to H is not injective.*

5.4 Product set growth in mapping class groups

We now move on to the more general case of subgroups of mapping class groups. We begin by considering only the pure subgroups, which allows us to get a product set growth result that we can then extend to all subgroups.

The general structure of the proof for pure subgroups is the same as in the previous section. That is, we consider our subgroup as a subgroup of a direct product. With the exception of certain cases, we show that each of the factor groups in the direct product satisfy uniform product set growth. We can then apply Corollary 3.6.14 to get growth for the whole subgroup.

Theorem 5.4.1. *Let $G \leqslant MCG(S)$ be a pure subgroup of a mapping class group. There exist $\alpha, \beta > 0$ such that for every finite symmetric $U \subset G$ at least one of the following must hold:*

(1) The centre of $\langle U \rangle$ is non-trivial.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Proof. Suppose $U \subset G$ is symmetric and finite, and that the centre $Z(\langle U \rangle)$ is trivial. Suppose that $\langle U \rangle$ is finite. As pure subgroups are torsion-free, this means that $\langle U \rangle$ is trivial. Therefore so long as we ensure that $\alpha \leqslant 1$, we will have that U satisfies $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

We can now assume that $\langle U \rangle$ is infinite. Let $\{\xi_1, \dots, \xi_l\}$ be the (possibly empty) set of curves in the boundary of S, and recall that T_{ξ_i} is the Dehn twist about ξ_i. Let S' be the surface obtained from S by capping each boundary component by a once punctured disc. By Proposition 5.1.48 we have a natural homomorphism $\psi : \langle U \rangle \to MCG(S')$, with kernel $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle \cap \langle U \rangle$.

Recall from Corollary 5.1.37 that the group $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle \cap \langle U \rangle$ is central in $\langle U \rangle$, and therefore by our initial assumption $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle \cap \langle U \rangle$ is trivial. This means that ψ is injective, and we can therefore suppose without loss of generality that S has no boundary.

Now let $\{\gamma_1, \dots, \gamma_k\}$ be the (possibly empty) set of curves in the canonical reduction multicurve for $\langle U \rangle$. Let $\mathcal{S}$ be the set of connected components of $S \setminus \bigcup_{i=1}^{k} \gamma_i$. By

Theorem 5.1.41 we have a natural homomorphism $\varphi : \langle U \rangle \to \Pi_{\Sigma \in \mathcal{S}} MCG(\Sigma)$, with kernel $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap \langle U \rangle$.

Recall from Corollary 5.1.37 that the group $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap \langle U \rangle$ is central in $\langle U \rangle$, and therefore by our initial assumption $\langle T_{\gamma_1}, \ldots, T_{\gamma_k} \rangle \cap \langle U \rangle$ is trivial. This means that φ is injective, and it therefore makes sense to talk about the projection of $\langle U \rangle$ to a factor of $\Pi_{\Sigma \in \mathcal{S}} MCG(\Sigma)$.

Consider the projection of $\langle U \rangle$ to such a factor $MCG(\Sigma)$. By Theorem 5.1.41, this projection is either trivial, or contains a pseudo-Anosov on Σ. If the projection is trivial, we can ignore this factor, and consider $\langle U \rangle$ as a subgroup of the direct product of the remaining factors. If the projection is $\mathbb{Z}$, then we can similarly ignore this factor, as the proof of Corollary 3.6.3 tells us that either the projection to the remaining factors is injective, or $\langle U \rangle$ has infinite centre.

Let $\mathcal{S}'$ be the surfaces Σ in $\mathcal{S}$ such that the projection of $\langle U \rangle$ to $MCG(\Sigma)$ is neither trivial nor $\mathbb{Z}$. The above reasoning tells us that we can see $\langle U \rangle$ as a subgroup of $\Pi_{\Sigma \in \mathcal{S}'} MCG(\Sigma)$. Let U_Σ be the projection of U to $MCG(\Sigma)$ for some $\Sigma \in \mathcal{S}'$.

We first suppose that Σ is non-sporadic. As $\langle U_\Sigma \rangle$ contains a pseudo-Anosov on Σ, Theorem 5.1.32 tells us that there is some $N \in \mathbb{N}$, dependent only on S, such that for some $k \leqslant N$ we have that U_Σ^k contains a pseudo-Anosov on Σ. Recall that the action of $\langle U_\Sigma \rangle$ on the curve complex associated to Σ is acylindrical, and this curve complex is hyperbolic. Therefore, as $\langle U_\Sigma \rangle$ is not $\mathbb{Z}$, Corollary 4.3.2 tells us that there exist $\alpha_\Sigma, \beta_\Sigma > 0$ determined by Σ such that $|U_\Sigma^n| \geqslant (\alpha_\Sigma |U_\Sigma|)^{\beta_\Sigma n}$ for every $n \in \mathbb{N}$.

Alternatively, if Σ is sporadic, then $MCG(\Sigma)$ is hyperbolic. Therefore, as $\langle U_\Sigma \rangle$ is not $\mathbb{Z}$, Theorem 3.4.9 tells us that again there exist $\alpha_\Sigma, \beta_\Sigma > 0$ determined by Σ such that $|U_\Sigma^n| \geqslant (\alpha_\Sigma |U_\Sigma|)^{\beta_\Sigma n}$ for every $n \in \mathbb{N}$.

There are only finitely many homeomorphism classes of essential subsurfaces of S, that is the subsurfaces with boundary components that are either in ∂S or are essential in S, so we can take $\alpha = \inf \alpha_\Sigma$ and $\beta = \inf \beta_\Sigma$ over all possible subsurfaces, and we will get that $\alpha, \beta > 0$.

Now let g be the genus of S, c be the number of connected components, and recall that if our original surface had b boundary components and p punctures, then by our use of Proposition 5.1.48 the surface that we are now considering has $b+p$ punctures. By Theorem 5.1.43, the maximal rank of a free abelian subgroup of $MCG(S)$ is $3g + b + p - 3c$. Therefore as each factor of $\Pi_{\Sigma \in \mathcal{S}'} MCG(\Sigma)$ contains an infinite order element, we must have that $|\mathcal{S}'| \leqslant 3g + b + p - 3c$. Let

$$\alpha' = \alpha^{3g+b+p-3c}$$

and let

$$\beta' = \frac{\beta}{3g + b + p - 3c}.$$

Then by Corollary 3.6.14 we have that $|U^n| \geqslant (\alpha'|U|)^{\beta' n}$ for every $n \in \mathbb{N}$.

Combining this with the case that $\langle U \rangle$ is finite, we see that if we let $\alpha'' = \min\{\alpha', 1\}$, then any set U that satisfies our initial hypotheses will also satisfy $|U^n| \geqslant (\alpha''|U|)^{\beta' n}$ for every $n \in \mathbb{N}$. □

By Proposition 3.6.2, this is a dichotomy of subgroups, and so answers Question 5 for the pure subgroups of mapping class groups.

Remark 5.4.2. As right-angled Artin groups embed as pure subgroups of mapping class groups, Theorem 5.4.1 reproves Theorem 5.3.6.

Theorem 5.1.39 tells us that every subgroup of a mapping class group has a pure subgroup of finite index. We can therefore combine Theorem 5.4.1 with Proposition 3.6.25 to get the following.

Corollary 5.4.3. *Let $MCG(S)$ be a mapping class group. There exist $\alpha, \beta > 0$ such that for every finite symmetric $U \subset MCG(S)$ at least one of the following must hold:*

(1) $\langle U \rangle$ has a finite index pure subgroup with non-trivial centre.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

Proof. Let $U \subset MCG(S)$ be finite and symmetric. Let P be a fixed finite index pure subgroup of $MCG(S)$, where $[MCG(S) : P] = d$. Then $H = P \cap \langle U \rangle$ is a finite index pure subgroup of $\langle U \rangle$. Suppose that H has trivial centre.

Let V be a finite symmetric generating set of H. Then by Theorem 5.4.1 there exist uniform $\alpha, \beta > 0$ such that $|V^n| \geqslant (\alpha|V|)^{\beta n}$ for every $n \in \mathbb{N}$. Let $m = 2d - 1$, then Proposition 3.6.25 tells us that

$$|U^n| \geqslant \left(\frac{\alpha}{2^{\frac{m}{\beta}} d}|U|\right)^{\frac{\beta n}{m}}$$

for every $n \in \mathbb{N}$. □

It does not immediately follow that this is a dichotomy of subgroups, as we have not yet shown that the groups in the first case do not have uniform product set growth. Although Proposition 3.6.25 tells us that uniform product set growth is inherited by finite index supergroups, we do not yet know if this is true in general for finite index subgroups. We can however still achieve a dichotomy in the particular case of mapping class groups, as shown by the following proposition.

Proposition 5.4.4. *Let $MCG(S)$ be a mapping class group, and let G be a finitely generated subgroup of $MCG(S)$. Suppose that G has a finite index pure subgroup with non-trivial centre. Then for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$.*

Proof. First, let $\{\xi_1, \dots, \xi_l\}$ be the (possibly empty) set of curves in the boundary of S. Recall from the proof of Theorem 5.4.1 that either $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle \cap G$ is non-trivial, in which case as this group is central in G we have that G does not have uniform product set growth by Proposition 3.6.2, or $\langle T_{\xi_1}, \dots, T_{\xi_l} \rangle \cap G$ is trivial, and we can assume without loss of generality that S has empty boundary. Here it does not matter that G might not be a pure subgroup, as no element in G permutes the boundary curves.

Now let P be a finite index pure subgroup of G with non-trivial centre. Let $\{\gamma_1, \dots, \gamma_k\}$ be the (possibly empty) set of curves in the canonical reduction multicurve for P. Recall from Corollary 5.1.37 that the group $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P$ is central in P. Suppose that $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P$ is non-trivial, and therefore infinite.

We can assume that $P \trianglelefteq G$ by Theorem 5.1.39, and by Remark 5.1.42 we have that $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap G \trianglelefteq G$, so $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P \trianglelefteq G$. In particular, recall that G acts on $\{\gamma_1, \dots, \gamma_k\}$ by permuting the curves, and so $g(\gamma_i) = \gamma_j$ for every $g \in G$. Hence $\{gT_{\gamma_i}g^{-1} : g \in G\} \subset \{T_{\gamma_1}, \dots, T_{\gamma_k}\}$, which tells us that the orbit of any element of $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P$ under the conjugation action of G is finite. As $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P$ is an abelian group, it does not have exponential growth, so we can apply Proposition 3.6.1 to show that G does not have uniform product set growth.

Now suppose that $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap P$ is trivial. This has finite index in the group $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap G$, and so as $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle$ is torsion-free we must also have that $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle \cap G$ is trivial. Recall from Proposition 5.1.35 that there is a natural homomorphism $\varphi : MCG(S, \{\gamma_1, \dots, \gamma_n\}) \to MCG(S \setminus \bigcup_{i=1}^k \gamma_i)$, with kernel $\langle T_{\gamma_1}, \dots, T_{\gamma_k} \rangle$, where $MCG(S, \{\gamma_1, \dots, \gamma_n\})$ is the subgroup of $MCG(S)$ that fixes the set of curves $\{\gamma_1, \dots, \gamma_n\}$. We have that the restriction of this homomorphism to $\varphi : G \to MCG(S \setminus \bigcup_{i=1}^k \gamma_i)$ is injective, so we can view G as a subgroup of $MCG(S \setminus \bigcup_{i=1}^k \gamma_i)$.

Let $\mathcal{S}$ be the collection of connected components of $S \setminus \bigcup_{i=1}^k \gamma_i$, and recall from Theorem 5.1.41 that we can view P as a subgroup of $\Pi_{\Sigma \in \mathcal{S}} MCG(\Sigma)$, where the projection of P to each factor is either trivial or contains a pseudo-Anosov on Σ. From the proof of Theorem 5.4.1 we know that if none of these projections are $\mathbb{Z}$, then P would have uniform product set growth. However this would contradict the

fact that P has non-trivial centre, so we know that at least one of these projections must be $\mathbb{Z}$.

Let $\mathcal{S}'$ be the connected surfaces $\Sigma \in \mathcal{S}$ such that the projection of P to $MCG(\Sigma)$ is neither trivial nor $\mathbb{Z}$. Note that an element $g \in G$ may permute the surfaces in $\mathcal{S}$, however g will stabilise the sets of subsurfaces $\mathcal{S}'$ and $\mathcal{S}\backslash\mathcal{S}'$. This is because for any $g \in G$ we have that $P = g^{-1}Pg$, and if we let P_Σ be the projection of P to $MCG(\Sigma)$ for $\Sigma \in \mathcal{S}$, then $P_\Sigma = (g^{-1}Pg)_\Sigma = P_{g(\Sigma)}$, noting that the projection of a pure mapping class to some $MCG(\Sigma)$ is just the restriction of that mapping class to Σ.

We can therefore write $g \in G$ as $g = (a_g, b_g)$, where a_g is the restriction of g to $\bigcup_{\Sigma\in\mathcal{S}'} \Sigma$, and b_g is the restriction of g to $\bigcup_{\Sigma\in\mathcal{S}\backslash\mathcal{S}'} \Sigma$. Let $A = \langle a_g \mid g \in G\rangle$ and $B = \langle b_g| \ g \in G\rangle$, and note that elements in these two groups commute with each other, as they are supported by disjoint subsurfaces. Therefore $G \leqslant A \times B$.

It follows that P is also a subgroup of $A \times B$. Let B' be the projection of P to B. As we know that the projection of P to $MCG(\Sigma)$ for $\Sigma \in \mathcal{S}\backslash\mathcal{S}'$ is trivial or $\mathbb{Z}$, and that at least one projection is $\mathbb{Z}$, we get that B' is an infinite subgroup of a free abelian group, and so is itself free abelian.

We would like to show that B' is a finite index subgroup of B. Let $u_1, \dots, u_d$ be coset representatives of P in G, so $u_i = (a_i, b_i)$. This means that for every $g \in G$ there exists $g' \in P$ and u_i such that $(a_i a_{g'}, b_i b_{g'}) = (a_g, b_g)$. It follows that $b_1, \dots, b_d$ are coset representatives of B' in B, so we have that B is virtually free abelian.

Suppose that the projection of G to A is finite-to-1, then as $P \leqslant G$ the projection of P to A is finite-to-1. We can view P as a subgroup of $A \times B'$, and as B' is torsion-free, Corollary 3.6.11 tells us that the projection of P to A is in fact injective. By our construction of A, this would mean that P is isomorphic to a pure subgroup of $\Pi_{\Sigma\in\mathcal{S}'}\Sigma$, where the projection of P to each factor contains a pseudo-Anosov, and is not $\mathbb{Z}$.

Following the proof of Theorem 5.4.1, this would mean that P has uniform product set growth, and therefore cannot have non-trivial centre. This is a contradiction, so the projection of G to A is infinite-to-1. As B is virtually free abelian, it does not have exponential growth, and so by Proposition 3.6.9 we have that G does not have uniform product set growth. In other words, for any constants $\alpha, \beta > 0$ we can find a finite generating set U of G such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$. □

6

Further questions

Never put off till tomorrow what you can do the day after tomorrow.

– Attributed to Mark Twain

Here we will briefly cover some questions that arise from the work in the preceding chapters. We begin with the topic of quasi-trees. Recall that in Section 2.4 it was shown that quasi-trees satisfy a uniform version of tree approximation. On the other hand, it was also shown that this does not hold if we ask that the approximating tree is a subtree of the space (see Example 2.4.8). Question 2 asked if there are any properties that characterise the quasi-trees which do admit such subtrees.

Question 6 (Question 2)**.** *Does there exist a characterisation of the quasi-trees for which there exists a constant $C \geqslant 0$, such that for any set of finite points $Y \subset X$ we can construct a subtree $T \subset X$ such that $Y \subset T$, and for every $x, y \in T$ we have that $d_T(x, y) \leqslant d(x, y) + C$?*

We note here that all of the work in this thesis has been for geodesic spaces. In particular, quasi-trees are defined as being the geodesic spaces that are quasi-isometric to trees. In Section 2.3 we showed that, assuming our space is geodesic, being quasi-isometric to a tree is the same as being $(1, C)$-quasi-isometric to some tree.

Recalling that $\mathbb{R}$-trees are simply the geodesic 0-hyperbolic spaces, the analogue to this for non-geodesic spaces would be to look at the spaces that are quasi-isometric to some 0-hyperbolic space. It is reasonable to wonder if a similar result could hold in this setting.

Question 7. *Let X be a metric space that is (L, C)-quasi-isometric to a 0-hyperbolic space for some $L \geqslant 1, C \geqslant 0$. Is X $(1, C')$-quasi-isometric to a 0-hyperbolic space for some $C' \geqslant 0$?*

Our proof of this statement for geodesic spaces relied heavily on Manning's bottleneck criterion giving us an equivalent, easier to work with definition of quasi-trees. To answer the above question we may therefore need a similar kind of criterion for spaces that are quasi-isometric to 0-hyperbolic spaces.

Question 8. *Does there exist an alternative characterisation of the metric spaces that are quasi-isometric to 0-hyperbolic spaces?*

Continuing the theme of considering non-geodesic spaces, it was shown in [13] by Bonk and Schramm that any non-geodesic δ-hyperbolic space isometrically embeds into a complete geodesic δ-hyperbolic space, and so it naturally follows that 0-hyperbolic spaces embed into $\mathbb{R}$-trees. It would be interesting to see if the same thing happens in the quasi-tree case, now that we have an equivalent definition of quasi-trees which is analogous to the Gromov product definition of a hyperbolic space.

Question 9. *Let X be a metric space. Suppose there exists $A > 0$ such that for every $x_0, x_1, \dots, x_n \in X$ we have that*

$$(x_1, x_n)_{x_0} \geqslant \min_{1 \leqslant i \leqslant n-1} (x_i, x_{i+1})_{x_0} - A.$$

Does X embed isometrically into a complete geodesic quasi-tree?

We now move on to questions related to product set growth. In Section 5.3 we found a dichotomy of subgroups in right-angled Artin groups, where the property in question was whether the finite symmetric generating sets of the subgroup satisfy uniform product set growth. The necessity of symmetry did not come from the result used to achieve the growth, as in Section 4.2 we showed that quickly generating loxodromic elements in an acylindrical action on a quasi-tree is sufficient regardless of symmetry, unlike the result for hyperbolic spaces in Section 4.3.

Instead, the requirement for symmetry came from the method we used in Section 5.2 to generate those loxodromic elements on the extension graph. More specifically, it came from the fact that we used Mangahas' result regarding the quick generation of pseudo-Anosovs in mapping class groups (see Theorem 5.1.47), which requires the sets to be symmetric. It may however be possible to prove a similar result in right-angled Artin groups without appealing to this result, perhaps in a more direct and combinatorial way. In particular, it may also be possible to find loxodromic elements without the set in question needing to be symmetric.

Question 10. *Suppose Γ is a finite graph. For $U \subset A(\Gamma)$, let $V_U \subset V$ be minimal under inclusion such that U is conjugate into $A(\Gamma(V_U))$. Does there exist $N = N(\Gamma) \in \mathbb{N}$ such that for every finite $U \subset A(\Gamma)$, where $\Gamma(V_U)$ is connected and $\langle U \rangle$ is neither cyclic nor contained non-trivially in a direct product, there exists $n \leqslant N$ such that U^n contains a loxodromic element on $\Gamma(V_U)^e$?*

If this does turn out to be the case, then we would automatically get a stronger result for growth in right-angled Artin groups, and a more complete answer to Question 5.

Question 11. *Let Γ be a finite graph. Do there exist constants $\alpha, \beta > 0$ such that for every finite $U \subset A(\Gamma)$, at least one of the following must hold?*

(1) The centre of $\langle U \rangle$ is non-trivial.

(2) $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$.

As mentioned above, in Section 4.2 we showed that quickly generating loxodromic elements in an acylindrical action on a quasi-tree is sufficient for product set growth, regardless of symmetry. However in Section 4.3 the version for hyperbolic spaces does require symmetry. Continuing the theme of the previous two questions, we can again ask if this is truly necessary in this case.

Question 12. *Let G be a group acting acylindrically on a hyperbolic space. Do there exist $\alpha, \beta > 0$ such that for every finite $U \subset G$ such that U contains a loxodromic element, and $\langle U \rangle$ is not virtually $\mathbb{Z}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$?*

Recall from Example 3.4.7 that the best possible exponent we could usually expect to get for these type of results is around $\beta = \frac{1}{2}$. The constants that we found for subgroups of mapping class groups are not close to this, as finding loxodromic elements, projecting to factors in direct products, and passing to finite index supergroups all have an effect on our choice of β. It is currently unknown what the optimal constants are in this case.

Question 13. *For a group G, what are the largest constants $\alpha, \beta > 0$ such that there exists a class of subgroups $\mathcal{H}$ such that for every finite (symmetric) $U \subset G$ where $\langle U \rangle \notin \mathcal{H}$, we have that $|U^n| \geqslant (\alpha|U|)^{\beta n}$ for every $n \in \mathbb{N}$, and for every $H \in \mathcal{H}$ there exists a finite (symmetric) generating set U such that $|U^n| < (\alpha|U|)^{\beta n}$ for some $n \in \mathbb{N}$?*

At the end of Section 5.2 we used a result of Fujiwara [37] to show that the set of growth rates for a finitely generated subgroup of a right-angled Artin group is well-ordered, assuming that this subgroup is not contained non-trivially in a direct product where its projection to more than one factor has exponential growth. This is because we know that we have this well-ordering for exponential factors by combining Proposition 5.2.14 and Theorem 5.2.15, and Lemma 3.1.15 tells us that we can ignore any factor without exponential growth. We are not yet able however to deal with the case of two factors with exponential growth.

Question 14. *Let $H \leqslant G_1 \times \cdots \times G_m$ be a finitely generated group, and let H_i be the projection of H to G_i. If every H_i has exponential growth, and the set of growth rates of H_i is well-ordered for every $1 \leqslant i \leqslant m$, then is the set of growth rates of H well-ordered?*

If the answer to this question turns out to be yes, then this would extend the conclusion of Theorem 5.2.16 to every subgroup of a right-angled Artin group.

The following question was asked by Anthony Genevois, who in [39] constructed an analogue of the extension graph for the more general case of graph products. This analogue is also a quasi-tree, and in most cases the action on it by the graph product is acylindrical [85].

Question 15. *Can we generalise Theorem 5.3.6 to graph products?*

To prove this we would need to first prove that we can quickly generate loxodromic elements in graph products, and in this scenario we cannot appeal to an existing result in mapping class groups. However, if a positive answer to Question 10 were found using more direct methods, then this could possible be extended to the graph product case.

We can also ask about a possible generalisation of our result for mapping class groups. The key ingredient there, apart from the ability to quickly generate pseudo-Anosovs on subsurfaces, is that not only is the group acylindrically hyperbolic, but that there is a system of subgroups that have non-elementary acylindrical actions on hyperbolic spaces. This can be more formally defined as the mapping class group being *hierarchically acylindrical*. For a definition of this we refer to [33].

Abbott, Ng, and Spriano proved the following result for hierarchically acylindrical groups.

Theorem 6.0.1. [1] *Let G be a virtually torsion-free hierarchically acylindrical group such that G is not quasi-isometric to $\mathbb{Z} \times E$ for any metric space E. Then for any generating set S of G, there exists a free subgroup of G generated by two elements whose word length with respect to S is uniformly bounded.*

Note that this proves uniform exponential growth for such groups. This does not however give any kind of dichotomy, because having a direct product with $\mathbb{Z}$ does not prevent uniform exponential growth. On the other hand, the groups being ruled out in the above theorem are similar to the classes of subgroups we ruled out in Chapter 3, which suggests that it may be possible to get some kind of answer to Question 5 in the broader case of hierarchically acylindrical groups.

Question 16. *Can we generalise Corollary 5.4.3 to virtually torsion-free hierarchically acylindrical groups?*

As we try to generalise our results to other groups, we may come across a scenario where the subgroups we ruled out in Section 3.6 are not sufficient to get a dichotomy. For right-angled Artin groups and pure subgroups of mapping class groups it was sufficient to rule out those subgroups with infinite centres (and in fact it was enough to simply consider the direct product case). When we extended our results to general subgroups of mapping class groups, however, we also needed to rule out subgroups containing an infinite normal subgroup without exponential growth, where that normal subgroup is a union of finite conjugacy classes. In generalisations of our results, this may again become insufficient.

Question 17. *Which groups do not have uniform product set growth, other than those groups without uniform exponential growth, and those groups containing an infinite normal subgroup without exponential growth that is a union of finite conjugacy classes?*

Bibliography

[1] Carolyn Abbott, Thomas Ng, and Davide Spriano. Hierarchically hyperbolic groups and uniform exponential growth. *Preprint, arXiv:1909.00439*, 2019.

[2] J. W. Anderson, J. Aramoyona, and K. J. Schackleton. Uniformly exponential growth and mapping class groups of surfaces. In *In the Tradition of Alfors-Bers, IV*, volume 432 of *Contemporary Mathematics*, pages 1–6. American Mathematical Society, 2007.

[3] Yago Antolín and Ashot Minasyan. Tits alternatives for graph products. *Journal für die reine und angewandte Mathematik*, 704:55–83, 2015.

[4] Tarik Aougab. Uniform hyperbolicity of the graphs of curves. *Geometry & Topology*, 17:2855–2775, 2012.

[5] G. N. Arzhantseva and I. G. Lysenok. A lower bound on the growth of word hyperbolic groups. *Journal of the London Mathematical Society*, 73:109–125, 2006.

[6] Sahana Balasubramanya. Acylindrical group actions on quasi-trees. *Algebraic & Geometric Topology*, 17:2145—2176, 2017.

[7] Gilbert Baumslag and Alexei Myasnikov. Algebraic geometry over groups I. algebraic sets and ideal theory. *Journal of Algebra*, 219:16–79, 1999.

[8] Mladen Bestvina. $\mathbb{R}$-trees in topology, geometry, and group theory. In *Handbook of Geometric Topology*, pages 55–91. North Holland, 2001.

[9] Mladen Bestvina, Ken Bromberg, and Koji Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publications mathématiques de l'IHÉS*, 122:1–64, 2015.

[10] Mladen Bestvina and Mark Feighn. A hyperbolic $\mathrm{Out}(F_n)$-complex. *Groups, Geometry, and Dynamics*, 4:31–58, 2010.

[11] Joan S. Birman, Alex Lubotzky, and John McCarthy. Abelian and solvable subgroups of the mapping class group. *Duke Mathematical Journal*, 50(4):1107–1120, 1983.

[12] O. V. Bogopolskii and V. N. Gerasimov. Finite subgroups of hyperbolic groups. *Algebra and Logic*, 34(6):343–345, 1995.

[13] Mario Bonk and Oded Schramm. Embeddings of Gromov hyperbolic spaces. *Geometric & Functional Analysis*, 10:266–306, 2000.

[14] Brian H. Bowditch. *A course on geometric group theory*, volume 16 of *MSJ Memoirs*. Mathematical Society of Japan, 2006.

[15] Brian H. Bowditch. Tight geodesics in the curve complex. *Inventiones mathematicae*, 171:281–300, 2008.

[16] Emmanuel Breuillard and Koji Fujiwara. On the joint spectral radius for isometries of non-positively curved spaces and uniform growth. *Annales de l'Institut Fourier*, 71(1):317–391, 2021.

[17] Emmanuel Breuillard and Tsachik Gelander. Cheeger constant and algebraic entropy of linear groups. *International Mathematics Research Notices*, 56:3511–3523, 2005.

[18] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer, 1999.

[19] J. O. Button. Explicit Helfgott type growth in free products and in limit groups. *Journal of Algebra*, 389:61–77, 2013.

[20] J. O. Button. Groups acting on hyperbolic spaces with a locally finite orbit. *Preprint, arXiv:2111.13427*, 2021.

[21] Sergei Buyalo and Viktor Schroeder. *Elements of Asymptotic Geometry*. EMS Monographs in Mathematics. European Mathematical Society, 2007.

[22] Mei-Chu Chang. Product theorems in SL_2 and SL_3. *Journal of the Institute of Mathematics of Jussieu*, 7:1–25, 2008.

[23] Ruth Charney and Karen Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bulletin of the London Mathematical Society*, 41:94–102, 2009.

[24] Matt Clay, Kasra Rafi, and Saul Schleimer. Uniform hyperbolicity of the curve graph via surgery sequences. *Algebraic & Geometric Topology*, 14:3325–3344, 2014.

[25] Matt T. Clay, Christopher J. Leininger, and Johanna Mangahas. The geometry of right angled Artin subgroups of mapping class groups. *Groups, Geometry, and Dynamics*, 6(2):249–278, 2012.

[26] Michel Coornaert, Thomas Delzant, and Athanase Papadopoulos. *Géométrie et théorie des groupes: Les groupes hyperboliques de Gromov.* Number 1441 in Lecture Notes in Mathematics. Springer-Verlag, 1990.

[27] Rémi Coulon and Markus Steenbock. Product set growth in Burnside groups. *Journal de l'École polytechnique - Mathématiques*, 9:463–504, 2022.

[28] Yu-miao Cui, Yue-ping Jiang, and Wen-yuan Yang. Lower bound on growth of non-elementary subgroups in relatively hyperbolic groups. *Preprint, arXiv:2103.02304*, 2021.

[29] François Dahmani, Vincent Guirardel, and Denis Osin. Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces. *Memoirs of the American Mathematical Society*, 245(1156), 2017.

[30] Pierre de la Harpe. *Topics in Geometric Group Theory.* Chicago Lectures in Mathematics. University of Chicago Press, 2000.

[31] Thomas Delzant and Markus Steenbock. Product set growth in groups and hyperbolic geometry. *Journal of Topology*, 13:1183–1215, 2020.

[32] Cornelia Druţu and Michael Kapovich. *Geometric Group Theory.* American Mathematical Society, 2018.

[33] Matthew G. Durham, Mark F. Hagen, and Alessandro Sisto. Boundaries and automorphisms of hierarchically hyperbolic spaces. *Geometry & Topology*, 21:3659–3758, 2017.

[34] Alex Eskin, Shahar Mozes, and Hee Oh. On uniform exponential growth for linear groups. *Inventiones mathematicae*, 160:1–30, 2005.

[35] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups.* Princeton University Press, 2011.

[36] Michael Fekete. Über die Verteilung der Wurzeln bei gewissen algebraischen Gleichungen mit ganzzahligen Koeffizienten. *Mathematische Zeitschrift*, 17:228–249, 1923.

[37] Koji Fujiwara. The rates of growth in an acylindrically hyperbolic group. *Preprint, arXiv:2103.01430*, 2021.

[38] Koji Fujiwara and Zlil Sela. The rates of growth in a hyperbolic group. *Preprint, arXiv:2002.10278*, 2020.

[39] Anthony Genevois. Cubical-like geometry of quasi-median graphs and applications to geometric group theory, 2017.

[40] Anthony Genevois and Camille Horbez. Acylindrical hyperbolicity of automorphism groups of infinitely-ended groups. *Journal of Topology*, 14:963–991, 2021.

[41] Etienne Ghys and Pierre de la Harpe. *Sur les Groupes Hyperboliques d'aprés Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser, 1990.

[42] Mikhail Gromov. Groups of polynomial growth and expanding maps. *Publications Mathématiques de l'IHÉS*, 53:53–78, 1981.

[43] Mikhail Gromov. Hyperbolic groups. In *Essays in Group Theory*, pages 75–263. Springer-Verlag, 1987.

[44] Daniel Groves, Jason Fox Manning, and Alessandro Sisto. Boundaries of Dehn fillings. *Geometry & Topology*, 23(6):2929–3002, 2019.

[45] Matthias Hamann. On the tree-likeness of hyperbolic spaces. *Mathematical Proceedings of the Cambridge Philosophical Society*, 164:345–361, 2018.

[46] H. A. Helfgott. Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$. *Annals of Mathematics*, 167:601–623, 2008.

[47] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb. 1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs. *Journal of the European Mathematical Society*, 17:755–762, 2015.

[48] Tim Hsu and Daniel T. Wise. On linear and residual properties of graph products. *The Michigan Mathematical Journal*, 46:251–259, 1999.

[49] Michael Hull and Denis Osin. Conjugacy growth of finitely generated groups. *Advances in Mathematics*, 235:361–389, 2013.

[50] Nikolai V. Ivanov. Algebraic properties of the Teichmüller modular group. *Doklady Akademii Nauk SSSR*, 275(4):786–789, 1984.

[51] Nikolai V. Ivanov. *Subgroups of Teichmüller Modular Groups*, volume 115 of *Translations of Mathematical Monographs*. American Mathematical Society, 1992.

[52] Ilya Kapovich and Nadia Benakli. Boundaries of hyperbolic groups. In *Combinatorial and Geometric Group Theory*, volume 296 of *Comtemporary Mathematics*. American Mathematical Society, 2002.

[53] Ilya Kapovich and Daniel T. Wise. The equivalence of some residual properties of word-hyperbolic groups. *Journal of Algebra*, 223:562–583, 2000.

[54] Alice Kerr. Tree approximation in quasi-trees. *Preprint, arXiv:2012.10741*, 2020.

[55] Alice Kerr. Product set growth in mapping class groups. *Preprint, arXiv:2103.12643*, 2021.

[56] Sang-Hyun Kim and Thomas Koberda. Embedability between right-angled Artin groups. *Geometry & Topology*, 17(1):493–530, 2013.

[57] Sang-Hyun Kim and Thomas Koberda. The geometry of the curve graph of a right-angled Artin group. *International Journal of Algebra and Computation*, 24(2):121–169, 2014.

[58] Sang-Hyun Kim and Thomas Koberda. An obstruction to embedding right-angled Artin groups in mapping class groups. *International Mathematics Research Notices*, 2014(14):3912–3918, 2014.

[59] Thomas Koberda. Right-angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups. *Geometric and Functional Analysis*, 22:1541–1590, 2012.

[60] Thomas Koberda. Geometry and combinatorics via right-angled Artin groups. *Preprint, arXiv:2103.09342*, 2021.

[61] John M. Mackay and Jeremy T. Tyson. *Conformal Dimension: Theory and Application*, volume 54 of *University Lecture Series*. American Mathematical Society, 2010.

[62] Johanna Mangahas. Uniform uniform exponential growth of subgroups of the mapping class group. *Geometric and Functional Analysis*, 19:1468–1480, 2010.

[63] Johanna Mangahas. A recipe for short-word pseudo-Anosovs. *American Journal of Mathematics*, 135(4):1087–1116, 2013.

[64] Avinoam Mann. *How Groups Grow*, volume 395 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 2012.

[65] Jason Fox Manning. Geometry of pseudocharacters. *Geometry & Topology*, 9:1147–1185, 2005.

[66] Jason Fox Manning. Quasi-actions on trees and property (QFA). *Journal of the London Mathematical Society*, 73:84–108, 2006.

[67] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138:103–149, 1999.

[68] John McCarthy. A "Tits-alternative" for subgroups of surface mapping class groups. *Transactions of the American Mathematical Society*, 291(2):583–612, 1985.

[69] Ashot Minasyan and Denis Osin. Acylindrical hyperbolicity of groups acting on trees. *Mathematische Annalen*, 362:1055–1105, 2015.

[70] Ashot Minasyan and Denis Osin. Acylindrically hyperbolic groups with exotic properties. *Journal of Algebra*, 522:218–235, 2019.

[71] James Munkres. *Topology*. Pearson, 2nd edition, 2000.

[72] Denis Osin. Algebraic entropy of elementary amenable groups. *Geometriae Dedicata*, 107:133–151, 2004.

[73] Denis Osin. Acylindrically hyperbolic groups. *Transactions of the American Mathematical Society*, 368:851–888, 2016.

[74] Alexander A. Razborov. A product theorem in free groups. *Annals of Mathematics*, 179(2):405–429, 2014.

[75] V. N. Remeslennikov. Finitely presented group whose center is not finitely generated. *Algebra and Logic*, 13:258–264, 1974.

[76] Stanislav R. Safin. Powers of sets in free groups. *Sbornik: Mathematics*, 202(11):1661–1666, 2011.

[77] Z. Sela. Acylindrical accessibility for groups. *Inventiones mathematicae*, 129:527–565, 1997.

[78] Jean-Pierre Serre. *Trees*. Springer-Verlag, 1980.

[79] Herman Servatius, Carl Droms, and Brigitte Servatius. Surface subgroups of graph groups. *Proceedings of the American Mathematical Society*, 106(3):573–578, 1989.

[80] Peter B. Shalen and Philip Wagreich. Growth rates, $\mathbb{Z}_p$-homology, and volumes of hyperbolic 3-manifolds. *Transactions of the American Mathematical Society*, 331(2):895–917, 1992.

[81] Alessandro Sisto. Contracting elements and random walks. *Journal für die reine und angewandte Mathematik*, 742:79–114, 2018.

[82] John Stallings. Groups of dimension 1 are locally free. *Bullletin of the American Mathematical Society*, 74(2):361–364, 1968.

[83] Terence Tao. Product set estimates for non-commutative groups. *Combinatorica*, 28:547–594, 2008.

[84] Jacques Tits. Free subgroups in linear groups. *Journal of Algebra*, 20:250–270, 1972.

[85] Motiejus Valiunas. Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products. *Groups, Geometry, and Dynamics*, 15:143–195, 2021.

[86] John S Wilson. On exponential growth and uniformly exponential growth for groups. *Inventiones mathematicae*, 155:287–303, 2004.

[87] Xiangdong Xie. Growth of relatively hyperbolic groups. *Proceedings of the American Mathematical Society*, 135(3):695–704, 2007.