Splittings, Tiles, and Planarity: A Trio of Problems in Geometric Group Theory

St Anne's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2025

For my father,

who had somewhere else to be.

Acknowledgements

First and foremost, I am indebted to my supervisor, Panos Papasoglu. This thesis would not have been possible without his continued support and inspirational insight. His guidance throughout is to blame for shaping me into the mathematician I am today.

There are many people I should thank for helpful conversations and feedback over the last four years which have contributed to this thesis. This includes (and is almost certainly not limited to) Azer Akhmedov, Dario Ascari, Júlia Baligács, Filippo Baroni, Martin Dunwoody, Agelos Georgakopoulos, Sam Hughes, Dawid Kielak, Marco Linton, John Mackay, Lawk Mineh, Bruce Richter, Michah Sageev, Misha Schmalian, Davide Spriano, Ric Wade, and Henry Wilton.

I must thank my mother, Juliet, for always believing in me and for showing me what real strength looks like. I also thank my sister Nicola and her husband Mark for their support.

I am grateful to my partner Holly; she is the smartest person I know and somebody to whom I look up immensely. My thanks also go to Holly's family—Jacqui, Matt, Danny, and Sarah—for their kindness and generosity over the past few years.

Finally, I am grateful to Panos' cat, Charlie Papasoglu, for many fruitful discussions. This gratitude does not extend to Charlie's sister, Dolce.

Abstract

This thesis aims to exposit three results in the field of geometric group theory (GGT). Each of these results has their own distinct flavour, and through this we intend to showcase some of the key themes of research within GGT.

The first of these problems ('Splittings') asks for an algorithm to detect whether a given group splits as an amalgamated free product or HNN extension in a certain way. More precisely, we will study the question of whether, given a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, one can effectively decide whether G splits over a subgroup commensurable with H. The answer is positive if we assume additionally that H is residually finite, though a small technicality seems to obstruct the general case. We also present an algorithm to compute the number $\mathrm{e}(G, H)$ of relative ends, as well as partial results towards computing the number of filtered ends $\tilde{\mathrm{e}}(G, H)$.

The second problem ('Tiles') relates to a long-standing question in group theory asked independently by C. Chou and B. Weiss, which is whether every group is *monotileable*. We present progress on this question by proving that every acylindrically hyperbolic group is monotileable. Our proof expands on the techniques of a recent paper of A. Akhmedov, who showed that every hyperbolic group is monotileable. This chapter is based on joint work with L. Mineh.

The third and final problem we consider ('Planarity') concerns a coarse characterisation of virtually planar groups. That is, those finitely generated groups containing a finite-index subgroup which admits a planar Cayley graph. We show that if a finitely generated group is quasi-isometric to a planar graph, then it is virtually planar. The main technical achievement of this chapter is showing that such a group is accessible, in the sense of Wall and Dunwoody, as we do not know, *a priori*, that our group is finitely presented. This is achieved through a careful study of quasi-actions on planar graphs. In particular, we 'split' our group into more and more 'highly connected' pieces, until this quasi-action becomes so controlled that we obtain a true action on a suitable 2-complex.

Table of contents

CHAPTER I

Introduction

He had said that the geometry of the dream-place he saw was abnormal,
non-Euclidean, and loathsomely redolent of spheres and dimensions apart from ours.

— H. P. Lovecraft, *The Call of Cthulu*

The goal of this thesis is to examine and answer three questions in the modern field of *geometric group theory*. First, however, it is only right that we spend some time on the question of how we arrived here. It will, of course, be impossible to present even a small fraction of key developments of this field within this introduction. Therefore, we will focus on a selection of key moments particularly relevant to the contents of this thesis. After spending some ink on the general historical development of geometric group theory, we will then exposit the background for each of the three problems.

I.1. Prologue

When faced with some mathematical object which we hope to understand, we may decide to start investigating its automorphisms. If we consider the set of all such automorphisms, we find that the binary operation of composition equips this set with a very natural algebraic structure. This is, of course, an example of a *group*, and the group-theoretic properties of this automorphism group often shed some light on the original object of study. This idea has found application throughout the mathematical sciences and beyond. One of the earliest and most notable examples of this philosophy is credited to É. Galois and his treatment of the roots of polynomials via their permutation groups in the early nineteenth century. While Galois died in a duel at the age of just 20 in 1832, his ideas led to the development of an entire branch of algebra which now carries his name. Perhaps more importantly, Galois' work meant that mathematicians were beginning to get interested in 'group theory'. The significance of this development was recognised almost contemporaneously, as can be seen in the following quote, taken from an 1854 article of A. Cayley [**25**]:

> *The idea of a group as applied to permutations or substitutions is due to Galois, and the introduction of it may be considered as marking an epoch in the progress of the theory of algebraical equations.*

To more and more problems did the theory of groups find itself being applied. Throughout the nineteenth century, geometry, in particular, was revolutionised by the genesis of group theory. One of the figures spearheading this revolution was F. Klein, through his Erlangen[1] programme [**83**]. This programme came not long after the discovery (and widespread acceptance) of non-Euclidean geometry. Klein was proposing a sort of unifying theory of geometry using the language of group theory where, roughly speaking, every 'geometry' is described entirely by its group of symmetries. This viewpoint allowed these new geometries to be 'compared' for the first time, and the utility of group theory to geometry had been demonstrated beyond any doubt. We must stress at this point that, while this is certainly group theory of a 'geometric flavour', most experts would agree that none of the above is *geometric group theory.*

In the second half of the nineteenth century, mathematicians also began studying groups 'for the sake of groups'. One notable example of this can be seen in the work of Cayley, who in 1854 began to classify certain finite groups through the introduction of his 'tables' for groups [**25**]. Twenty-four years later, Cayley made yet another contribution to 'abstract' group theory when he demonstrated how to represent groups as certain graphical 'diagrams' [**26**]. Both the aforementioned tables and diagrams now bear his name—so-called *Cayley tables* and *Cayley graphs* respectively—and are central tools used throughout modern group theory. Cayley graphs in particular play a special role in geometric group theory, and indeed in this thesis.

Another key step in the development of abstract group theory was the introduction of groups defined not as some sort of collection of automorphisms, but by a collection of generators and defining relations. In more modern terminology, this is the introduction of *group presentations.* This change of tack is often credited to a paper of Walther von Dyck from 1882 [**49**]. Dyck was a student of Klein. In the introduction of this paper, Dyck remarks that the study of groups had centred heavily around geometry in recent years, following work of Klein, Fuchs, Poincaré, and so on. Dyck believed that a more abstract and combinatorial paradigm would help to 'separate the essence of a group from the properties that are accidentally brought into it by its special form of appearance'. This paper is sometimes highlighted as the birth of what is known as *combinatorial group theory*; the study of groups via the combinatorics of their presentations.

Much of the development of combinatorial group theory was motivated by advancements in topology. In 1895, Poincaré [**112**] gave the first definition of the *fundamental group* of a topological space. Poincaré proved that this algebraic invariant was strictly more powerful at distinguishing spaces than those numerical invariants defined earlier by Betti. Later work of Tietze [**130**] clarified and expanded upon the ideas of Poincaré, further cementing the place of group theory as a central tool in topology, and a series of papers of Max Dehn [**34, 35, 36, 37**] developed these ideas even further. The work of Dehn was particularly influential on the emergence of geometric group theory. In [**34**], Dehn reintroduced and popularised the graphical methods of Cayley. Secondly, in [**35**], Dehn introduced his three 'fundamental problems' for finitely presented groups—the *word*,

[1]The Erlangen programme is named after the University of Erlangen, where the programme was first announced in 1872.

conjugacy, and *isomorphism* problems—which went on to become central topics in the study of combinatorial (and indeed geometric) group theory. Thirdly, in [**36**], DEHN gave a combinatorial algorithm to solve the word problem in the fundamental group of a closed orientable surface, now known as *Dehn's algorithm.* It was later observed by GREENDLINGER that this algorithm is applicable to a wide range of groups [**65**], particularly *small cancellation groups.* Small cancellation theory is one of the great successes of combinatorial group theory, and a good survey can be found in [**88**, § 5]. Another fruitful direction was the study of *one-relator groups*, spearheaded by WILHELM MAGNUS, a student of DEHN, throughout the 1940s. The recent survey article [**87**] by LINTON and NYBERG-BRODDA is a good introduction to this rich and interesting sub-field. There is much, much more to be said about this period in the history of group theory. We direct the interested reader to [**27**, § I.4] for a more detailed discussion on the work of POINCARÉ, TIETZE, and DEHN mentioned above.

Throughout the middle of the twentieth century, combinatorial group theory reigned as one of the dominant paradigms through which to study groups, and the standard methods used remained combinatorial and algebraic. While it was clear that DEHN's work had made great use of some geometric methods (as had many others), there was a certain amount of debate as to whether these 'visual aids' were 'just for show', and were perhaps nothing more than a facade hiding what were ultimately algebraic results. Some discussion of this debate appears in [**27**, § I.5], and as noted there, possibly the first convincing piece of evidence that DEHN's approach was truly helpful came from STALLINGS' theorem on the ends of groups. The 'ends' of a finitely generated group were first considered by FREUDENTHAL in 1931 [**56**], and STALLINGS proved the remarkable result that a finitely generated group has more than one end if and only if it splits as an amalgam or HNN extension over a finite subgroup [**125, 126**]. This result gave due credit to the geometric methods pioneered by DEHN, as it is hard to imagine a purely algebraic proof of STALLINGS' theorem. Some argue that STALLINGS' theorem should be considered the very first 'big' result of geometric group theory. Note that 'splittings' of groups, like those mentioned here, will play an important role throughout this thesis, and are discussed further below.

Geometric methods in group theory saw a rise in popularity throughout the second half of the twentieth century. By the 1980s, there was no denying that geometric group theory had well and truly arrived. This can especially be seen with the work of GROMOV. For example, the year 1981 saw the proof of GROMOV's theorem on groups of polynomial growth [**66**], which equates the purely geometric class of groups of *polynomial growth* with the entirely algebraically-defined class of *virtually nilpotent* groups. Another highlight was his 1987 essay [**67**]. Here, GROMOV introduces the class of *hyperbolic groups.* These are groups whose Cayley graphs exhibit a coarse form of negative curvature, and they generalise ideas of 'combinatorial negative curvature' previously present in the study of small cancellation groups. It is interesting to note a group is hyperbolic if and only if it admits a finite presentation for which Dehn's algorithm, discussed in the previous paragraph, solves the word problem, thus creating a tangible link between the work of Dehn and modern geometric group theory. Note that Gromov's contributions to the genesis and popularisation of geometric group theory certainly do not stop there; see [**68**], for example. Today, geometric group theory is a vast and active field, drawing on tools and techniques from all across mathematics. For a more comprehensive survey of the modern subject, we recommend [**43**].

The goal of this thesis is to exposit three problems in geometric group theory. These problems are relatively distinct in flavour, although their backgrounds do overlap to a certain extent. We will now discuss the background of each of these problems individually.

I.2. 'Splittings'

The first question we will investigate in this thesis asks for a solution to a certain *decision problem*. This is based on an article of the present author [**90**].

The study of decision problems within group theory is almost as old as the definition of an abstract group itself, dating back to Dehn's three fundamental problems. The classical theorems of Novikov–Boone [**106, 12**] and Adian–Rabin [**1, 113**] state that these problems (as well as many others) turn out to be undecidable within the class of finitely presented groups. Thus, if we wish to search for effective solutions to group theoretical problems, we must first restrict our view to some 'nice' class of groups. Geometric hypotheses are often very helpful here, and as such the study of decision problems within geometric group theory is a rich and fascinating theory. Following Gromov, the class of *hyperbolic* groups is a common backdrop in the study of algorithmic problems. For example, the class of hyperbolic groups has uniformly[2] solvable word and conjugacy problems [**19**, ch. III.H], and more recently it was shown that one can effectively distinguish isomorphism classes of hyperbolic groups [**121, 32**].

In this chapter we will be studying a particular decision problem which asks about the existence of certain *splittings* of our group. Splittings have played an important role in group theory for almost a century, ever since the introduction of the amalgamated free product of two groups by O. Schreier in 1927 [**116**]. Later, in 1949, came HNN extensions, carrying the initials of their discoverers G. Higman, B. H. Neumann, and H. Neumann [**74**]. Although these two constructions were originally introduced as ways of 'gluing groups together' to form bigger groups, it soon became clear that there was much to be gained in 'splitting' given groups into smaller pieces. The first to make use of this technique was likely Magnus in his study of one-relator groups [**93**]. However, the power of this idea was only fully realised with the work of H. Bass and J.-P. Serre. In his monograph [**123**], Serre developed a structure theory for a group G acting on a simplicial tree T.

[2]The term *uniformly* here refers to the existence of a universal algorithm which solves the word/conjugacy problems for any hyperbolic group (and elements therein) given as part of the input. This uniformity should not be taken for granted. For example, the word problem is not uniformly solvable amongst the class of those groups with solvable word problems. One way of interpreting this fact is that 'knowing the word problem has a solution' is not actually sufficient to solve the word problem itself.

Such groups could be described as the fundamental group of a *graph of groups* structure applied to the quotient graph T/G (cf. § I.5.5). This was later substantially developed and formalised by BASS in [**7**], and now bears the name *Bass–Serre theory* in recognition of this. Bass–Serre theory unifies amalgamated free products and HNN extensions into a single, much more general framework, and is a central tool within geometric group theory.

In all that follows, when we say that a group G *splits* over a subgroup H, we mean that G admits a minimal simplicial action on a tree T without inversions, and H stabilises an edge in this action. Understanding the possible splittings of a given group can provide a deep insight into its structural properties. As such, it is incredibly natural to ask whether such splittings can be detected algorithmically. The first instance of this question comes from 3-manifold theory, dating back to the algorithm of JACO and OERTEL [**78**] which decides if a given closed irreducible 3-manifold M is *Haken*, or equivalently if $\pi_1(M)$ splits over an infinite surface group.

Splittings over finite subgroups are called *finite splittings*, and a celebrated theorem of STALLINGS [**125, 126**] states that a finitely generated group admits a finite splitting if and only if it has more than one geometric end. This demonstrates a powerful link between the coarse geometry of a group and its splitting properties. Returning to the realm of hyperbolic groups, we have the following unpublished result attributed to GERASIMOV.

THEOREM I.2.1 (GERASIMOV). *There is an algorithm which, upon input of a presentation of a hyperbolic group G, will compute the number of ends of G.*

In particular, one can effectively detect finite splittings of hyperbolic groups. This result was later extended by DIAO and FEIGHN to finite graphs of finitely generated free groups [**38**], by DAHMANI and GROVES [**31**] to relatively hyperbolic groups, and by TOUIKAN [**131**] to finitely presented groups with a solvable word problem and no 2-torsion. Also worthy of mention is the algorithm by JACO, LETSCHER, and RUBINSTEIN for computing the prime decomposition of a closed orientable 3-manifold [**77**], as well as the classical description of the free splittings of a one-relator groups [**88**, Prop. II.5.13].

Finite splittings aside, the next logical step is to detect splittings over two-ended subgroups. This was achieved for hyperbolic groups independently by BARRETT [**6**] and TOUIKAN [**131**] using quite distinct approaches. Note that TOUIKAN's algorithm here only applies in the torsion-free case.

THEOREM I.2.2 (BARRETT, TOUIKAN). *There is an algorithm which, upon input of a presentation of a hyperbolic group G, will decide if G splits over a two-ended subgroup.*

The algorithm of BARRETT, which is of particular interest to us, makes use of BOWDITCH's deep theorem on two-ended splittings of hyperbolic groups [**14**]. This theorem states that a one-ended hyperbolic group G which is not virtually Fuchsian will admit such a splitting if and only if its Gromov boundary $\partial_\infty G$ contains a cut pair. In fact, BARRETT applies this result to construct effectively

BOWDITCH's canonical JSJ decomposition of a hyperbolic group. Note that BOWDITCH's theorem provides a purely geometric characterisation of the existence of such splittings, and thus proves this property is invariant under quasi-isometries between hyperbolic groups. In [**110**], Papasoglu extends this fact to all finitely presented groups.

In Chapter II, we will expand on the techniques of BARRETT and apply them to larger splittings. We will restrict our attention to splittings over *quasi-convex* subgroups, i.e. those subgroups whose inclusion maps are quasi-isometric embeddings, since distorted subgroups exhibit global geometry which is harder to understand on a local scale. If a group G splits over a subgroup commensurable with H, we say H is *associated with a splitting*. Finding sufficient conditions for a subgroup to be associated to a splitting is a problem which has received a great amount of interest (e.g. [**118, 119, 104, 120, 105**]). Applying the results of [**105**] to the setting of quasi-convex subgroups of hyperbolic groups, we are able to prove the following decidability result.

THEOREM I.2.3 (cf. II.4.11). *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasi-convex, residually finite subgroup H. This algorithm will then decide if H is associated with a splitting, and will output such a splitting if one exists.*

It is possible to weaken somewhat the residual finiteness assumption placed on H in the theorem above and give a more general (but more involved) result. We will postpone this more technical statement until § II.4.2 (see Theorems II.4.9, II.4.10).

In light of STALLINGS' theorem, it is a natural generalisation to define the number of ends of a pair of groups (G, H) where $H \leq G$. This definition was first introduced by HOUGHTON [**76**] and later explored in more depth in the context of discrete groups by SCOTT [**117**]. The number of ends of the pair (G, H), denoted $\mathrm{e}(G, H)$, can be identified with the number of geometric ends of the quotient of the Cayley graph of G by the left action of H. This quotient graph is sometimes called the *coset graph* or *Schreier graph* of H. It is not hard to show that if G splits over H then $\mathrm{e}(G, H) \geq 2$, but the converse does not hold. Our methods give a new proof of the following theorem, originally due to VONSEEL [**134**].

THEOREM I.2.4 (VONSEEL, cf. II.4.1). *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, will output* $\mathrm{e}(G, H)$.

There is a competing notion of 'ends' of a pair of a groups which goes by several names in the literature. This idea was considered independently by BOWDITCH [**15**], KROPHOLLER–ROLLER [**85**], and GEOGHEGAN [**59**], who refer to this invariant as *coends*, *relative ends*, and *filtered ends* respectively. See [**120**, ch. 2] for a discussion on the equivalence of these three definitions. In this paper we will adopt the terminology and notation of GEOGHEGAN, and denote the number of filtered ends of the pair (G, H) by $\tilde{\mathrm{e}}(G, H)$. This value appears to be more resilient to calculation without extra hypotheses, but nonetheless we have some partial results. Recall that the *generalised word problem* for a finitely generated group H is the problem of, given words $w_0, \dots w_n$ in the generators of H, deciding whether $w_0 \in \langle w_1, \dots w_n \rangle_H$. We then have the following statement.

THEOREM I.2.5 (cf. II.4.2, II.4.3). *There is an algorithm which takes in as input a one-ended hyperbolic group G, and generators of a quasi-convex subgroup H. This algorithm will terminate if and only if $\tilde{e}(G, H)$ is finite, and if it terminates will output the value of $\tilde{e}(G, H)$.*

Furthermore, if one is also given a solution to the generalised word problem for H, then there is an algorithm which decides whether $\tilde{e}(G, H) \geq N$ for any given $N \geq 0$.

We remark in § II.2.2 that $\tilde{e}(G, H)$ can be identified with the number of components of $\partial_\infty G - \Lambda H$. Thus, the above algorithm allows us to decide if $\partial_\infty G - \Lambda H$ is disconnected. It's also worth noting that if we know a priori that $\tilde{e}(G, H)$ is finite, for example if H is two-ended, then $\tilde{e}(G, H)$ is fully computable. It does not seem possible to decide in general if $\tilde{e}(G, H) = \infty$ using our machinery for an arbitrary quasi-convex subgroup, without assuming further hypotheses. We discuss this limitation in § II.4.1.

I.3. 'Tiles'

The next chapter, Chapter III, concerns an entailment between two important classes of groups, and makes progress on a long-standing problem in group theory. It is based on an article jointly written by the present author and L. MINEH [**92**].

This chapter centres around the following basic definitions. For us, a *tile* of a group G is a subset T such that G can be covered by a disjoint union of (left) translates of T. The group G is then said to be *monotileable*[3] if every finite subset is contained in a finite tile. A long-standing question, asked independently by CHOU [**29**] and WEISS [**136**], is the following.

QUESTION I.3.1. *Is every group monotileable?*

CHOU and WEISS independently introduced monotileability; CHOU doing so in order to study *almost convergent sets* in amenable groups, while WEISS took a view towards applications to ergodic theory, via Rohklin sets in amenable groups. CHOU showed that being monotileable is preserved under extensions, directed unions, being residually monotileable, and under free products [**29**]. Moreover, the class of monotileable groups straightforwardly includes finite groups and countable abelian groups. It follows that elementary amenable groups are monotileable. It remains open whether all countable amenable groups are monotileable.

In [**124**], SEWARD explores various related tiling properties of groups, including the so-called 'CCC' property (standing for 'coherent, centred, cofinal'), which is a stronger, uniform version of

[3] In previous literature, the terms *monotileable*, *MT*, and *Property (P)* have variously been used to refer to the same property.

being monotileable. Gao, Jackson, and Seward have shown that many classes of groups are CCC, including countable free products of countable groups [**58**, Thm. 4.5.7]. Whether all countable groups have the CCC property also remains open. We would also like to mention a result of Akhmedov and Fulghesu that any subset of a free group that is connected with respect to a free generating set forms a tile [**3**, Prop. 5.1].

In Chapter III, we are interested in answering the above question for certain geometric classes of groups. This builds upon recent work of Akhmedov [**2**], who proved the following result.

Theorem I.3.2 (Akhmedov). *Hyperbolic groups are monotileable.*

This is the first time that methods from geometric group theory have been applied to the question of monotileability. However, it is currently unclear whether this theorem truly provides any new examples of monotileable groups. Indeed, it is currently unknown whether these exists a hyperbolic group which is not residually finite. Our goal will be to push these methods further, and extend Akhmedov's result to a much broader class of groups—so-called *acylindrically hyperbolic groups*—and present explicit new examples of monotileable groups.

Many groups admit interesting actions on hyperbolic metric spaces, without being hyperbolic themselves. One notable and famous example is the mapping class group $\mathrm{Mod}(\Sigma)$ of a non-exceptional surface Σ. The group $\mathrm{Mod}(\Sigma)$ acts on its curve graph $\mathcal{C}(\Sigma)$, and it is a remarkable fact, due to Masur and Minsky, that the curve graph is a hyperbolic metric space [**97**]. Of course, the curve graph is, in general, locally infinite, and the action of the mapping class group on this graph is not geometric. However, it is still somehow 'nice enough' to expect its study to be fruitful. This 'niceness' was first captured by Bowditch in [**17**], who showed that the action is *acylindrical*, in a certain sense. While acylindrical actions on trees had been studied earlier by Sela [**122**], it was Bowditch's work which truly initiated the general study of acylindrical actions on hyperbolic spaces.

Many related ideas appeared in the literature in the surrounding years. For example, the *WPD elements* of Bestvina and Fujiwara; the *weakly acylindrical actions* of Hamenstädt; the *weakly contracting elements* of Sisto; the *hyperbolically embedded subgroups* of Dahmani, Guirardel, and Osin [**33**]. Only recently has the dust settled on a unifying framework for all of the above, due to Osin. These conditions are now studied under the umbrella of *acylindrically hyperbolic groups*. These groups are, for example, commonly characterised by the property of admitting a (possibly locally infinite) hyperbolic Cayley graph upon which they act acylindrically; see § III.1.2 for details.

The mapping class group is, of course, not the only instance of this phenomenon. For example, Bestvina and Feighn have also constructed acylindrical actions of the outermorphism groups $\mathrm{Out}(F_n)$, $n \geq 2$, on hyperbolic graphs [**10**]. Overall, the class of acylindrically hyperbolic groups is incredibly vast. We refer the reader to [**108**] for a survey of examples and prominent features of the theory.

In Chapter III, we answer Question I.3.1 positively for acylindrically hyperbolic groups.

THEOREM I.3.3 (cf. III.5.1). *Acylindrically hyperbolic groups are monotileable.*

We note that there are examples of non-residually finite acylindrically hyperbolic groups that are not free products.[4] In particular, our theorem provides genuinely new examples of monotileable groups.

As mentioned earlier, the class of monotileable groups is known to satisfy certain closure properties. For example, it is closed under taking group extensions. This leads to yet more new examples of monotileable groups. We discuss two particular applications below.

Firstly, recall that a *one-relator group* is a group admitting a finite presentation with a single relator. By a result of MINASYAN and OSIN [**100**], many one-relator groups are known to be acylindrically hyperbolic, and the structure of non-acylindrically hyperbolic one-relator groups is somewhat constrained. Hence we are able to use Theorem I.3.3 to deduce the following.

COROLLARY I.3.4 (cf. III.5.2). *One-relator groups are monotileable.*

Secondly, let Γ be a simple graph with integer edge labels $m_{u,v} \geq 2$ for each $uv \in E(\Gamma)$. The associated *Artin group* is the group with the presentation

$$A_\Gamma = \left\langle V(\Gamma) \;\middle|\; \overbrace{uvu\ldots}^{m_{u,v}} = \overbrace{vuv\ldots}^{m_{u,v}},\ uv \in E(\Gamma) \right\rangle$$

Many Artin groups are known to be acylindrically hyperbolic, such as those of Euclidean type [**23**] and those whose graph Γ is not a join of two subgraphs [**28**]. For these groups Theorem I.3.3 applies straightforwardly. We record an application to the well-known family of *two-dimensional* Artin groups that is not entirely immediate.

COROLLARY I.3.5 (cf. III.5.3). *Two-dimensional Artin groups are monotileable.*

We would also like to note that it is conjectured to be the case that the central quotient $A_\Gamma/Z(A_\Gamma)$ of any irreducible Artin group A_Γ is acylindrically hyperbolic. If this were true, Theorem I.3.3 would show that all Artin groups are monotileable, as they would be products of extensions of acylindrically hyperbolic groups by abelian groups. However, this conjecture is likely quite difficult.

We briefly outline the argument behind the proof of Theorem I.3.3. Let G be an acylindrically hyperbolic group and let Γ be a (possibly locally infinite) hyperbolic Cayley graph of G, upon which G acts acylindrically. Let $F \subseteq G$ be an arbitrary finite subset. The main idea is to find an element

[4] See, for example, the groups $G(S)$ in [**21**] when S is periodic.

$z \in G$ with very large translation length that in some sense does not interact with F. In particular, multiplying elements of F on the left and right by powers of z should give one elements far away from F in Γ (with respect to the edge-path metric).

We show that the existence of such elements, which we call *swingers*, together with Γ being hyperbolic, implies that G is monotileable. In fact, taking z to be a sufficiently large swinger will make $F \cup \{z\}$ into a tile. One may think of z as 'swinging' the set F around by large enough distances, so that z can 'plug the gaps' and slowly tile G, starting from the identity element and working outwards. This idea is due wholly to Akhmedov, who presented the argument in the setting of hyperbolic groups in [**2**].

The aforementioned hyperbolic Cayley graph Γ has an ideal boundary at infinity, upon which the group acts. We exploit the dynamics of this action on the boundary to find swingers. The aim here is essentially to find z that acts by translation along an axis sufficiently far from fixed subspaces corresponding to a large ball about the identity in Γ. We recast this as a statement about fixed points of the action of G on the boundary of Γ. The main difficulty here comes from the fact that Γ may not be locally finite.

I.4. 'Planarity'

The final problem we will tackle in thesis, in Chapter IV, presents a characterisation of so-called *virtually planar groups* in terms of the coarse geometry of their Cayley graphs. Indeed, the activity of seeking to deduce algebraic information about the structure of a finitely generated group from the geometric structure of its Cayley graphs is a classic and common pastime of geometric group theorists. Amongst the oldest examples of this is a paper of Maschke from 1896 [**95**], published just eighteen years after Cayley introduced his eponymous graph [**26**]. This work features a complete classification of those finite groups admitting a Cayley graph which is *planar*. That is, it admits some embedding into the plane. In particular, such a group is one of

$$1, \quad \mathbf{Z}_n, \quad \mathbf{Z}_n \times \mathbf{Z}_2, \quad D_{2n}, \ A_4, \quad S_4, \quad A_5.$$

In other words, they are precisely the finite subgroups of the group of homeomorphisms of the 2-sphere $\mathbf{S}^2$. Of course, the aforementioned result is less about these groups themselves, and more a statement about finite, connected, planar, transitive graphs. Such graphs often form the skeleta of certain types of uniform polyhedra, which are famously well studied. A complete classification of those finite, connected, planar, transitive graphs can be found in [**55**].

The study of infinite planar Cayley graphs began much later. Results due to Wilkie [**137**] and Zieschang–Vogt–Coldewey [**139**] classify those groups with Cayley graphs that can be embedded in the plane with no accumulation points of vertices. These results show that a one-ended group admits a planar Cayley graph only if it is either a wallpaper group or a non-Euclidean crystallographic group. See [**88**, §3] for a good discussion.

In order to study those infinite-ended planar groups, the first barrier one needs to cross is that of *accessibility*. As mentioned earlier, the celebrated theorem of STALLINGS [**125, 126**] states that a finitely generated group has more than one end if and only if it splits over a finite subgroup. The definition of an accessible group is due to WALL [**135**] and says that a group is *accessible* if it splits as a graph of groups with finite edge groups, where each vertex group has at most one end. In other words, if the process of iteratively taking a finite splitting and passing to a vertex group necessarily terminates. It is a consequence of the GRUSHKO–NEUMANN theorem [**70, 103**] that all finitely generated torsion-free groups are accessible. It was conjectured by WALL in 1971 that all finitely generated groups are accessible. In 1985, DUNWOODY proved his famous theorem that every finitely presented group is accessible [**46**] and eight years later presented the first known example of an inaccessible finitely generated group [**47**], effectively closing the book on WALL's conjecture.

Returning to the plane, it was noted by LEVINSON–MASKIT [**86**] that a consequence of MASKIT's planarity theorem [**96**] is that if a Cayley graph of a given group admits a 'point-symmetric embedding' in the plane then the aforementioned group is finitely presented (and thus accessible). This was later extended to all planar Cayley graphs by DROMS in [**42**]. A recent account of MASKIT's theorem has been given by BOWDITCH in [**18**]. It is shown in [**5**] by ARZHANTSEVA–CHERIX that 'most' Cayley graphs of finitely presented groups are non-planar, in a certain statistical sense. They also prove that the admission of a planar Cayley graph is a property preserved by free products. An enumeration of planar Cayley graphs is given by GEORGAKOPOULOS–HAMANN in [**61, 62**]. Several other results relating to planar Cayley graphs are given in [**60**], including the observation that any group acting properly discontinuously on a planar manifold admits a planar Cayley graph. With all that said, it is not unreasonable to say now that the planar Cayley graphs are well understood.

On the topic of accessibility, we would also like to mention a recent theorem of ESPERET, GIOCANTI, and LEGRAND-DUCHESNE [**54**], which states that if a finitely generated group admits a minor-excluded[5] Cayley graph, then it is accessible.

We now turn to so-called *virtually planar groups*. These are finitely generated groups containing a finite-index subgroup admitting a planar Cayley graph. Here we are more concerned with questions of *rigidity*. That is, understanding which properties enjoyed by virtually planar groups actually serve to characterise this class of groups. Among the most famous theorems of this flavour is the following, originating in the work of MESS [**99**]. The proof of this result involves key contributions from CASSON–JUNGREIS [**24**], GABAI [**57**], and TUKIA [**132**]. Summarising, we state the following.

THEOREM I.4.1 (MESS *et al.*). *Let G be a finitely generated group. Then G is quasi-isometric to a complete Riemannian plane if and only if G is a virtual surface group.*

[5]Recall that a graph Γ is *minor-excluded* if there exists a finite graph which is not a minor of Γ.

In the above, a *virtual surface group* is one containing a finite index subgroup isomorphic to the fundamental group of a closed orientable surface of positive genus. In particular, virtual surface groups are virtually planar. Another proof of Theorem I.4.1 was presented by MAILLOT in [**94**], along with the following extension.

THEOREM I.4.2 (MAILLOT). *Let G be a finitely generated group. Then G is quasi-isometric to a complete, simply connected, planar Riemannian surface with non-empty geodesic boundary if and only if G is virtually free.*

Yet another proof of Theorem I.4.1 was given by BOWDITCH in [**16**]. In this paper, BOWDITCH extends this rigidity even further by presenting a selection of other characterising properties of virtual surface groups. We state some of these properties below.

THEOREM I.4.3 (BOWDITCH). *Let G be a finitely generated group. Then the following are equivalent:*

(1) *G is a virtual surface group.*
(2) *G is one-ended, contains an infinite order element, and every infinite-cyclic subgroup is codimension-1.*
(3) *G is $PD(2)$ over $\mathbf{Q}$.*
(4) *G is FP_2 over $\mathbf{Q}$, and $H^2(G; \mathbf{Q}G)$ contains a 1-dimensional G-invariant subspace.*
(5) *G is finitely presented, one-ended, semistable at infinity, and $\pi_1^\infty(G) \cong \mathbf{Z}$.*

This list is not exhaustive and one could go on here, but the point has been made. The class of virtually planar groups seems to be a very rigid class indeed. We remark that condition (3) is reminiscent of an earlier result of ECKMANN, LINNELL, and MÜLLER [**51, 50**], which characterises surface groups as those which are $PD(2)$ over the integers.

The goal of Chapter IV will be to push this rigidity even further. In particular, we will study those finitely generated groups which are quasi-isometric to planar graphs. Of course, virtually planar groups are examples of such groups, and our goal will be to show that these are the only examples. The biggest barrier to proving such a theorem is dealing with the issue of accessibility. Virtually planar groups are finitely presented and thus accessible, but *a priori*, we have no reason to believe that a finitely generated group which is quasi-isometric to a planar graph is accessible. That being said, we now state the main theorem of Chapter IV.

THEOREM I.4.4 (cf. IV.7.8). *If a finitely generated group is quasi-isometric to a planar graph then it is accessible.*

Our proof of the above is inspired by DUNWOODY's proof that planar groups are accessible [**48**]. Given Theorem I.4.4 above, our study of these groups quasi-isometric to planar graphs is reduced to the one-ended case. We deal with this case separately, and prove the following.

Theorem I.4.5 (cf. IV.2.7). *Let G be a finitely generated, one-ended group. If G is quasi-isometric to a planar graph then G is quasi-isometric to a complete Riemannian plane.*

Note that this is not immediately obvious. There exist one-ended planar graphs which are not quasi-isometric to complete Riemannian planes; see § IV.2 for an example. Theorems I.4.4, I.4.5, and I.4.1 combine to give the following corollary, which further illustrates the rigidity of virtually planar groups.

Corollary I.4.6 (cf. IV.7.9). *A finitely generated group is quasi-isometric to a planar graph if and only if it is virtually planar.*

It is interesting to note that the above, together with a result of Papasoglu–Whyte [**111**], imply that there are precisely eight quasi-isometry classes of finitely generated groups quasi-isometric to planar graphs, since every surface group is quasi-isometric to either the Euclidean plane $\mathbf{R}^2$ or the hyperbolic plane $\mathbf{H}^2$. In particular, every such group is quasi-isometric to one of the following groups:

$$1, \quad \mathbf{Z}, \ F_2, \quad \mathbf{Z}^2, \quad \Sigma, \quad \mathbf{Z}^2 * \mathbf{Z}^2, \quad \Sigma * \Sigma, \quad \mathbf{Z}^2 * \Sigma,$$

where F_2 is the free group of rank two, and Σ denotes the fundamental group of the closed orientable surface of genus 2. In fact, one can upgrade this from quasi-isometric to commensurable, via the argument given in [**9**].

Finally, we wish to highlight that the results of this chapter apply more generally to 'quasi-transitive graphs' which are quasi-isometric to planar graphs, and not just Cayley graphs. Some discussion of this generalisation is given in § IV.7.2.

I.5. Conventions

We now standardise our notation and terminology. Tools described here will be relevant throughout this thesis. Anything which only finds use within one specific chapter shall be postponed to the relevant chapter.

Throughout this thesis, the following basic standards shall be followed:

- We take 0 to be a natural number.
- Groups will typically be written multiplicatively.
- To signify that A is a strict subset of B we will write $A \subsetneq B$, and will write $A \subset B$ if the inclusion is not necessarily strict.
- If G is a group acting on the right (resp. left) on a set Z, we denote by Z/G (resp. $G\backslash Z$) the quotient object.

I.5.1. Metric and coarse geometry. Let X be a metric space. We will typically denote this metric by d_X unless otherwise stated. If there is no risk of confusion, we may abbreviate $\mathrm{d}_X = \mathrm{d}$. Given $x \in X$ and $r \geq 0$, denote by $B_X(x; r)$ the closed r-neighbourhood about $x \in X$. We similarly define $B_X(S; r)$ for subsets $S \subset X$. Let $x \in X$ and $A \subset X$. We define the distance between x and A as

$$\mathrm{d}_X(x, A) := \inf\{\mathrm{d}_X(x, a) : a \in A\}.$$

If $B \subset X$ is another subset, the *Hausdorff distance between A and B* is defined as

$$\mathrm{Haus}_X(A, B) := \max\left\{\sup_{a \in A} \mathrm{d}_X(a, B)\ ,\ \sup_{b \in B} \mathrm{d}_X(b, A)\right\}.$$

Note that $\mathrm{Haus}_X(A, B)$ is finite if and only if A is contained in a finite neighbourhood of B and vice versa.

Definition I.5.1 (Quasi-isometry). Let X, Y be metric spaces, $\lambda \geq 1$, $\varepsilon \geq 0$. Then a map $\psi : X \to Y$ is a (λ, ε)-*quasi-isometric embedding* if

$$\tfrac{1}{\lambda}\,\mathrm{d}_X(x, y) - \varepsilon \leq \mathrm{d}_Y(\psi(x), \psi(y)) \leq \lambda\,\mathrm{d}_X(x, y) + \varepsilon,$$

for all $x, y \in X$. We call ψ a (λ, ε)-*quasi-isometry* if, in addition to the above, we have that

$$\mathrm{Haus}_Y(Y, \psi(X)) \leq \varepsilon.$$

If $\lambda = \varepsilon$, we may simply call ψ a λ-*quasi-isometry.* A map satisfying this second condition is said to be *coarsely surjective*. If there exists a quasi-isometry $\psi : X \to Y$, we say that X and Y are *quasi-isometric.*

Let $\eta \geq 0$. If $\varphi : X \to Y$ is a quasi-isometry such that $\mathrm{d}_X(\varphi \circ \psi(x), x) \leq \eta$ for all $x \in X$, then we call φ a *η-quasi-inverse to ψ*. Note that every quasi-isometry has a quasi-inverse.

If a map $\gamma : I \to X$ is a (λ, ε)-quasi-isometric embedding where $I \subset \mathbf{R}$ is some interval, then we will refer to γ (or its image) as a (λ, ε)-*quasi-geodesic*.

Definition I.5.2 (Quasi-actions). Let X be a metric space, G be a group, and $\lambda \geq 1$. Then a λ-*quasi-action* of G on X is an assignment to each $g \in G$ a λ-quasi-isometry $\varphi_g : X \to X$ such that the following hold:

(1) For every $g, h \in G$, $x \in \Gamma$, we have that

$$\mathrm{d}_X(\varphi_{gh}(x), \varphi_g \circ \varphi_h(x)) \leq \lambda,$$

(2) For every $g \in G$, φ_g and $\varphi_{g^{-1}}$ are λ-quasi-inverses.

Given $C \geq 0$, we say that a quasi-action as above is C-*cobounded* if for all $x, y \in X$ there exists some $g \in G$ such that $\mathrm{d}_X(x, \varphi_g(y)) \leq C$.

Suppose X is a metric space with an action $G \curvearrowright X$ by isometries, Y is another metric space, and $\varphi : X \to Y$ is a quasi-isometry with quasi-inverse $\psi : Y \to X$. Then this induces a quasi-action of G on Y via $\varphi_g := \varphi \circ g \circ \psi$. It is easy to check that if the original action is cobounded then the resulting quasi-action is cobounded.

I.5.2. Ends. The following definitions are due to Freudenthal [**56**]. We suggest [**19, 59**] as references for this topic.

Let X be a proper[6] geodesic metric space. A continuous map $\gamma : [0, \infty) \to X$ is called a *ray*. We call a ray γ *proper* if for every bounded set $B \subset X$ we have that $\gamma^{-1}(B)$ is also bounded. We say that two proper rays γ_1, γ_2 in X are *end-equivalent* if for every bounded subset $B \subset X$ there exists some $y > 0$ such that $\gamma_1([y, \infty))$ and $\gamma_2([y, \infty))$ are contained in the same connected component of $X \setminus B$. This is easily seen to be an equivalence relation on the set of proper rays. Denote by $\Omega(X)$ the set of equivalence classes. Elements of $\Omega(X)$ are called *ends* of X. If ω is an end of X and $\gamma \in \omega$, we may call ω an *endpoint* of X. Similarly, we may refer to a bi-infinite ray as a path between two ends.

Write $|X|_{\mathrm{Fr}} = X \sqcup \Omega(X)$. Equip $|X|_{\mathrm{Fr}}$ with the topology generated by the open sets

$$V_B(x) = \{y \in |X|_{\mathrm{Fr}} : \text{there is a path from } x \text{ to } y \text{ contained in } X \setminus B\},$$

where $x \in |X|_{\mathrm{Fr}}$ and B ranges over all compact subsets of X. It is easy to check that $|X|_{\mathrm{Fr}}$ is compact, Hausdorff, and locally path connected [**40**, §8.6] and that the natural inclusion $X \hookrightarrow |X|_{\mathrm{Fr}}$ is a topological embedding. The space $|X|_{\mathrm{Fr}}$ is known as the *Freudenthal* or *end-point compactification* of X. We quickly record the following property of the Freudenthal compactification, which sets it apart from other compactifications.

Proposition I.5.3 ([**114**, Thm. 1.5(f)]). *Let X be a proper geodesic metric space. Let $U \subset |X|_{\mathrm{Fr}}$ be a connected open set. Then $U \setminus \Omega X$ is also connected.*

This property will play an important part in Chapter IV.

I.5.3. Graphs. A *graph* Γ will consist of a set $V(\Gamma)$ called the *vertex set*, and a set $\vec{E}(\Gamma)$ of directed edges equipped with an involution $e \mapsto e^{-1}$, and a map $\tau : \vec{E}(\Gamma) \to V(\Gamma)$ mapping each directed edge to its terminus. Let $\iota(e) := \tau(e^{-1})$ denote the *origin* of a directed edge. The quotient $E(\Gamma) := \vec{E}(\Gamma)/\ ^{-1}$ refers to the set of *undirected edges*. The *endpoints* of an undirected edge are the termini of its elements. We may sometimes specify an undirected edge by concatenating its endpoints, e.g. $e = uv$.

A *combinatorial path* (or just *path*) in Γ is a sequence $p = e_1, \ldots, e_n$ of $e_i \in \vec{E}(\Gamma)$ such that $\tau(e_i) = \iota(e_{i+1})$ share an endpoint every for $1 \leq i < n$. We say that p is a path from $\iota(e_1)$ to $\tau(e_n)$, and these two vertices are called the *endpoints* of p, writing $\iota(p) := \iota(e_1)$ and $\tau(p) = \tau(e_n)$ The

[6]Recall that a metric space is *proper* if closed balls are compact.

length of this path p is defined to be $\mathrm{length}(p) := n$. We say that a path is *simple* if any given vertex is visited at most once. If p is a path from u to v, and q is a path from v to w, then the concatenation pq is a path from u to w. Denote by p^{-1} the reversal of p. A *loop* is a path as above starting and ending at the same vertex, and we say that a loop is *simple* if every vertex is visited at most once, except the common initial/terminal vertex which is visited exactly twice. We define (simple) (bi-)infinite combinatorial paths similarly. Such an infinite path is sometimes called a *one-* or *two-ended ray.*

The graph Γ is said to be *connected* if there exists a path connecting any two points. If Γ is connected, we metrise $V(\Gamma)$ by defining $\mathrm{d}_\Gamma(u, v)$ as the minimal length of a path connecting $u, v \in V(\Gamma)$.

The *degree* or *valence* of a vertex $v \in V(\Gamma)$ is the number of $e \in \vec{E}X$ such that $\tau(e) = v$. We call Γ *locally finite* if every vertex has finite degree, and *bounded valence* if there exists $m \geq 0$ such that every vertex has degree at most m.

Given a subset $F \subset E(\Gamma)$, we denote by $\Gamma \setminus F$ the subgraph of Γ obtained by removing the edges in F (but not their terminal vertices). Given a subgraph $Y \subset \Gamma$, we denote by $\Gamma \setminus Y$ the subgraph of X obtained by removing all vertices of Γ contained in Y, as well as all edges with at least one vertex in Y. Given a subset $U \subset V(\Gamma)$, we denote by $\Gamma[U]$ the *subgraph induced by* U. That is, the subgraph of Γ with vertex set precisely U, where we include $e \in E(\Gamma)$ if and only if both endpoints of e lie in U. An *induced subgraph* is a subgraph which is equal to the subgraph induced by its vertex set.

Given a graph Γ, one can form its *geometric realisation* by identifying each edge $e \in E(\Gamma)$ with a copy of the unit interval. If Γ is connected then the resulting one-dimensional cell complex is equipped with a geodesic metric space. **Throughout this thesis, we will abuse terminology and often conflate a graph with its geometric realisation**. We will not mention this technicality again, as it will have no impact on the mathematics of this thesis.

I.5.4. Cayley graphs. Let G be a group and $S \subset G$. Let $S^{\pm} = \{s, s^{-1} : s \in S\}$. Assume that S is *symmetric*, that is, $S = S^{\pm}$. Then the *Cayley graph of G associated to S*, denoted by $\mathrm{Cay}(G, S)$, is defined as follows. Writing $X = \mathrm{Cay}(G, S)$, we have that $V(X) = G$ and for every $g \in G$, $s \in S$ there is a directed edge $e \in \vec{E}(X)$ such that $o(e) = g$, $t(e) = gs$. This edge is typically labelled by s. Note that $\mathrm{Cay}(G, S)$ is connected if and only if $G = \langle S \rangle$, and there is always a natural left-action of G upon $\mathrm{Cay}(G, S)$, which is free on the vertices. Unless otherwise stated, when we speak of a Cayley graph $\mathrm{Cay}(G, S)$ it shall be taken for granted that S generates G. We will usually take S to be finite, except for within Chapter III where S will typically be infinite.

If S and S' are two finite generating sets of a group G, it is a standard fact that $\mathrm{Cay}(G, S)$ and $\mathrm{Cay}(G, S')$ are quasi-isometric. Given a finitely generated group G and a finite generating set S, we write $\mathrm{e}(G) = |\Omega(\mathrm{Cay}(G, S))|$. Note that this is well-defined independently of S, since the number of ends of a metric space is a quasi-isometry invariant. It is a standard fact, due to Hopf [**75**], that $\mathrm{e}(G)$ can only take the values 0, 1, 2, or ∞.

I.5.5. Bass–Serre theory. We now standardise notation and terminology related to Bass–Serre theory. We assume that the reader is familiar with the rudiments of this topic. Standard references include [**7, 123**].

Let G be a group and T a simplicial tree. Then T is called a *G-tree* if it is equipped with an (left) action by G which is *without inversions*. That is, if $g \in G$ fixes an edge setwise then it also fixes it pointwise. We call T *non-trivial* if this action has no global fixed point. We say that T is *minimal* if there does not exist a G-invariant proper subtree. Finally, we say that T is *reduced* if it is minimal and for every edge $e = uv \in E(T)$, either u and v are orbit-equivalent, or $\mathrm{Stab}(e)$ properly includes into $\mathrm{Stab}(u)$ and $\mathrm{Stab}(v)$.

Definition I.5.4 (Graphs of groups). Let Γ be a connected graph. Then a *graph of groups* $G(\Gamma) = (G_x, \varphi_e)$ over Γ consists of the following data:

(1) For every $x \in V(\Gamma) \sqcup \vec{E}(\Gamma)$, we assign a group G_x. Given $e \in \vec{E}(\Gamma)$, this assignment is such that $G_e = G_{e^{-1}}$. The group G_x is called the *local group at x*, or perhaps the *vertex* or *edge* group at x, depending on whether x is itself a vertex or edge of Γ.
(2) For every $e \in \vec{E}(\Gamma)$ we assign a monomorphism $\varphi_e : G_e \to G_{t(e)}$.

Given $v_0 \in V(\Gamma)$ denote by $\pi_1(G(\Gamma), v_0)$ the *fundamental group of $G(\Gamma)$ based at v_0*.

The choice of the basepoint v_0 will usually be inconsequential for our purposes, and so we shall suppress this from our notation and write $\pi_1(G(\Gamma)) = \pi_1(G(\Gamma), v_0)$.

Given a G-tree T, we write $G\backslash T$ for the quotient space. Note that since G acts without inversions, the graph structure of T descends to a graph structure on $G\backslash T$.

Theorem I.5.5 (Bass–Serre). *Let G be a group and T be a G-tree. Then the quotient graph $\Gamma = G\backslash T$ admits a natural graph of groups structure $G(\Gamma) = (G_x, \varphi_e)$ such that $\pi_1(G(\Gamma))$ is naturally isomorphic to G.*

Conversely, given a graph of groups $G(\Gamma)$ with fundamental group G, there is a G-tree T, unique up to G-equivariant isomorphism, such that $\Gamma \cong G\backslash T$, and the natural graph of groups structure $G'(\Gamma)$ on Γ induced by this action is isomorphic (in an appropriate sense[7]) to $G(\Gamma)$. The tree T is called the universal covering tree *of $G(\Gamma)$, and is sometimes denoted by $T = \widetilde{G(\Gamma)}$.*

The above correspondence will play a key role throughout this thesis, and we will often play fast-and-loose with the details when switching between the 'upstairs' and 'downstairs' paradigms. This treatment will be more than sufficient for our purposes.

[7] See [**7**] for definition of a(n iso)morphism between graphs of groups.

Throughout this thesis, a *splitting* of a group G shall be taken to mean an action G on a simplicial T without inversions and without global fixed point. We say that G splits *over* the edge stabilisers. We now state STALLINGS' theorem for reference later on.

THEOREM I.5.6 (STALLINGS). *Let G be a finitely generated group. Then $\mathrm{e}(G) > 1$ if and only if G splits over a finite subgroup.*

I.5.6. Hyperbolic spaces and the Gromov boundary. We now recall some basic definitions and facts about hyperbolic spaces and groups acting on them. Let X be a geodesic space. A *geodesic triangle* with vertices x_1, x_2, x_3 is a union $\gamma_1 \cup \gamma_2 \cup \gamma_3$, where γ_i is a geodesic between x_i and x_{i+1}, and indices are taken modulo 3. Given $\delta \geq 0$, we say that such a triangle is *δ-slim* if

$$\gamma_i \subset B_X(\gamma_{i+1} \cup \gamma_{i+2}; \delta)$$

for each $i \in \{1, 2, 3\}$. The space X is now said to be *δ-hyperbolic* if every geodesic triangle in X is δ-slim. We will often omit the hyperbolicity constant δ, and simply refer to such a space as *hyperbolic*. Recall that hyperbolicity is easily seen to be a quasi-isometry invariant amongst geodesic[8] metric spaces.

We now introduce the Gromov boundary of a hyperbolic space. Much of this discussion is based on [**80**]. Let X be a δ-hyperbolic geodesic metric space. Note that we will need to consider non-proper spaces in Chapter III, so we do not assume that X is proper unless explicitly stated. Given $x, y, z \in X$, let

$$(x \cdot y)_z := \tfrac{1}{2}(\mathrm{d}_X(x, z) + \mathrm{d}_X(y, z) - \mathrm{d}_X(x, z))$$

This is known as the *Gromov product* of x and y with respect to z. Fix a basepoint $z \in X$. We say that a sequence (p_n) in X *tends to infinity* if $\liminf_{n,m\to\infty}(p_n \cdot p_m)_z = \infty$. We say two such sequences (p_n), (q_m) are *equivalent* if $\liminf_{n,m\to\infty}(p_n \cdot q_m)_z = \infty$. It is easy to check that equivalence and 'tending to infinity' do not depend on the choice of basepoint. Write $[(p_n)]$ for the equivalence class of (p_n). We then define

$$\partial_\infty X := \{[(p_n)] : (p_n) \text{ tends to infinity}\}.$$

This is called the *Gromov boundary* of X. We now topologise $\partial_\infty X$ as follows. Given $p \in \partial_\infty X$, $r \geq 0$, define

$$U(p, r) = \left\{ q \in \partial_\infty X : \exists (x_n) \in p, (y_n) \in q \text{ such that } \liminf_{i,j\to\infty} (x_i \cdot y_j)_z \geq r \right\}.$$

This defines a topology on $\partial_\infty X$ where the collection $\{U(p, r) : r \geq 0\}$ specifies a basis of neighbourhoods.

[8] The definition of hyperbolicity we give here, due to RIPS, only works for geodesic metric spaces. There is an equivalent definition in terms of the *Gromov product*, due to GROMOV, which applies more generally. One should be careful here, however, as hyperbolicity in this sense is not a quasi-isometry invariant amongst arbitrary metric spaces. See [**43**, § 11] a discussion on the different definitions.

Write $\widehat{X} = X \cup \partial_\infty X$. We topologise $\widehat{X}$ similarly. Given $p \in \partial_\infty X$ and $r \geq 0$, let

$$U'(p,r) = U(p,r) \cup \left\{ y \in X : \exists (x_n) \in p \text{ such that } \liminf_{i,j\to\infty} (x_i \cdot y)_z \geq r \right\}.$$

Again, we now define a topology on $\widehat{X}$ by specifying a basis of neighbourhoods for any given point in $\widehat{X}$; for $x \in X$, we use a basis given by the topology of X, and for p in $\partial_\infty X$ we use the $U'(p,r)$ defined above. Again, it is easy to check that these topologies on $\partial_\infty X$ and $\widehat{X}$ do not depend on the choice of basepoint z. We sometimes refer to $\widehat{X}$ as the *completion* of X. Also, the inclusion map $X \hookrightarrow \widehat{X}$ is a topological embedding with dense image. If X is proper then $\widehat{X}$ is compact, and so this gives rise to a natural compactification of X.

If $\gamma : [0,\infty) \to X$ is a quasi-geodesic ray, it is immediate that the sequence $(\gamma(n))$ tends to infinity, in the above sense. In particular, we may define the *endpoint of* γ *at infinity* as $\gamma(\infty) := [(\gamma(n))]$. Given a bi-infinite quasi-geodesic $\gamma : \mathbf{R} \to X$, we may similarly define its two endpoints at infinity, denoted $\gamma(\pm\infty)$. The following is helpful to record, especially for non-proper spaces.

Proposition I.5.7 ([**80**, Rmk. 2.16]). *Let X be a geodesic hyperbolic space. Then for any distinct $x, y \in \widehat{X}$ there exists a $(1, 20\delta)$-quasi-geodesic γ with endpoints x and y. If X is proper then we may take γ to be a geodesic.*

If a group G acts on a (not necessarily proper) hyperbolic space X by isometries, then this induces an action on the completion $\widehat{X}$. Given some $x \in X$, we consider the orbit Gx. The intersection

$$\Lambda G := \overline{Gx} \cap \partial_\infty X$$

is called the *limit set* of G. This does not depend on the choice of x. The limit set of G has either zero, one, two, or infinitely many points. In the first three cases, the action of G on X is called *elementary*; in the last, *non-elementary.* Note that when the action of G is non-elementary, ΛG is a perfect set (i.e. it has no isolated points) [**73**, Thms. 2.10, 4.5]. If the action of G on X is cobounded, then it is straightforward to see that $\Lambda G = \partial_\infty X$.

We now record that if X is proper, then there is a natural way to metrise $\partial_\infty X$. First, one extends the definition of the Gromov product to $\partial_\infty X$, by defining

$$(p \cdot q)_z = \sup \left\{ \liminf_{n,m\to\infty} (p_n \cdot q_m)_z : (p_n) \in p,\ (q_m) \in q \right\},$$

for $p, q \in \partial_\infty X$, $z \in X$. Then, we define the following.

Definition I.5.8 (Visual metric). Given a proper geodesic metric space X, a *visual metric* on $\partial_\infty X$ with parameter a and multiplicative constants k_1, k_2 is a metric ρ on $\partial_\infty X$ satisfying

$$k_1 a^{-(p\cdot q)_z} \leq \rho(p,q) \leq k_2 a^{-(p\cdot q)_z}$$

for every $p, q \in \partial_\infty X$.

We do not prove the existence of such a metric here. See [**19**, p. 434] for how to construct a visual metric given suitable parameters.

Finally, for ease of reference, we record some standard facts which will be of relevance to both Chapters II and III. The following is commonly referred to as the *Morse lemma*. This result is named after H. M. MORSE, who proved a version of this result for Riemannian surfaces [**102**]. See [**43**, Rmk. 11.41] for further discussion of the history of the Morse lemma.

LEMMA I.5.9 ([**43**, Thm. 11.72, Lem. 11.75]). *Let X be a δ-hyperbolic geodesic metric space. For every $\lambda \geq 1$, $c \geq 0$, there exists an explicit constant $\mu = \mu(\delta, \lambda, c) \geq 0$ such that the following holds.*

If p is a geodesic with endpoints $a, b \in \widehat{X}$ and q is a (λ, c)-quasi-geodesic with the same endpoints as p, then the Hausdorff distance between p and q is bounded above by μ.

The constant $\mu(\delta, \lambda, c)$ in Lemma I.5.9 is sometimes referred to as the (λ, c)-*Morse constant*.

For the next statement we need an extra definition. Let $x_0, x_1, x_2 \in \widehat{X}$ be distinct points. Then a *generalised triangle* with vertices x_0, x_1, x_2 is a union of three, possibly infinite, geodesics p_0, p_1, p_2 such that the endpoints p_i are x_i and x_{i+1}, where indices are taken modulo 3. It is a standard fact that generalised triangles are 5δ-slim, in the obvious sense [**43**, Exc. 11.86]. More generally, if we relax the p_i to be (λ, c)-quasi-geodesic triangles, then we call the resulting figure a (λ, c)-*quasi-geodesic generalised triangle*. Combining the Morse lemma I.5.9 with the fact that generalised (geodesic) triangles are uniformly slim, one immediately deduces the following which we state for ease of reference later.

LEMMA I.5.10. *Let X be a hyperbolic geodesic metric space. For every $\lambda \geq 1$, $c \geq 0$, there exists a constant $\delta' = \delta'(\lambda, c) \geq 0$ such that every (λ, c)-quasi-geodesic generalised triangle is δ'-slim.*

There is more to be said about hyperbolicity, which we defer to later chapters. In Chapter II, we will discuss some concepts specifically relevant to hyperbolic groups, such as quasi-convexity. In Chapter III, we will give some background on acylindrical actions upon hyperbolic spaces.

CHAPTER II

Detecting splittings over prescribed subgroups

Then, when they came to the evil-smelling throat of Avernus, first they soared and then they swooped down through the clear air and settled where Aeneas had prayed they would settle, on top of the tree that was two trees, from whose green there gleamed the breath of gold along the branch.

— VIRGIL, *The Aenied*

The first problem we will tackle in this thesis is that of the existence of an algorithm to detect certain splittings of a given group, under specific hypotheses. We will consider a one-ended hyperbolic group G, and ask whether there is a splitting over a prescribed quasi-convex subgroup H given as part of the input. This will is achieved by understanding the dynamics of the action of G on the limit set complement $\partial_\infty G \setminus \Lambda H$ via the local geometry of the Cayley graph. This chapter is based on the article [**90**] by the present author.

II.1. Preliminaries

II.1.1. Almost invariant subsets and (filtered) ends of pairs. We will need the idea of an *almost invariant subset*. A very good introduction to the upcoming definitions can be found in [**118**], which features many helpful examples.[1]

In what follows, G will be a finitely generated group and $H \leq G$ a finitely generated subgroup. Given two sets U and V, denote by $U \triangle V$ the *symmetric difference* of U and V. That is,

$$U \triangle V := (U \setminus V) \cup (V \setminus U).$$

We say that two sets U, V are *almost equal* if $U \triangle V$ is finite.

DEFINITION II.1.1 (Almost invariant subset). Let G act on the right on a set Z. We say that $U \subset Z$ is *almost invariant* if for all $g \in G$, Ug is almost equal to U.

DEFINITION II.1.2. We say a subset $U \subset G$ is *H-finite* or *small*, if U projects to a finite subset of $H \backslash G$. If U is not H-finite, then we say U is *H-infinite*, or *large*.

Throughout this chapter, given a subset $U \subset G$ we will write $U^* := G \setminus U$.

[1]The reader should note however that this paper contains an error, a correction of which can be found in [**120**].

Definition II.1.3. We say that a subset $U \subset G$ is *H-almost invariant* if it is invariant under the left action of H, and $H\backslash U$ is almost invariant under the right action of G on $H\backslash G$. We say that U is *non-trivial* if both U and U^* are H-infinite.

Let U and W be two non-trivial H-almost invariant subsets of G. We say that U and W are *equivalent*, if $U \triangle W$ is H-finite.

Definition II.1.4. Let U be an H-almost invariant subset. Given $g \in G$, we say that gU *crosses* U if all of

$$gU \cap U, \quad gU \cap U^*, \quad gU^* \cap U, \quad gU^* \cap U^*$$

are large. If there exists $g \in G$ such that gU crosses U then we say that U *crosses itself*. If U does not cross itself, we say it is *almost nested*. If one of the above intersections is necessarily empty for each $g \in G$, we say U is *nested*.

It is easy to see that if U and W are equivalent H-almost invariant sets, then U crosses itself if and only if W crosses itself.

Example II.1.5. Suppose a group G splits as an amalgam or HNN extension over a subgroup H. Then one can construct a non-trivial nested H-almost invariant subset $U \subset G$ as follows. Let T be the Bass–Serre tree of this splitting. Construct an equivariant map $\phi : G \to V(T)$ by mapping the identity element to some arbitrary $w \in V(T)$ and extending equivariantly. Since G acts upon itself freely and transitively, ϕ is well defined. Now, let $e \in E(T)$ be the edge stabilised by H with endpoints $u, v \in V(T)$. Deleting the interior of e separates T into two components, T_u and T_v containing u and v respectively. Set $U = \phi^{-1}(V(T_u))$, then it is a simple exercise to check that U is a non-trivial nested H-almost invariant subset of G.

The following terminology will be used throughout this chapter.

Definition II.1.6. Given a group G, a subgroup $H \leq G$ is said to be *associated to a splitting* if G splits over a subgroup commensurable with H.

We can now state the following key theorem due to Scott–Swarup [**119**], which is in some sense a converse to Example II.1.5.

Theorem II.1.7 ([**119**, Thm. 2.8]). *Let G be a finitely generated group, H a finitely generated subgroup, and U an H-almost invariant subset of G. Suppose that U is almost nested, then H is associated to a splitting.*

There is a generalisation of the above which will be important to us. Firstly, we must further loosen our requirements for nesting. Denote by $\mathrm{Comm}_G(H)$ the commensurator of H in G. That is,

$$\mathrm{Comm}_G(H) = \{g \in G : |H : H \cap H^g| < \infty,\ |H^g : H \cap H^g| < \infty\}.$$

Then we have the following definition.

DEFINITION II.1.8. Let $U \subset G$ be H-almost invariant. We say that U is *semi-nested* if $\{g \in G : gU \text{ crosses } U\}$ is contained in $\mathrm{Comm}_G(H)$.

Informally, we relax our definition to allow crossings of U by gU on the condition that gH is 'very close' to H. The following useful result, due to NIBLO–SAGEEV–SCOTT–SWARUP [**105**], says that this relaxation still produces splittings.

THEOREM II.1.9 ([**105**, Thm. 4.2]). *Let G be a finitely generated group and H a finitely generated subgroup. Suppose that there exists a non-trivial H-almost invariant subset which is semi-nested. Then H is associated to a splitting.*

This idea of 'crossings' of almost invariant sets is much deeper than what is presented here, and leads to a rich theory of *intersection numbers*. The interested reader should consult [**119**] as a starting point.

There is a natural way to 'count' these H-almost invariant subsets, which provides a useful integer invariant of the subgroup. Let $\mathcal{P}(H\backslash G)$ denote the power set of $H\backslash G$. Let $\mathcal{F}(H\backslash G)$, denote the set of finite subsets. Under the operation of symmetric difference, $\mathcal{P}(H\backslash G)$ can be seen as a $\mathbf{Z}_2$-vector space, and $\mathcal{F}(H\backslash G)$ a subspace. The quotient space $\mathcal{E}(H\backslash G) := \mathcal{P}(H\backslash G)/\mathcal{F}(H\backslash G)$ can be identified naturally with the set of H-almost invariant sets of G, modulo equivalence.

DEFINITION II.1.10. Let G be a group and $H \leq G$. We define the *number of ends of the pair* (G, H), denoted $\mathrm{e}(G, H)$ as the dimension of $\mathcal{E}(H\backslash G)$ as a $\mathbf{Z}_2$-vector space.

If G is finitely generated, it is a standard fact that $\mathrm{e}(G, \{1\})$ recovers usual geometric ends of G (see Section I.5.2 for this definition).

There is also the following characterisation of ends of pairs, which will be helpful later. This result motivates the earlier description that $\mathrm{e}(G, H)$ 'counts' H-almost invariant subsets.

PROPOSITION II.1.11 ([**117**, Lem. 1.6]). *Let G be a group, H a subgroup, and $n \geq 0$. Then $\mathrm{e}(G, H) \geq n$ if and only if there exists a collection of n pairwise disjoint H-almost invariant subsets of G.*

The number $\mathrm{e}(G, H)$ can also be seen geometrically within the coset graph.

PROPOSITION II.1.12 ([**117**, Lem. 1.1]). *Let X be a Cayley graph of G, and $H \subset G$. Then $\mathrm{e}(G, H)$ is equal to the number of ends of the coset graph $H\backslash X$.*

There is another competing, but equally interesting notion of ends of a pair of groups, namely the idea of *filtered ends*. This was considered independently by GEOGHEGAN [**59**, § 14], KROPHOLLER–ROLLER [**85**], and BOWDITCH [**15**]. We now summarise the definition as it appears in [**59**]. We first need the following technical preliminaries relating to filtrations.

Let Y be a connected, locally finite cell complex. A *filtration* $\mathcal{K} = \{K_i\}$ of Y is an ascending sequence of subcomplexes $K_1 \subset K_2 \subset \ldots \subset Y$ such that $\bigcup_i K_i = Y$. We say that this filtration is *finite* if each K_i is finite. We call the pair $(Y, \mathcal{K})$ a *filtered complex*, and a map $f : (Y, \mathcal{K}) \to (X, \mathcal{L})$ between filtered complexes is called a *filtered map* if the following four conditions hold:

(1) $\forall i$, $\exists j$ such that $f(K_j) \subset L_i$,
(2) $\forall i$, $\exists j$ such that $f(K_i) \subset L_j$,
(3) $\forall i$, $\exists j$ such that $f(Y - K_j) \subset X - L_i$,
(4) $\forall i$, $\exists j$ such that $f(Y - K_i) \subset X - L_j$.

We say that a homotopy H_t between two filtered maps is a *filtered homotopy* if H_t is a filtered map for each t. Fix a basepoint $b \in Y$, then a *filtered ray based at b* is a map $\gamma : [0, \infty) \to Y$ with $\gamma(0) = b$, which is filtered with respect to the filtration $\{[0, i] : i \in \mathbb{N}\}$ of $[0, \infty)$. We say that two filtered rays based at b are *equivalent* if there is a filtered homotopy between them fixing b. A *filtered end* of $(Y, \mathcal{K})$ is an equivalence class of filtered rays based at b. It is easy to see that the choice of b does not affect the number of filtered ends.

Let G be a finitely presented group and H a finitely generated subgroup. Let X_G be a Cayley complex for G. That is, X_G is the universal cover of some finite cell complex with exactly one vertex, and fundamental group isomorphic to G. Let $p : X_G \to H\backslash X_G$ denote quotient map. Choose a finite filtration $\mathcal{K} = \{K_i\}$ of $H\backslash X_G$, and lift this to a filtration $\mathcal{L} = \{L_i\}$ of X_G, where $L_i = p^{-1}(K_i)$.

Definition II.1.13. The *number of filtered ends of the pair* (G, H), denoted by $\tilde{\mathrm{e}}(G, H)$, is defined as the number of filtered ends of the filtered complex $(X_G, \mathcal{L})$.

Though it seems as though this definition depends on a choice of the filtration $\mathcal{K}$, it is a helpful fact that it does not; see [**59**, § 14] for details. While this definition is a little technical, we will see later on that, in the case of quasi-convex subgroups of hyperbolic groups, these filtered ends can be clarified by looking at the Gromov boundary. We conclude this section with two results relating ends of pairs and filtered ends of pairs.

Proposition II.1.14 ([**59**, Prop. 14.5.3]). *Let G be a group and H be a finitely generated subgroup. Then* $\mathrm{e}(G, H) \leq \tilde{\mathrm{e}}(G, H)$.

Equality is certainly possible but not true in general, see [**120**, p. 32] for a counterexample. The following now demonstrates the common thread between the two competing definitions.

Proposition II.1.15 ([**120**, Lem. 2.40]). *Let G be a group and H be a finitely generated subgroup of infinite index. Then* $\tilde{\mathrm{e}}(G, H) > 1$ *if and only if there exists $K \leq H$ such that* $\mathrm{e}(G, K) > 1$.

II.1.2. Hyperbolic groups. Recall the definition of a hyperbolic metric space from § I.5.6. The following definition is due to GROMOV [**67**].

DEFINITION II.1.16 (Hyperbolic group). A finitely generated group G is said to be a *hyperbolic group* if $\mathrm{Cay}(G, S)$ is hyperbolic for some (and hence any) finite generating set $S \subset G$.

For the rest of this subsection, let G be a hyperbolic group with finite generating set S. Let $X = \mathrm{Cay}(G, S)$, and fix δ such that X is δ-hyperbolic. We write

$$\partial_\infty G := \partial_\infty \mathrm{Cay}(G, S).$$

Up to homeomorphism, this does not depend on the choice of S. Note that there is a natural way to identify the connected components of $\partial_\infty X$ with the set of ends $\Omega(X)$. In particular, the boundary $\partial_\infty X$ is connected if and only if G is one-ended. Thus by STALLINGS' theorem for ends of groups (I.5.6) we have that $\partial_\infty G$ is connected if and only if G admits no finite splitting.

Throughout this chapter, we will make great use of the following basic result, often known as the *visibility* property for hyperbolic groups.

LEMMA II.1.17 (Visibility). *For every $g, g' \in V(X)$, there is a geodesic ray γ based at g' which passes within distance at most $c = 3\delta$ of g.*

Viewing geodesic rays as 'lines of sight' from the basepoint, we can imagine this theorem as saying that every point in X is (nearly) visible from every other point.

We will need the following deep result due to BESTVINA and MESS [**11**], BOWDITCH [**14**], and SWARUP [**128**]. This result essentially quantifies the connectedness of the boundary and relates it to the local geometry of the Cayley graph. Let $c = 3\delta$ as in Lemma II.1.17, and let $M = 6c + 2\delta + 3$. We then have the following.

DEFINITION II.1.18. We say that X satisfies $\ddagger_n$ for $n \geq 1$ if for every $R \geq 0$ and every $x, y \in X$ such that $\mathrm{d}_X(x, 1) = \mathrm{d}_X(y, 1) = R$ and $\mathrm{d}_X(x, y) \leq M$ we have that there is a path through $X \setminus B_X(1; R)$ connecting x to y, of length at most n.

THEOREM II.1.19 ([**11**, Props. 3.1, 3.2]). *Let G be a hyperbolic group, with δ-hyperbolic Cayley graph X. Then $\partial_\infty G$ is connected if and only if there exists $n \geq 1$ such that X satisfies $\ddagger_n$.*

Note that through the algorithm presented in [**31**] by DAHMANI and GROVES, one can decide if a given hyperbolic group G is one-ended, and if so this algorithm will output n such that $\ddagger_n$ holds in the Cayley graph associated to the given generators of G. Their algorithm also applies to relatively hyperbolic groups, but we will not use that here.

II.1.3. Quasi-convex subgroups and splittings. Recall that a subset of a geodesic space is termed *convex* if it contains any geodesic between any two of its points. This notion is too precise for the setting of groups, so we must 'quasi-fy' it.

Let $Q \geq 0$, and let G be a hyperbolic group with Cayley graph X. Then a subgroup $H \leq G$ is called Q*-quasi-convex* if for every $h, h' \in H$, any geodesic path between h, h' in X is contained in the closed Q-neighbourhood of H. The quasi-convexity of H does not depend on the choice of Cayley graph for G, given that G is hyperbolic, and quasi-convex subgroups of hyperbolic groups are also hyperbolic. The following characterisation of quasi-convexity will be important.

LEMMA II.1.20. *Let G be a hyperbolic group, $H \leq G$ a finitely generated subgroup, and fix word metrics on these groups. Then H is quasi-convex if and only if the inclusion map $H \hookrightarrow G$ is a (λ, ε)-quasi-isometric embedding for some $\lambda \geq 1$, $\varepsilon \geq 0$.*

Moreover, given a presentation of G and generators of H, if H is quasi-convex then these constants can be found algorithmically.

The above follows immediately from KAPOVICH's algorithm [**79**] for computing the quasi-convexity constant Q, together with [**19**, § III.Γ, Lem. 3.5].

Quasi-convex subgroups interact with the Gromov boundary in a very controlled way. Recall the definition of a *limit set* from § I.5.6. By considering the action of H on the Cayley graph of G, we define $\Lambda H \subset \partial_\infty G$ in the obvious way. It is a standard fact that if H is quasi-convex then $\Lambda H \cong \partial_\infty H$, and ΛH is a closed H-invariant subset of $\partial_\infty G$ (see e.g. [**80**]). Combining Lemmas II.1.17 and I.5.9 we immediately arrive at the following, which will play a big role in the next section.

LEMMA II.1.21. *Let X be a δ-hyperbolic Cayley graph of a hyperbolic group G, and let H be a Q-quasi-convex subgroup. There exists some computable constant $\eta = \eta(\delta, Q) > 0$ with the following property. If γ is a geodesic ray in X based at 1 such that $\gamma(\infty) \in \Lambda H$, then γ is contained in the closed η-neighbourhood of H.*

Another important fact about quasi-convex subgroups of hyperbolic groups is that they are 'almost' self-commensurating. More formally we have the following classical result due to ARZHANTSEVA [**4**] and KAPOVICH–SHORT [**82**].

THEOREM II.1.22 ([**4**, Thm. 2]). *Let G be a hyperbolic group and H a quasi-convex subgroup. Then H has finite index in* $\mathrm{Comm}_G(H)$.

COROLLARY II.1.23. *Let H be a quasi-convex subgroup of a hyperbolic group G. Then the commensurator* $\mathrm{Comm}_G(H)$ *is the unique maximal finite index overgroup of H.*

PROOF. Let $K \leq G$ be some finite index overgroup of H, and let $k \in K$. Clearly $H^k \subset K$, so $H^k \cap H \leq H \cap K = H$ and H^k has finite index in K. Then

$$|K : H \cap H^k| = |K : H| \cdot |H : H \cap H^k|,$$

so $|H : H \cap H^k| < \infty$. Similarly, $|H^k : H \cap H^k| < \infty$, and so $k \in \mathrm{Comm}_G(H)$. $\square$

We conclude this section with a description of the vertex groups of a splitting in which the edge groups are quasi-convex. We say that a subgroup is *full* if it does not have any finite index overgroups.

PROPOSITION II.1.24 ([**14**, Prop. 1.2]). *Let G be a one-ended hyperbolic group, and suppose that G splits over a quasi-convex subgroup as an amalgam or an HNN extension. Then the vertex groups of this splitting are full quasi-convex.*

PROOF. The fact that they are quasi-convex is proven in [**14**], so we will just show here that they are necessarily also full.

Let G act simplicially, minimally, and without inversions on a tree T with a single edge orbit, such that the edge stabilisers are quasi-convex. Let v be a vertex of T, and let G_v denote the stabiliser of v. Suppose that G_v is not full, so there is finite index overgroup $K > G_v$. Then the K-orbit of v is finite, so there is some finite subtree $\Sigma \subset T$ stabilised by K. Let u be the geometric center of Σ. If u is a midpoint of an edge e, then it must be the case that the action of K inverts this edge. We assumed that G acts without inversions so this cannot happen, and so u is a vertex of T. Since u is distinct from v (else $K = G_v$), we see that there is a vertex w adjacent to v such that

$$G_v \leq G_w \leq K.$$

In particular, G_v has finite index in G_w.

Let e be the edge connecting v to w. We split into two cases. Firstly, suppose that this splitting is an amalgam. Then $G_e = G_v \cap G_w = G_v$, but then the given splitting is of the form $G = G_w *_{G_v} G_v$, and so $G_w = G$ and G fixes a point on T.

Suppose instead that G acts on T with a single vertex orbit, i.e. that this splitting is an HNN extension. Then there is some $g \in G$ such that $G_w = G_v^g$. Suppose that $G_w > G_v$, then we find the following sequence

$$G_v < G_v^g < G_v^{g^2} < G_v^{g^3} < G_v^{g^4} < \dots,$$

where each has finite index in the last. But then G_v has no maximal finite index overgroup, which contradicts Corollary II.1.23 since G_v is quasi-convex. Thus $G_w = G_v = G_e$. In particular, by translating this picture around we see that the stabiliser of every edge in T is equal to the stabiliser of either of its endpoints. Since T is connected, it follows that G_v actually fixes T pointwise and is in fact the kernel of the action of G on T. Thus, G_v is a normal subgroup of G. Since G_v is quasi-convex, this can only happen if G_v is finite or has finite index in G [**82**]. By assumption, G

is one-ended and so G_v has finite index in G. But then G must act on T with a global fixed point. With this contradiction, we conclude that G_v must be full. □

The following definition will be helpful.

DEFINITION II.1.25 (Lonely subgroups). Let G be a group, then a subgroup $H \leq G$ is said to be *lonely* if there does not exist another subgroup H' distinct from H such that H is commensurable with H'.

We then have the following dichotomy, which will play a central role in our algorithmic results.

PROPOSITION II.1.26. *Let G be a hyperbolic group and H a quasi-convex subgroup such that either*

(1) *$H \neq \mathrm{Comm}_G(H)$, or*
(2) *H is lonely.*

Then H is associated to a splitting if and only if there exists a semi-nested H-almost invariant set.

PROOF. Applying Theorem II.1.9, it is clear that if we have a semi-nested H-almost invariant set then H is associated to a splitting.

Conversely, suppose that G splits over a subgroup H' commensurable with H. If H is lonely then the result is clear as we necessarily split over H, thus there is a nested H-almost invariant set as in Example II.1.5. So assume instead that $H \neq \mathrm{Comm}_G(H)$. We now split into two cases.

Firstly, suppose that H' has infinite index in the vertex group(s) of its splitting. We have by Theorem II.1.22 that $\mathrm{Comm}_G(H)$ is commensurable with H', so it will be strictly contained in a vertex group of this splitting, up to conjugacy. We can then transform our splitting over H' easily into a splitting over $K := \mathrm{Comm}_G(H)$ via e.g.

$$G = A *_{H'} B = A *_{H'} (K *_K B) = (A *_{H'} K) *_K B.$$

The HNN case is similar. Since we have a non-trivial splitting over K, there is an almost nested K-almost invariant set by Example II.1.5. Clearly such a set is also H-almost invariant, and so we are done.

Suppose instead that H' has finite index in one of the vertex stabilisers. By Proposition II.1.24 we know that the vertex groups of the splitting over H' are full, so by Corollary II.1.23 we deduce that one vertex group of the splitting over H' is precisely $\mathrm{Comm}_G(H)$. Let T be the Bass–Serre tree of this splitting. Then H fixes a vertex of T, namely the vertex stabilised by $\mathrm{Comm}_G(H)$. Let $e \in E(T)$ be the edge stabilised by H' with endpoints $v, u \in V(T)$, and suppose without loss of generality that H fixes v. Then the H-orbit of e is a finite collection of edges abutting v. Since $H \neq \mathrm{Comm}_G(H)$ there are multiple H-orbits of edges abutting v. Let $F = \{he \ : \ h \in H\}$. Then deleting the interior of the edges in F separates T into a disjoint collection of subtrees. Let T_v be the subtree which contains v, then set $U = \phi^{-1}(V(T_v))$, where $\phi : G \to T$ is the map constructed in Example II.1.5. It is clear that U is a non-trivial H-almost invariant set, and

$$\{g \in G : gU \text{ crosses } U\} \subset G_v = \mathrm{Comm}_G(H),$$

where $G_v \leq G$ is the stabiliser of v. Thus U is semi-nested by definition. □

Remark II.1.27. It might seem strange that in the hypothesis of Proposition II.1.26 we ask that H be distinct from its commensurator. This rules out, for example, maximal cyclic subgroups (which are often associated to splittings). This exclusion can be explained by an example.

Suppose our group G splits as an amalgam $G = A *_H B$ where $|B : H| < \infty$. Then the vertex group B is associated to a splitting, by definition. However, it is certainly not clear if G contains some semi-nested B-invariant subset (it is helpful to check where the construction given in the above proof fails in this case). This technicality can be overcome by replacing our given quasi-convex subgroup with one of its finite index subgroups (if one exists), which explains why it is helpful to assume that our given subgroup is residually finite.

II.2. Limit set complements

The key aim of this section is to understand the connectivity of $\partial_\infty G \setminus \Lambda H$ via local geometry. In particular, we will generalise a toolbox introduced by Barrett in [**6**].

II.2.1. Annular neighbourhoods. The Gromov boundary encodes a wide variety of data about our group, but we cannot access it directly. Thus, we must characterise the presence of the topological features we care about via some kind of equivalent local geometric feature in the Cayley graph.

For what follows we will need to fix some notation. Throughout this section G will be a one-ended hyperbolic group with δ-hyperbolic Cayley graph X, and H will be a Q-quasi-convex subgroup with a fixed finite generating set Y. Note that given a presentation of a hyperbolic group, one can effectively compute the hyperbolicity constant δ of the corresponding Cayley graph via [**109**]. Given $x, y \in X$, we may denote by $[x, y]$ a choice of geodesic path between x and y. Let $c = 3\delta$ as in Lemma II.1.17, $\mu = \mu(\delta, 1, 0)$ denote the $(1, 0)$-Morse constant, as in Lemma I.5.9. Fix a visual metric ρ on $\partial_\infty G$ (see Definition I.5.8), with parameter a and multiplicative constants k_1, k_2. We will make use of some methods from [**11**], so recall Theorem II.1.19 and fix $n \geq 0$ such that $\ddagger_n$ holds in X. Finally, choose $\lambda \geq 1$, $\varepsilon \geq 0$ such that the inclusion map $H \hookrightarrow G$ is a (λ, ε)-quasi-isometric embedding as in Lemma II.1.20, and $\eta = \eta(\delta, Q)$ as in Lemma II.1.21.

Lemma II.2.1. *Let $p \in \partial_\infty G \setminus \Lambda H$. Then for $\gamma \in p$, we have that*

$$\lim_{t \to \infty} \mathrm{d}_X(\gamma(t), H) = \infty.$$

PROOF. Suppose not, so $\liminf_{t\to\infty} \mathrm{d}_X(\gamma(t), H) < \infty$. Then there exists a sequence $t_i \to \infty$, and some $A \geq 0$ such that $\mathrm{d}_X(\gamma(t_i), H) < A$ for all i. Let $h_i \in H$ be such that $\mathrm{d}_X(\gamma(t_i), h_i) < A$.

Recall that we fixed a generating set Y of H, and thus a word metric on H. Let δ_H be such that H is δ_H-hyperbolic. By Lemma II.1.17, choose geodesic rays ζ_i in H based at 1 such that ζ_i passes within $3\delta_H$ of h_i. Let $\iota : H \hookrightarrow G$ be the inclusion map, which is a (λ, ε)-quasi-isometric embedding. Then $\iota(\zeta_i)$ is a (λ, ε)-quasi-geodesic ray at 1. Let ν_i be a geodesic ray in X with the same endpoints as $\iota(\zeta_i)$. Then

$$\mathrm{d}_X(\nu_i, \gamma(t_i)) < A + 3\lambda\delta_H + \varepsilon + \mu(\lambda, \varepsilon),$$

for every i. It is now routine to see that $\nu_i(\infty) \to p$ in $\partial_\infty G$. But ΛH is a closed subset of $\partial_\infty G$, so $p \in \Lambda H$. This is a contradiction, so the lemma follows. □

We now introduce some notation borrowed from [**6**]. This notation will be used throughout the rest of this chapter. Let $0 \leq r \leq K \leq R$ be constants, where R is possibly infinite. Firstly, to ease notation let us denote by $N_R(H)$ the closed R-neighbourhood of H in X, i.e. $N_R(X) := B_X(H; R)$. Denote by $N_{r,R}(H)$ the annular region

$$N_{r,R}(H) = \{x \in X \ : \ r \leq \mathrm{d}_X(x, H) \leq R\}.$$

Let $C_K(H) = \{x \in X \ : \ \mathrm{d}_X(x, H) = K\}$, and finally let $A_{r,R,K}(H)$ be the union of the components of $N_{r,R}(H)$ which intersect $C_K(H)$. The geometric significance of $A_{r,R,K}(H)$ is that we take an annular region about H, but discard the components which are 'too close' to H. This is illustrated in Figure II.1. Following Barrett, we will demonstrate a correspondence between the components of $\partial_\infty G \setminus \Lambda H$ and the components of $A_{r,R,K}(H)$, provided r, R, and K are sufficiently large.

Next, we introduce 'shadows'. Given a component Z of $A_{r,\infty,K}(H)$, define its *shadow* $\mathcal{S}Z$ as the set of points $p \in \partial_\infty G$ such that for any ray $\gamma \in p$ based at 1, we have that $\gamma(t)$ is in Z for all

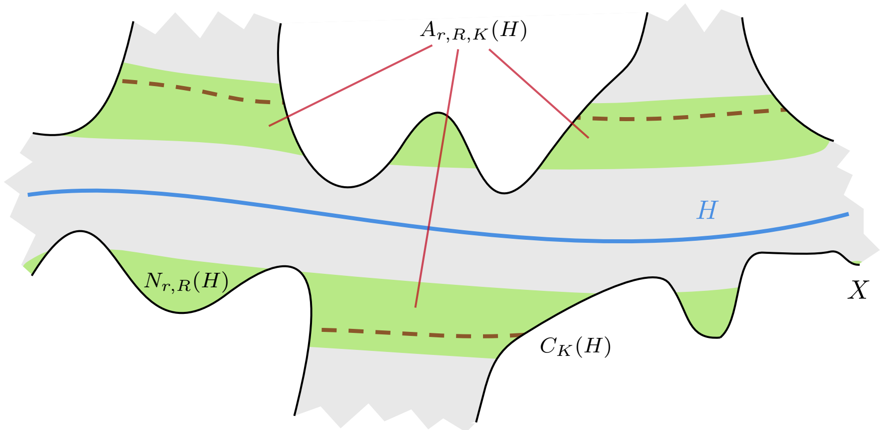

FIGURE II.1. The region $A_{r,R,K}(H)$ surrounding H. In this cartoon H is depicted as a line for illustrative purposes.

sufficiently large t. We also extend the definition of a shadow to a union of components $W = \bigcup_i Z_i$ by defining $\mathcal{S}W = \bigcup_i \mathcal{S}Z_i$.

We now prove five lemmas, generalised from [**6**]. Originally, these results related to bi-infinite geodesics, but careful inspection of their proofs reveals that we only really need the fact that a geodesic path in a hyperbolic space is a quasi-convex set. This allows us to generalise some of the machinery in this paper from geodesics to quasi-convex subgroups. Apart from some small changes, the following proofs are practically reproduced verbatim from [**6**].

Lemma II.2.2 ([**6**, Lem. 1.17]). *For $r > \eta$, $K \geq r$, we have that*

$$\partial_\infty G \setminus \Lambda H = \bigcup_Z \mathcal{S}Z,$$

where Z ranges over the connected components of $A_{r,\infty,K}(H)$. Moreover, $\mathcal{S}Z_1 \cap \mathcal{S}Z_2 = \emptyset$ for distinct components Z_1, Z_2.

Proof. It is clear from the definition of $\mathcal{S}Z$ that distinct components have disjoint shadows. Since $r > \eta$, we have by Lemma II.1.21 that $\mathcal{S}Z$ and ΛH are disjoint. Now let $p \in \partial_\infty G \setminus \Lambda H$. We need to check that $p \in \mathcal{S}Z$ for some component Z of $A_{r,\infty,K}(H)$. Let $\gamma \in p$, then there is some t_0 such that for $t \geq t_0$ we have that $\mathrm{d}_X(\gamma(t), H) \geq r + \mu$.

Let Z be the component of $A_{r,\infty,K}(H)$ containing $\gamma(t)$ for $t \geq t_0$. We claim that if γ' is another ray in p, then, $\gamma'(t')$ lies in Z for sufficiently large t'. We know that the Hausdorff distance between γ and γ' is at most μ. In particular, for every t' there is some t such that $\mathrm{d}_X(\gamma(t), \gamma'(t')) \leq \mu$, and by the triangle inequality t and t' satisfy $|t - t'| \leq \mu$. Then if $t' \geq t_0 + \mu$, we must have that $t \geq t_0$, and so the segment between $\gamma'(t')$ and $\gamma(t)$ lies in $A_{r,\infty,K}(H)$, and thus $\gamma'(t') \in Z$. It follows $p \in \mathcal{S}Z$. $\square$

Lemma II.2.3 ([**6**, Lem. 1.18]). *Let $r > \eta$ and $K \geq r + Q + \delta + c$. Then for every component Z of $A_{r,\infty,K}(H)$, we have that $\mathcal{S}Z$ is non-empty.*

Proof. Let $x \in C_K(H) \cap Z$, and using Lemma II.1.17 choose a geodesic ray γ based at 1 which passes within c of x, say $\mathrm{d}_X(\gamma(t), x) \leq c$. Since $K > r + c$, we have that $\gamma(t) \in Z$. In particular,

$$\mathrm{d}_X(\gamma(t), H) > r + Q + \delta.$$

Suppose now that for some $t' \geq t$, we have that $\mathrm{d}_X(\gamma(t'), H) \leq r$. Let $y \in H$ be a nearest point projection of $\gamma(t')$. Then consider the geodesic triangle $[1, y, \gamma(t')]$, where the segment $[1, \gamma(t')]$ is precisely an initial segment of γ. Then $\gamma(t)$ is either δ-close to $[y, \gamma(t')]$ or $[1, y]$. In any case,

$$\mathrm{d}_X(\gamma(t), H) \leq r + Q + \delta,$$

which is a contradiction. Thus $\gamma(t') \in Z$ for every $t' \geq t$. It now follows by an argument similar to that in the proof of Lemma II.2.2 that $\gamma(\infty) \in \mathcal{S}Z$. $\square$

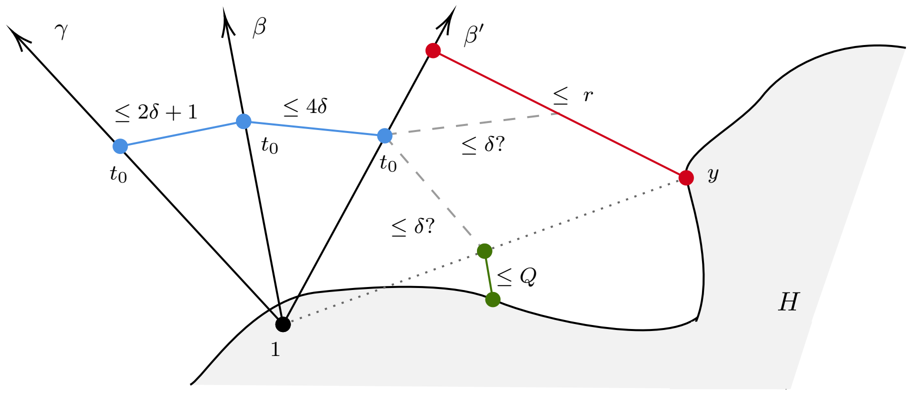

FIGURE II.2. Bounding the distance between $\gamma(t_0)$ and H.

LEMMA II.2.4 ([**6**, Lem. 1.19]). *Let $r > \eta$, $K > r$, then for every component Z of $A_{r,\infty,K}(H)$, we have that $\mathcal{S}Z$ is open and closed in $\partial_\infty G \setminus \Lambda H$.*

PROOF. Let $\gamma \in p \in \mathcal{S}Z$, and for $t \geq 0$ let

$$V_t(\gamma) = \{\beta(\infty) : \beta \text{ is a ray based at } 1,\ \mathrm{d}_X(\gamma(t), \beta(t)) < 2\delta + 1\}.$$

As t varies, this forms a basis of neighbourhoods about p. By Lemma II.2.1 there exists a t_0 such that for all $t \geq t_0$, we have that

$$\mathrm{d}_X(\gamma(t), H) > r + Q + 7\delta + 1.$$

We claim that $V_{t_0}(\gamma) \subset \mathcal{S}Z$, which implies that $\mathcal{S}Z$ is open.

Let $\beta \in q \in V_{t_0}(\gamma)$, so by definition we have $\mathrm{d}_X(\beta(t_0), \gamma(t_0)) < 2\delta + 1$. Then consider another ray $\beta' \in q$ also, so $\mathrm{d}_X(\beta(t_0), \beta'(t_0)) < 4\delta$. We need to show that $\beta'(t) \in Z$ for all $t \geq t_0$. Suppose this is not the case. We have that $\beta'(t_0)$ is in Z, so for this not to hold we must have that β' 'leaves' Z at some point. Thus for some $t' \geq t_0$, we have that $\mathrm{d}_X(\beta'(t'), H) \leq r$, say for some $y \in H$ we have $\mathrm{d}_X(\beta'(t'), y) \leq r$. Consider a geodesic path $[1, y]$. Since H is Q-quasi-convex this path is contained in the Q-neighbourhood of H. Combine this with the fact that the triangle $[1, \beta'(t'), y]$ is δ-slim, then inspection of Figure II.2 reveals that

$$\mathrm{d}_X(\gamma(t_0), H) \leq r + Q + 7\delta + 1.$$

This is a contradiction, so $\mathcal{S}Z$ is open.

To conclude, note that the shadows $\mathcal{S}Z$ form a disjoint open cover of $\partial_\infty G \setminus \Lambda H$. It follows that that they must also be closed. $\square$

For the next lemma, recall that k_1, k_2 and a denote the parameters of our fixed visual metric on $\partial_\infty G$ (recall Definition I.5.8). Also recall that n has been fixed so that $\ddagger_n$ holds in X, and $M = 6c + 2\delta + 3$ as defined in § II.1.2.

Lemma II.2.5 ([**6**, Lem. 1.20]). *Suppose that r satisfies*

$$r > 2\log_a\left(\frac{k_2(n-1)}{k_1(1-a^{-1})}\right) + M + 8\delta + \eta + C,$$

and let $K > r$. Then for every component Z of $A_{r,\infty,K}(H)$, $\mathcal{S}Z$ is contained within exactly one connected component of $\partial_\infty G \setminus \Lambda H$. Moreover, every component of $\partial_\infty G \setminus \Lambda H$ is path connected.

Proof. Let $p, q \in \mathcal{S}Z$, and let $\alpha_1 \in p$, $\alpha_2 \in q$. Then there is some t_1, t_2 such that $\alpha_1(t_1), \alpha_2(t_2) \in Z$. Let $\phi : [0,\ell] \to Z$ be a path connecting $\alpha_1(t_1)$ to $\alpha_2(t_2)$, parameterised by arc length.

We follow the methodology of Bestvina–Mess [**11**] and 'project' this path to the boundary. For every $i \in \mathbf{Z} \cap [0,\ell]$, let β_i be a ray from 1, passing through a point $\beta_i(m_i) = z_i$ such that $\mathrm{d}_X(z_i, \phi(i)) \leq c$. We will show that $\beta_i(\infty)$ and $\beta_{i+1}(\infty)$ can be connected by a path in $\partial_\infty G$ avoiding ΛH. This implies the result. Note that each $\beta_i(\infty)$ is not in ΛH, since $r \geq \eta + c$.

For notational convenience, assume that $i = 0$. Let n be such that $\ddagger_n$ holds in X. For every n-adic $t \in [0,1]$, by induction on the power k of the denominator we define β_t, satisfying

$$\mathrm{d}_X(\beta_{j/n^k}(m_i + k), \beta_{j+1/n^k}(m_i + k)) \leq M,$$

where $M = 6c + 2\delta + 3$, for each $0 \leq j < n^k$. Note that the base case of this induction holds because $M \geq 2c + 1$. Secondly, the triangle inequality then gives the following lower bound:

$$\begin{aligned}(\beta_{j/n_k}(\infty) \cdot \beta_{j+1/n_k}(\infty))_1 &\geq \liminf_{n_1,n_2}(\beta_{j/n_k}(n_1) \cdot \beta_{j+1/n_k}(n_2))_1\\ &\geq (\beta_{j/n_k}(m_0 + k) \cdot \beta_{j+1/n_k}(m_0 + k))_1\\ &= m_0 + k - M/2.\end{aligned}$$

We therefore deduce that

$$\rho(\beta_{j/n_k}(\infty), \beta_{j+1/n_k}(\infty)) \leq k_2 a^{-m_0 - k + M/2}.$$

Inductively applying the triangle inequality and summing the geometric series, we thus obtain the bound that for every n-adic rational $t \in [0,1]$, we have

$$\rho(\beta_0(\infty), \beta_t(\infty)) \leq \frac{k_2(n-1)a^{-m_0+M/2}}{1-a^{-1}}.$$

We then define a path $\psi : [0,1] \to \partial_\infty G$ by $\psi(t) = \beta_t(\infty)$, for every n-adic t, and extending to $[0,1]$ continuously. We have shown that this path is contained within a ball of radius

$$\frac{k_2(n-1)a^{-m_0+\frac{M}{2}}}{1-a^{-1}} \tag{1}$$

around $\beta_0(\infty)$.

We now seek to place a lower bound on $\rho(\beta_0(\infty), \Lambda H)$ in terms of m_0. Let γ be a ray from 1 with limit in ΛH. So γ is contained in the η-neighbourhood of H. By [**19**, § III.H, 3.17] we have that

$$(\beta_0(\infty) \cdot \gamma(\infty))_1 \leq \liminf_{n_1,n_2}(\beta_0(n_1) \cdot \gamma(n_2)) + 2\delta. \tag{2}$$

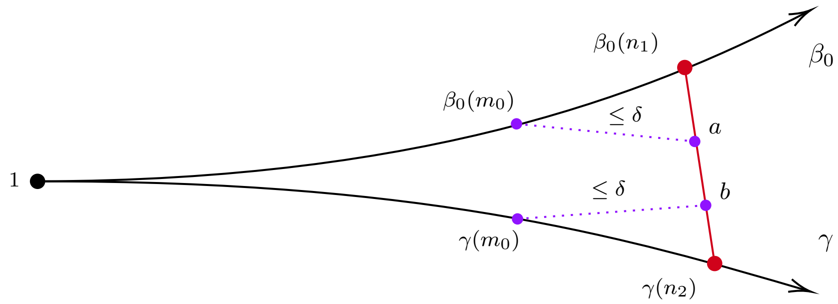

FIGURE II.3. The point a is closer to $\beta_0(n_1)$ than b.

So let $n_1, n_2 \geq m_0$. We will show that $(\beta_0(n_1) \cdot \gamma(n_2))_1$ can be approximated by $(\beta_0(m_0) \cdot \gamma(m_0))_1$. Since $\mathrm{d}_X(\beta_0(m_0), H) \geq r > \delta + \eta$, we have that $\mathrm{d}_X(\beta_0(m_0), \gamma) > \delta$. So there exists some $a \in [\beta_0(n_1), \gamma(n_2)]$ such that $\mathrm{d}_X(\beta_0(m_0), a) \leq \delta$. Similarly, we have that $\mathrm{d}_X(\gamma(m_0), \beta_0) > \delta$, as $r > 2\delta + \eta$. Thus again there exists a point $b \in [\beta_0(n_1), \gamma(n_2)]$ such that $\mathrm{d}_X(\gamma(m_0), b) \leq \delta$.

We now check that a and b are 'in order'. Suppose that b is closer to $\beta_0(n_1)$ than a. Then by considering the geodesic triangle $[a, \beta_0(m_0), \beta_0(n_1)]$, we see that $\mathrm{d}_X(b, \beta_0) \leq 2\delta$. But this implies that $\mathrm{d}_X(\beta_0(m_0), \gamma(m_0)) \leq 6\delta$, which is a contradiction since $r > 6\delta + \eta$. Therefore the picture we have looks something like Figure II.3.

We can thus conclude that

$$\mathrm{d}_X(\beta_0(n_1), \gamma(n_2)) = \mathrm{d}_X(\beta_0(n_1), a) + \mathrm{d}_X(a, b) + \mathrm{d}_X(b, \gamma'(n_2)).$$

We can then compute the following inequality.

$$\begin{aligned}
(\beta_0(n_1) \cdot \gamma(n_2))_1 - (\beta_0(m_0)) \cdot \gamma(m_0))_1 &= \frac{1}{2}[n_1 + n_2 - \mathrm{d}_X(\beta_0(n_1), \gamma(n_2))] \\
&\quad - \frac{1}{2}[2m_0 - \mathrm{d}_X(\beta_0(m_0), \gamma(m_0))] \\
&= \frac{1}{2}[(n_1 - m_0) + (n_2 - m_0) \\
&\quad + \mathrm{d}_X(\beta_0(m_0), \gamma(m_0)) - \mathrm{d}_X(\beta_0(n_1), \gamma(n_2))] \\
&= \frac{1}{2}[\mathrm{d}_X(\beta_0(n_1), \beta_0(m_0)) - \mathrm{d}_X(\beta_0(n_1), a) \\
&\quad + \mathrm{d}_X(\gamma(n_2), \gamma(m_0)) - \mathrm{d}_X(\gamma(n_2), b) \\
&\quad + \mathrm{d}_X(\beta_0(m_0), \gamma(m_0)) - \mathrm{d}_X(a, b)] \\
&\leq \frac{1}{2}[\delta + \delta + 2\delta] = 2\delta.
\end{aligned}$$

Applying the bound (2), it follows that

$$\begin{aligned}
(\beta_0(\infty) \cdot \gamma(\infty))_1 &\leq (\beta_0(m_0) \cdot \gamma(m_0))_1 + 4\delta \\
&= \frac{1}{2}[2m_0 - \mathrm{d}_X(\beta_0(m_0), \gamma(m_0))] + 4\delta \\
&\leq m_0 - \frac{1}{2}(r - \eta) + 4\delta.
\end{aligned}$$

This then implies the following lower bound on the distance between $\gamma(\infty)$ and $\beta_0(\infty)$:

$$\rho(\beta_0(\infty), \gamma(\infty)) \geq k_1 a^{-m_0 - 4\delta + \frac{r-\eta}{2}}.$$

We combine this inequality with (1), and a simple calculation reveals that our choice of r ensures that the path constructed between p and q avoids ΛH. The result follows. $\square$

LEMMA II.2.6 ([**6**, Lem. 1.21]). *If $R > 4\delta + Q + \max\{r + 4\delta + 1, K\}$, then the inclusion*

$$A_{r,R,K}(H) \hookrightarrow A_{r,\infty,K}(H)$$

induces a bijection between connected components.

PROOF. Surjectivity is obvious since $R \geq K$, so we need only show injectivity. Let $x, y \in C_K(H)$ lie in the same component of $A_{r,\infty,K}(H)$. We claim that the shortest path between x and y in $N_{r,\infty}(H)$ actually lies inside $N_{r,R}(H)$, which implies the result.

Let $\phi : [0, \ell] \to X$ be the shortest such path, parameterised by length. Then suppose that for some $s \in [0, \ell]$, we have that $\mathrm{d}_X(\phi(s), H) > R$. Let $[t_0, t_1] \subset [0, \ell]$ be the maximal subinterval containing s, such that for all $t \in [t_0, t_1]$, we have

$$\mathrm{d}_X(\phi(t), H) \geq r + 4\delta + 1.$$

We claim that $\phi|_{[t_0,t_1]}$ is an $(8\delta + 2)$-local geodesic. Indeed, for any $t \in [t_0, t_1]$, we have that the segment

$$\phi|_{[t-4\delta-1, t+4\delta+1] \cap [t_0,t_1]}$$

is contained in $N_{r+4\delta+1,\infty}(H)$. Therefore, any segment from $\phi(\max\{t_0, t - 4\delta - 1\})$ to $\phi(\min\{t_1, t + 4\delta + 1\})$ lies in $N_{r,\infty}(H)$. By the minimality of ϕ, it follows that $\phi|_{[t_0,t_1]}$ is an $(8\delta + 2)$-local geodesic. Apply [**19**, § III.H, 1.13], we then have that any geodesic path between $\phi(t_0)$ and $\phi(t_1)$ is contained in the 2δ-neighbourhood of $\phi|_{[t_0,t_1]}$. Maximality of $[t_0, t_1]$ means that either $t_0 = 0$, or $\mathrm{d}_X(\phi(t_0), H) = r + 4\delta + 1$. In both cases, it follows that $\mathrm{d}_X(\phi(t_0), H) \leq \max\{r + 4\delta + 1, K\}$. Identical reasoning also shows that $\mathrm{d}_X(\phi(t_1), H) \leq \max\{r + 4\delta + 1, K\}$.

Let s_0, s_1 be closest-point projections of $\phi(t_0)$, $\phi(t_1)$ respectively. Now, consider the quadrilateral $[s_0, s_1, \phi(t_1), \phi(t_0)]$. Any quadrilateral in X is 2δ-thin, so by inspecting Figure II.4 we see that in any case, there is a path from $\phi(s)$ to H of length at most $4\delta + Q + \max\{r + 4\delta + 1, K\} = R$. This contradicts our choice of s, and so the result follows. $\square$

In summary, the above lemmas give us the following.

PROPOSITION II.2.7. *Let G be a one-ended hyperbolic group with δ-hyperbolic Cayley graph X, and let H be a Q-quasi-convex subgroup. Then there exist computable values $r \leq K \leq R$, such that the map*

$$\pi_0(A_{r,R,K}(H)) \to \pi_0(A_{r,\infty,K}(H)) \xrightarrow{\mathcal{S}} \pi_0(\partial_\infty G \setminus \Lambda H)$$

is a well-defined bijection.

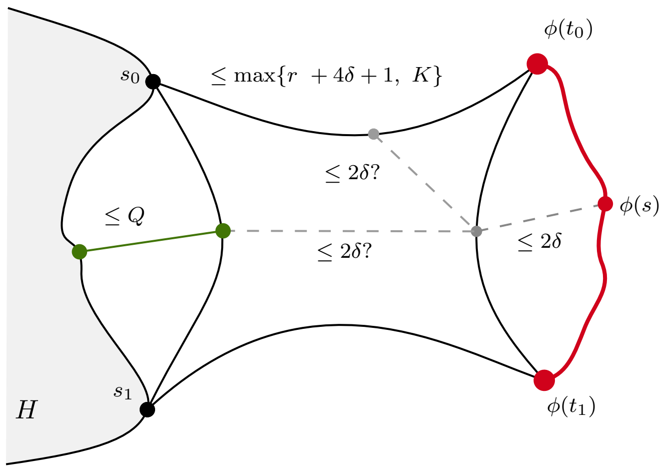

FIGURE II.4. The point $\phi(s)$ cannot be too far from H.

For the remainder of this chapter, we fix r, R, K such that the above is satisfied. We conclude this section with a final remark that H acts on $\pi_0(A_{r,R,K}(H))$ and $\pi_0(\partial_\infty G \setminus \Lambda H)$ by permutations, and the above bijection is easily seen to be equivariant with respect to this action.

II.2.2. Ends of pairs, revisited. We will conclude this section by using the above machinery to relate the number of (filtered) ends of a quasi-convex subgroup to its limit set complement. We begin with the following characterisation of filtered ends.

THEOREM II.2.8. *Let G be a one-ended hyperbolic group and H a quasi-convex subgroup. Then $\tilde{e}(G, H)$ is equal to the number of components of $\partial_\infty G \setminus \Lambda H$.*

PROOF. Let $N \in \mathbb{N} \cup \{\infty\}$ denote the number of components of $\partial_\infty G \setminus \Lambda H$. Choose a ray γ_i based in each component U_i of $\partial_\infty G \setminus \Lambda H$. If $i \neq j$, it is easy to see that there is no filtered homotopy between γ_i and γ_j. This proves that $\tilde{e}(G, H) \geq N$.

Conversely, fix a component U of $\partial_\infty G \setminus \Lambda H$. It is clear if γ is a ray such that $\gamma(\infty) \in U$, then γ is a filtered ray. Now, if γ' is another ray such that $\gamma'(\infty) \in U$, then by Lemma II.2.5 we have that U is path connected. So consider a path through U between $\gamma(\infty)$ and $\gamma'(\infty)$, then it is easy to see that this induces a filtered homotopy between γ and γ' in the Cayley complex. Thus, $\tilde{e}(G, H) \leq N$. □

Note that H acts on $\partial_\infty G \setminus \Lambda H$ by homeomorphisms. In particular, H permutes the connected components of this set. This gives us the following.

THEOREM II.2.9. *Let G be a one-ended hyperbolic group and H a quasi-convex subgroup. Then $e(G, H)$ is equal to the number of H-orbits of components of $\partial_\infty G \setminus \Lambda H$.*

PROOF. Recall by Proposition II.1.12 that $e(G, H)$ is equal to the number of ends of the coset graph $H\backslash X$. Note that each right coset Hg satisfies the following: $d_X(x, H)$ is constant across all

$x \in Hg$. We call this value the *height* of a coset. It is easy to see that at any given height there is only finitely many cosets.

Now, we will form the coset graph of H by a sequence of identifications, in a way which will make the conclusion clear. First, identify all cosets which lie in $N_r(H)$ to points. Next, identify all points in $N_{r,R}(H) \setminus A_{r,R,K}(H)$ which lie in the same coset. Now, in each component of $A_{r,\infty,K}(H)$, identify vertices which lie in the same coset of H. At this stage, it is clear we have a graph with $\tilde{\mathrm{e}}(G,H)$ ends. The only remaining identifications to be made are identifying some of these 'ends'. We have that the number of ends of the final graph is then equal to the number of H-orbits of components of $A_{r,\infty,K}(H)$, which is equal to the number of H-orbits of components of $\partial_\infty G \setminus \Lambda H$. But the number of ends of this graph is equal to $\mathrm{e}(G,H)$ by Proposition II.1.12, and so the result follows. □

The results of this section lead us to a straightforward proof of the following corollary.

Corollary II.2.10. *Let G be a one-ended hyperbolic group, and H a quasi-convex subgroup. Then $\mathrm{e}(G,H)$ is finite. In particular, $\mathrm{e}(G,H) \leq |B_{2l+R}(1)|$, where l is the longest length of a given generator of H.*

Proof. The action of H on the components of $\partial_\infty G \setminus \Lambda H$ is induced by the action of H on the components of $A_{r,R,K}(H)$. The latter is a locally finite graph on which H acts cocompactly, and so this must have finitely many H-orbits of connected components, since the quotient graph is finite. It follows that there can only be finitely many H-orbits of components of $\partial_\infty G \setminus \Lambda H$, and the result follows by Theorem II.2.9. The precise bound follows from just bounding above the size of a fundamental domain for the action of H on $A_{r,R,K}(H)$. □

The above results thus give us a full description of all possible H-almost invariant subsets of G, up to equivalence: any such set is equivalent to some union of H-orbits of components of $A_{r,\infty,K}(H)$. To ease notation, if Z is a union of connected components of $A_{r,\infty,K}(H)$, then let

$$Z^* := A_{r,\infty,K}(H) \setminus Z.$$

It is easy to see that the symmetric difference $Z^* \triangle (G \setminus Z)$ is H-finite, so this overloading of notation is not a problem for our purposes.

Now, note that if Z is a connected component of $A_{r,\infty,K}(H)$, then gZ is a connected component of $A_{r,\infty,K}(gH)$. The connected components of the latter are in one-to-one correspondence with the components of $\partial_\infty G \setminus \Lambda gH$, and we can define $\mathcal{S}(gZ)$ in the obvious way to realise this correspondence. Note that one has to be careful with the choice of basepoint, but this doesn't really matter as if we increase r accordingly for each g, then it is easy to see that the choice of basepoint does not affect the properties of $\mathcal{S}$. We conclude this section by characterising via the boundary what it means for an H-almost invariant set X to cross itself, in the sense of Definition II.1.4.

PROPOSITION II.2.11. *Let $X \subset G$ be an H-almost invariant set, and let Z be the unique union of connected components of $A_{r,\infty,K}(H)$ which is equivalent to X. Then X crosses itself if and only if there exists $g \in G$ such that*

$$\mathcal{S}Z \cap \mathcal{S}(gZ^*), \quad \mathcal{S}Z \cap \mathcal{S}(gZ), \quad \mathcal{S}(Z^*) \cap \mathcal{S}(gZ), \quad \mathcal{S}(Z^*) \cap \mathcal{S}(gZ^*)$$

are all non-empty.

PROOF. It can be seen that $Z \cap gZ$ is H-infinite if and only if it contains arbitrarily large balls [**105**, Rmk. 1.13]. This clearly implies that $\mathcal{S}Z \cap \mathcal{S}(gZ) \neq \emptyset$.

Conversely, $\mathcal{S}Z \cap \mathcal{S}(gZ)$ is an open subset of $\partial_\infty G$, so if it is non-empty then it contains some open ball $B \subset \mathcal{S}Z \cap \mathcal{S}(gZ)$. It is follows quickly from this that $Z \cap gZ$ must contain balls of arbitrary diameter. By symmetry in the other four cases, the result follows. □

II.3. Detection via digraphs

So far, we have related the connectivity of the limit set complement to the connectivity of a certain subgraph of the Cayley graph, but this subgraph is still infinite. In order to obtain any sort of algorithm we will need some way of understanding the global connectivity properties by looking only at a finite piece.

In this section, we will make headway towards this goal by constructing a certain *labelled digraph*, whose connectivity and language encodes a great deal of information about the properties of the subgraph $A_{r,R,K}(H)$.

II.3.1. Digraphs and their languages. We begin by briefly introducing labelled digraphs. These have many applications to the study of subgroups of free groups, and a good survey of their rich theory can be found in [**81**].

Now could be a good time to review the graph-theoretical conventions set out in Section I.5.3.

DEFINITION II.3.1 (Labelled digraphs). Let S be a finite set of symbols, closed under taking formal inverses, i.e. $S = S^{-1}$. An S-*digraph* Θ is a finite graph where oriented edge $e \in \vec{E}(\Theta)$ is labelled by some $s \in S$, such that if e labelled by $s \in S$ then e^{-1} is labelled by s^{-1}.

Note that in this definition we allow the possibility of single-edge loops, and we do not require our graphs to be connected. We also do not require these graphs to be *folded*, in the sense of STALLINGS. That is, a vertex may have multiple outgoing edges with the same label. We now set up some notation.

Denote by $\mathrm{Lab}(e)$ the label of the oriented edge e, so $\mathrm{Lab}(e^{-1}) = \mathrm{Lab}(e)^{-1}$. Given a set of symbols S, denote by S^* the set of finite strings in S. If $p = e_1 \dots e_n$ is a path through Θ, define its label $\mathrm{Lab}(p) \in S^*$ as the formal string $\mathrm{Lab}(p) = \mathrm{Lab}(e_1) \dots \mathrm{Lab}(e_n)$. We say that two combinatorial paths p, q are *freely equal* if $\iota(p) = \iota(q)$, $t(p) = t(q)$, and the words $\mathrm{Lab}(p)$ and $\mathrm{Lab}(q)$ are equal in the free group $F(S)$. We have the following definition.

Definition II.3.2. Let Θ be an S-digraph, and let v, v' be vertices of Θ. Define the *language of* Θ *from* v *to* v' as the set

$$\mathcal{L}(\Theta, v, v') = \{\mathrm{Lab}(p) \ : \ p \text{ is a path in } \Theta \text{ from } v \text{ to } v'\} \subset S^*.$$

Let H be a group and $S \subset H$ a finite symmetric subset. Denote by $\pi : S^* \to H$ the obvious projection. Then we have the following key lemma.

Lemma II.3.3. *Let* Θ *be an* S*-digraph, where* S *is a finite subset of a group* H*. Then for every pair* v_0, v_1 *of vertices of* Θ*, there is a finitely generated subgroup* $K_{v_0} \leq H$ *and a finite subset* $T_{v_0,v_1} \subset H$ *such that*

$$\pi(\mathcal{L}(\Theta, v_0, v_1)) = \bigcup_{t \in T_{v_0,v_1}} K_{v_0} t.$$

Moreover, T_{v_0,v_1} *and generators of* K_{v_0} *can be computed effectively.*

Proof. To ease notation, let $\mathcal{L} = \mathcal{L}(\Theta, v_0, v_1)$. We first construct K_{v_0} and T_{v_0,v_1}. Consider the following sets of paths through Θ.

- $\mathscr{P}(v, v') = \{p : p \text{ is a simple path through } \Theta \text{ with } \iota(p) = v,\ t(p) = v'\}$,
- $\mathscr{L}(v) = \{\ell : \ell \text{ is a simple loop through } \Theta \text{ with } \iota(\ell) = t(\ell) = v\}$,
- $\mathscr{Q}(v) = \{p\ell p^{-1} : v' \in V(\Theta),\ \ell \in \mathscr{L}(v'),\ p \in \mathscr{P}(v, v')\}$.

Intuitively, $\mathscr{Q}(v)$ is the set of 'lollipop loops' based at v, which are formed by attaching a simple loop to a simple path and its inverse. Let $Y = \{\pi(\mathrm{Lab}(\gamma)) : \gamma \in \mathscr{Q}(v_0)\}$, and set $K_{v_0} = \langle Y \rangle$. Also define

$$T_{v_0,v_1} = \{\pi(\mathrm{Lab}(p)) : p \in \mathscr{P}(v_0, v_1)\}.$$

The rest of this proof will be dedicated to showing the following statement. We claim that every path γ through Θ with $\iota(\gamma) = v_0$ is freely equal to a path of the form $\alpha_1 \dots \alpha_r q$, where each $\alpha_i \in \mathscr{L}(v_0)$ and $q \in \mathscr{P}(v_0, t(\gamma))$. From this the lemma follows immediately.

With the above claim in mind, let γ be a path in Θ from v_0 to v_1 and assume without loss of generality that $\mathrm{Lab}(\gamma)$ is freely reduced. If γ is a simple path or simple loop then the statement is trivial, so assume γ self-intersects somewhere away from $\iota(\gamma) = v_0$. We then proceed by induction on the length of γ. We may decompose γ into the form

$$\gamma = p_0 \ell_1 p_1 \ell_2 p_2 \dots \ell_k p_k,$$

where the p_j are (possibly trivial) simple paths such that $t(p_j) = \iota(p_{j+1})$ and each ℓ_j is a non-trivial simple loop in $\mathscr{L}(\iota(p_j))$. For each $1 \leq j \leq k$ let

$$\ell_j' := p_0 p_1 \dots p_{j-1} \ell_j p_{j-1}^{-1} \dots p_1^{-1} p_0^{-1}.$$

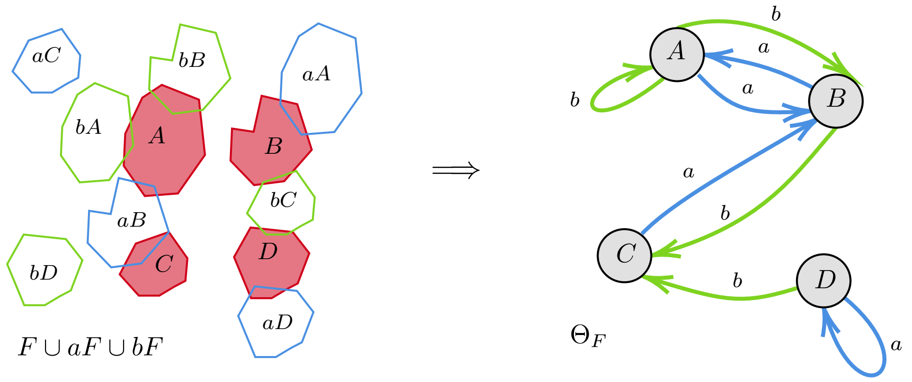

FIGURE II.5. An example of the construction of an adjacency digraph. Here, $\pi_0(F) = \{A, B, C, D\}$ and $S_F = \{a, b, a^{-1}, b^{-1}\}$. Due to symmetry we have not drawn the translates of F by a^{-1} or b^{-1}, nor the corresponding edges of Θ_F.

Inspection reveals then that γ is freely equal to the path $\gamma' := \ell_1'\ell_2'\dots\ell_k'p_0p_1\dots p_k$. Consider the subpath $\gamma_j = p_0\dots p_j$. Each γ_j is a path of strictly shorter length than γ. By the inductive hypothesis we have that γ_j is of the form $s_1\dots s_k q$ where each $s_i \in \mathscr{Q}(v_0)$ and $q \in \mathscr{P}(v_0, t(p_j))$. Now we have that

$$\ell_{j+1}' = \gamma_{j-1}\ell_j\gamma_{j-1}^{-1} = s_1\dots s_i q\ell_j q^{-1}s_i^{-1}\dots s_1^{-1}.$$

Since $q\ell_j q^{-1} \in \mathscr{Q}(v_0)$, we have that ℓ_j' is thus freely equal to a path of the form $\alpha_1\dots\alpha_r$ where each $\alpha_i \in \mathscr{Q}(v_0)$. If we apply the inductive hypothesis once more to γ_k, then the claim follows. □

II.3.2. Adjacency digraphs. Let G be a finitely generated group and X a Cayley graph of G. Let H be a finitely generated subgroup of G. Let $F \subset X$ be a finite subgraph, and consider $X = HF$. Then X is a subgraph of X on which H acts cocompactly. We say that F is a *finite H-cover* of X. In principle, X will have many finite H-covers.

Given such a finite set of vertices F, we form its *H-adjacency set*, or just its *adjacency set* S_F as follows. Let

$$S_F = \{s \in H \ : \ s \neq 1,\ sF \cap F \neq \emptyset\}.$$

Informally, S_F is the finite set containing $s \in H$ such that sF intersects F. Note that S_F is symmetric, i.e $S_F = S_F^{-1}$. The fact that S_F is finite follows from F being finite and X being locally finite.

We form the S_F-digraph Θ_F, called the *F-adjacency graph*, as follows. The vertex set $V(\Theta_F)$ is precisely the set $\pi_0(F)$ of connected components of F, and there is a directed edge labelled by $s \in S_F$ from v to v' if $v \cap sv' \neq \emptyset$. In particular, this implies that v and sv' are contained in the same connected component of X. It is easy to check that Θ_F is indeed a well-defined digraph. See Figure II.5 for an example of this construction.

The relevance of the adjacency digraph Θ_F to our problem is given by the following lemma.

Lemma II.3.4. *Let $v, v' \in \pi_0(F)$, and $h \in H$. Then v and hv' are contained in the same connected component of $X = HF$ if and only if $h \in \pi(\mathcal{L}(\Theta_F, v, v'))$.*

Proof. Firstly, suppose that $h = \pi(w)$, where $w = s_1 \dots s_n \in \mathcal{L}(\Theta_F, v, v')$. Let $\gamma = e_1 \dots e_n$ be a path through Θ_F between v and v' labelled by w, so e_i is labelled by s_i. Let $v = v_0$ and $v_i = t(e_i)$ for each $1 \leq i \leq n$. Observe that by definition we have that v_i intersects $s_{i+1}v_{i+1}$ in X, and so $s_{i+1}v_{i+1}$ and v_i must be contained within the same component of X. Since $w \in \mathcal{L}(\Theta_F, v, v')$, we have that $v_n = v'$. Note also that $s_{i+1} \dots s_n \in \mathcal{L}(\Theta_F, v_i, v_n)$. We proceed by induction on n. Clearly if $n = 1$, then the result is true by the definition of Θ_F.

Suppose then that $n > 1$. Then $w' = s_2 \dots s_n \in \mathcal{L}(\Theta_F, v_2, v_n)$. Let $h' = \pi(w')$, then by the inductive hypothesis we have that v_1 and $h'v_n$ are contained within the same component of X. Then v_0 is contained within the same component as $s_1 v_1$ by definition. But this is contained in the same component as $s_1 h' v_n = hv'$.

Conversely, suppose that v and hv' lie in the same component of X for some $h \in H$. Let γ be a path through X between v and hv'. The path γ will pass through a sequence of translates of components of F, say

$$v = h_0 v_0,\ h_1 v_1,\ h_2 v_2,\ \dots,\ h_l v_l = hv',$$

where $h_i \in H$, $h_0 = 1$, $v_i \in \pi_0(F)$, and $h_{i-1}v_{i-1} \cap h_i v_i \neq \emptyset$ for every $1 \leq i \leq n$. By definition, we then have that $s_i := h_{i-1}^{-1} h_i \in S_F$, and there is an edge in Θ_F from v_{i-1} to v_i labelled by s_i. We thus deduce that $s_1 \dots s_n \in \mathcal{L}(\Theta_F, v, v')$. Since $h = s_1 \dots s_n$, the result follows. $\square$

The next result is a characterisation of when the graph $X = HF$ has only finitely many components. Given $v, v' \in \pi_0(F)$, recall the definition of K_v and $T_{v,v'}$ from Lemma II.3.3.

Lemma II.3.5. *Let F be a finite H-cover of $X = HF$. Then X has finitely many connected components if and only if for every $v \in \pi_0(F)$ we have that K_v has finite index in H.*

Proof. By Lemmas II.3.3 and II.3.4, note that $h \in K_v$ if and only if there is a path through X from v to hv, and thus the same also holds for $h'v$ and $h'hv$ for any $h' \in H$. So suppose first that some K_v has finite index in H and let T be a finite left-transversal of K_v. Then $TK_v v = Hv$, and this intersects at most $|T|$ components of X. Repeat this for every $v \in \pi_0(F)$ and see that X has at most

$$\sum_{v \in \pi_0(F)} |H : K_v| < \infty$$

components.

Conversely, fix $v \in \pi_0(F)$ such that $|H : K_v| = \infty$ and let $T = \{h_i : i \in \mathbf{N}\}$ be an infinite left-transversal in H of K_v. We claim that every $h_i v$ lies in a distinct component of X. Indeed, suppose that $h_i v$ and $h_j v$ were contained in the same component of X for some $i \neq j$. Then by

Lemmas II.3.3 and II.3.4 we have that $h_i^{-1}h_j \in K_v$. This contradicts our choice of T, so the lemma follows. $\square$

It is clear that H permutes the connected components of X. Note that if X is connected then Θ_F is certainly connected, though the converse is not necessarily true. Instead, the number of connected components of Θ_F actually encodes the following.

Lemma II.3.6. *The number of connected components of Θ_F is equal to the number of H-orbits of components of $X = HF$.*

Proof. By Lemma II.3.4 we see that $v, v' \in \pi_0(F)$ are connected by a path in Θ_F if and only if v is joined by a path through X to some H-translate of v'. The lemma then follows immediately. $\square$

II.3.3. First algorithms. The first application of our digraph machinery is the following. Throughout this subsection we fix constants $0 \leq r \leq K \leq R < \infty$ and a Cayley graph X of our one-ended hyperbolic group G.

Proposition II.3.7. *Let H be a quasi-convex subgroup of a one-ended hyperbolic group G. Then there is an algorithm which, given $x \in G$, will decide if $x \in A_{r,R,K}(H)$.*

Proof. First, we show that one can decide membership of $N_{r,R}(H)$. Given $x \in G$, compute the finite balls $U_1 = B_r(x)$, $U_2 = B_R(x)$ in X. Then $x \in N_{r,R}(H)$ if and only if $U_1 \cap H = \emptyset$ and $U_2 \cap H \neq \emptyset$. This is decidable, since membership of H is decidable. Similarly, we can decide membership of $C_K(H)$. If x is not in $N_{r,R}(H)$, then terminate and return 'no'.

Now, note that H acts cocompactly on $N_{r,R}(H)$, so compute a finite H-cover F containing x. This can be achieved by, for example, letting l be the length of the longest given generator of H, and choosing $F = N_{r,R}(H) \cap B_{2l+R}(x)$. Let $v_x \in \pi_0(F)$ be the component of F containing x. We now form the adjacency digraph Θ_F, and mark the vertices $v \in \pi_0(F)$ which intersect $C_K(H)$ in X. One can then check that $x \in A_{r,R,K}(H)$ if and only if v_x lies in a connected component of Θ_F containing a marked vertex. $\square$

We now have the following algorithms, which will allow us to distinguish (filtered) ends from one another.

Proposition II.3.8. *There is an algorithm which, given $x, y \in A_{r,R,K}(H)$, will decide if x and y lie in distinct H-orbits of connected components of $A_{r,R,K}(H)$.*

Proof. Use Proposition II.3.7 to compute a finite H-cover F of $A_{r,R,K}(H)$, and let $v_x, v_y \in \pi_0(F)$ be such that $x \in v_x$, $y \in v_y$. Then, as remarked in the proof of Lemma II.3.6, we need only form the adjacency digraph Θ_F and check whether v_x and v_y lie in the same connected component of Θ_F. $\square$

Proposition II.3.9. *There is an algorithm which, given $x, y \in A_{r,R,K}(H)$ and a solution to the generalised word problem for H, will decide if x and y lie in distinct connected components of $A_{r,R,K}(H)$.*

Proof. Find a finite H-cover F for $A_{r,R,K}(H)$ containing both x and y. Let $v_x, v_y \in \pi_0(F)$ be the components of F containing x and y respectively. By Lemma II.3.4, we have that x and y are contained in the same component if and only if $1 \in \pi(\mathcal{L}(\Theta_F, v_x, v_y))$. Form the subgroup K_{v_x} and set of words T_{v_x,v_y} as in Lemma II.3.3, then x and y lie in different components if and only if $T_{v_x,v_y} \cap K_{v_x}$ is empty. Using our given solution to the generalised word problem in H, this is decidable. □

We now turn to reproving Proposition II.3.9, but we drop the hypothesis that H has a solvable generalised word problem and replace it with the condition that H has finitely many filtered ends in G. The key observation that makes this problem tractable in this case is the following.

Proposition II.3.10. *If $A_{r,R,K}(H)$ has finitely many components, then there is an algorithm which, given $x, y \in A_{r,R,K}(H)$, will decide if x and y lie in distinct connected components of $A_{r,R,K}(H)$.*

Proof. It is known by Lemma II.3.5 that, for every $v \in \pi_0(F)$, the subgroup $K_v \leq H$ constructed in Lemma II.3.3 has finite index in H. In particular, K_v is quasi-convex in H, and so we can decide membership of K_v. The algorithm then proceeds exactly as in Proposition II.3.9. □

II.4. Main results

II.4.1. Counting (filtered) ends of pairs. We now give algorithms to compute $\mathrm{e}(G, H)$ and $\tilde{\mathrm{e}}(G, H)$. Fix constants $0 \leq r \leq K \leq R < \infty$ such that Proposition II.2.7 is satisfied. Let F be a finite H-cover for $A_{r,R,K}(H)$, and form the adjacency digraph Θ_F. Note that the construction of this digraph is completely effective. We immediately have the following new algorithm for computing $\mathrm{e}(G, H)$, which was first shown to be computable by Vonseel [**134**].

Theorem II.4.1. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, will output $\mathrm{e}(G, H)$.*

Proof. One simply counts connected components of Θ_F. The construction of Θ_F is completely effective, so the result follows immediately from Lemma II.3.6 and Theorem II.2.9. □

Computing $\tilde{\mathrm{e}}(G, H)$ poses more problems. In particular, we need to somehow be able to decide if K_v has finite index in H, which can be seen to be undecidable for an arbitrary choice of H via the Rips construction (see e.g. [**8**]). Thus, in light of Lemma II.3.5 we cannot expect to be able to decide if $\tilde{\mathrm{e}}(G, H) = \infty$ without either adding further hypotheses to H (e.g. this is decidable if

H is free [**81**]), or somehow further controlling the structure of $A_{r,R,K}(H)$. It is not clear whether the latter of these is even possible, which suggests this problem may be undecidable for arbitrary choices of quasi-convex $H \leq G$. We can however at least give the following two algorithms.

THEOREM II.4.2. *There is an algorithm which, upon input of a one-ended hyperbolic group G, generators of a quasi-convex subgroup H, a solution to the generalised word problem for H, and an integer $N \geq 0$, will decide whether $\tilde{e}(G, H) \geq N$. In particular, we can decide if $\partial_\infty G \setminus \Lambda H$ is connected.*

PROOF. Compute a finite H-cover F_0 of $A_{r,R,K}(H)$, then inductively define $F_{i+1} := YF_i \cup F_i$, where Y is a symmetric generating set for H. Thus we have an increasing sequence of H-covers (F_i), where each strictly contains the last. For each i, let N_i denote the number of components of $A_{r,R,K}(H)$ which intersect F_i. This number is computable by Proposition II.3.9, and we can conclude that $e(G, H) \geq N_i$. If there is some i such that $N_i = N_{i+1}$, then since Y is a generating set, it follows that $N_j = N_i$ for all $j > i$. Given N as input, our algorithm will run until $N_i \geq N$ for some i, or terminate if the sequence $(N_i)_i$ stabilises. By the above, this will always halt. □

THEOREM II.4.3. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, will terminate if and only if $\tilde{e}(G, H)$ is finite. Moreover, upon termination it will output the value of $\tilde{e}(G, H)$.*

PROOF. We proceed as before and compute a finite H-cover F_0 of $A_{r,R,K}(H)$, then inductively define $F_{i+1} := YF_i \cup F_i$, where Y is a symmetric generating set for H. Thus, we have an increasing sequence of H-covers

$$F_0 \subset F_1 \subset F_2 \subset \dots,$$

where each strictly contains the last. Moreover, $\bigcup_i F_i = A_{r,R,K}(H)$. For each $i \geq 0$ we now run the following search. For each component of YF_i, search for a path through $A_{r,R,K}(H)$ into F_i. If this process terminates for a given $i \geq 0$ then it follows from an easy induction argument that there is a path from any point in $A_{r,R,K}(H) = HF_i$ back to F_i travelling through $A_{r,R,K}(H)$. In particular, this means that every connected component of $A_{r,R,K}(H)$ intersects F_i.

Clearly such an F_i exists if and only if $A_{r,R,K}(H)$ has finitely many components, which is equivalent to the condition that $\tilde{e}(G, H) < \infty$ by Proposition II.2.7 and Theorem II.2.8. If we do find such an F_i then to compute the exact value of $\tilde{e}(G, H)$ we may use Proposition II.3.10 to decide how many distinct components of $A_{r,R,K}(H)$ intersect F_i. By our choice of F_i, this will then be precisely the total number of components of $A_{r,R,K}(H)$. □

II.4.2. Splitting detection. In this section we apply the above tools to the problem of deciding if a given quasi-convex subgroup is associated to a splitting. In short, we run two searches in parallel – one search for a splitting and another search for obstructions to splittings. The first step in searching for splittings over subgroups commensurable with H is to be able to recognise such subgroups. We will achieve this by deciding membership of $\mathrm{Comm}_G(H)$.

PROPOSITION II.4.4. *Let G be a hyperbolic group, then there is an algorithm which, on input of generators of a quasi-convex subgroup H, will decide if $|G : H| < \infty$.*

PROOF. We have that $|G : H|$ is finite if and only if $\partial_\infty G = \Lambda H$. By Proposition II.2.7 this is true if and only if $A_{r,R,K}(H) = \emptyset$ for suitably chosen r, R, K. This can be decided by computing a finite H-cover F of $N_{r,R}(H)$ as in the proof of Proposition II.3.7 and checking if F intersects $C_K(H)$. □

PROPOSITION II.4.5. *Let G be a hyperbolic group. Given generators of a quasi-convex subgroup $H \leq G$, then membership of the commensurator $\mathrm{Comm}_G(H)$ is decidable.*

PROOF. Let $g \in G$, then since G is hyperbolic we have that H^g is quasi-convex. Moreover, $H \cap H^g$ is quasi-convex and we can compute an explicit generating set for this group via [**64**]. We then use Proposition II.4.4 to decide if $|H : H \cap H^g|$ and $|H^g : H \cap H^g|$ are finite. This decides whether $g \in \mathrm{Comm}_G(H)$. □

Note that $\mathrm{Comm}_G(H)$ is itself quasi-convex, and so given a generating set of this subgroup we would have that the membership problem would be decidable via Kapovich's algorithm. However, we are not given generators of $\mathrm{Comm}_G(H)$, but of H. So, what the above proposition tells us is that we can decide membership of the commensurator in spite of this problem. Applying this observation, we produce the following algorithm which searches for splittings where the edge group is commensurable with H.

PROPOSITION II.4.6. *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, and terminates if and only if H is associated to a splitting.*

PROOF. Enumerate presentations of G via Tietze transformations. If a given presentation has the general form of an amalgam or HNN extension, run Kapovich's algorithm [**79**] on the generators of the edge group, which terminates if and only if this subgroup is quasi-convex and outputs a quasi-convexity constant Q if it does terminate. This procedure enumerates splittings of G over quasi-convex subgroups.

Given a particular quasi-convex splitting of G, say over H', we can decide if H' is commensurable with H as follows. Using Proposition II.4.5 we decide if $H' \leq \mathrm{Comm}_G(H)$ and $H \leq \mathrm{Comm}_G(H')$. It is easy to check these two relations hold if and only if H is commensurable with H'. This completes the algorithm. $\square$

Recall Proposition II.2.11, which characterised crossings via intersections of shadows in $\partial_\infty G$. We now characterise these intersections via local geometry, and present an algorithm which terminates if and only if such a crossing exists. We first need the following technical lemma.

Lemma II.4.7. *Let U_1, U_2 be unions of connected components of $A_{r,\infty,K}(H)$. The intersection $\mathcal{S}U_1 \cap \mathcal{S}(gU_2)$ is non-empty if and only if there exists some $x \in U_1 \cap gU_2$ such that*

$$\mathrm{d}_X(x, H) > K \quad \text{and} \quad \mathrm{d}_X(x, gH) > K + 5\delta + |g|.$$

Proof. Firstly, suppose that such an x exists, then let γ be a ray based at 1 passing within $c = 3\delta$ of x, as in Lemma II.1.17. As in the proof of Lemma II.2.3 we see that $\gamma(\infty) \in \mathcal{S}U_1$ since γ passes within c of $C_K(H)$. Secondly, let γ' be a geodesic ray based at g such that $\gamma'(\infty) = \gamma(\infty)$. Then the Hausdorff distance between γ and γ' is at most $5\delta + |g|$ (apply e.g. [**43**, Exc. 11.86]). Again, as in the proof of Lemma II.2.3 we see that $\gamma'(\infty) \in \mathcal{S}(gU_2)$, and we're done.

Conversely, let $\gamma \in p \in \mathcal{S}U_1 \cap \mathcal{S}(gU_2)$. By the definition of $\mathcal{S}$ we have that there is some t_0 such that for all $t \geq t_0$, $\gamma(t) \in U_1 \cap gU_2$. Moreover, by Lemma II.2.1 we have for $i = 1, 2$ that $\mathrm{d}_X(\gamma(t), H_i) \to \infty$ as $t \to \infty$. Thus, by setting $x = \gamma(t)$ for some sufficiently large t, we are done. $\square$

We're now ready to search for crossings. This algorithm will check every choice of H-almost invariant set and search for any crossings. It will terminate if and only if it finds a crossing for every such choice. Note that since $\mathrm{e}(G, H) < \infty$ by Corollary II.2.10, there is only finitely many possible H-almost invariant sets to check, up to equivalence.

Proposition II.4.8. *There exists an algorithm which, on input of a one-ended hyperbolic group G and generators of a quasi-convex subgroup H, will terminate if and only if for every choice $U \subset G$ of H-almost invariant set, we have that U is not semi-nested.*

Proof. We begin by picking representative choices for every equivalence class of non-trivial H-almost invariant subsets $U_1, \ldots, U_n$. In particular, each U_i is a union of H-orbits of connected components of $A_{r,\infty,K}(H)$. For notational convenience we identify U_i^* with its representative in this list.

Enumerate elements $g \in G \setminus \mathrm{Comm}_G(H)$ via Proposition II.4.5. For each $i = 1, \ldots, n$ search for some $x_{i,1} \in U_i \cap gU_i$, $x_{i,2} \in U_i \cap g(U_i^*)$, $x_{i,3} \in U_i^* \cap gU_i$, and $x_{i,4} \in U_i^* \cap g(U_i^*)$ such that the conditions in Lemma II.4.7 are met for each $x_{i,j}$. We terminate our search if and only if we find such an $x_{i,j}$ for every i, j. By Proposition II.2.11 and Lemma II.4.7 this will terminate if and only if every H-almost invariant subset is not semi-nested. $\square$

We now present the final result of this chapter, an algorithm to detect splittings over quasi-convex subgroups. We split this algorithm into two cases, and firstly we consider the situation that we know *a priori* that our subgroup has finitely many filtered ends.

THEOREM II.4.9. *There is an algorithm which, upon input of a one-ended hyperbolic group G and generators of a quasi-convex subgroup H such that $\tilde{\mathrm{e}}(G, H) < \infty$, will decide whether H is associated to a splitting. Furthermore, if such a splitting exists then the algorithm will output this splitting.*

PROOF. Let $U_1, \ldots, U_n$ be the components of $A_{r,\infty,K}(H)$, and let $H' \leq H$ be a finite index subgroup which fixes each individual component, so each U_i is an H'-almost invariant subset. Any other subgroup which is commensurable with H will have the same set of filtered ends, so this is the 'finest' set of H''-almost invariant subsets for any subgroup H'' commensurable with H. It therefore follows that if H' does not admit a semi-nested H'-almost invariant set, then neither does any other subgroup which is commensurable with H.

With the above in mind we run two algorithms in parallel. We search for a splitting over a subgroup commensurable with H via Proposition II.4.6, and concurrently run the algorithm in Proposition II.4.8 on H'. By the above discussion exactly one of these will terminate, and if the former algorithm terminates then it will output a presentation of a splitting over a subgroup commensurable with H. $\square$

Indeed, if $\tilde{\mathrm{e}}(G, H)$ is not finite then we cannot rely on the machinery used above, as the stabiliser of some component of $\partial_\infty G \setminus \Lambda H$ may have infinite index in H. Moreover, referring back to the discussion in Section II.4.1, we are unlikely to be able to decide if $\tilde{e}(G, H)$ is finite for arbitrary quasi-convex H, at least with just the current tools presented in this chapter.

Recall that a subgroup H in G is said to be lonely if there is no subgroup $H' \neq H$ such that H is commensurable to H'. For a quasi-convex subgroup H of a hyperbolic group G, this condition is equivalent to saying that $H = \mathrm{Comm}_G(H)$ and H has no finite quotients.

THEOREM II.4.10. *There is an algorithm which, upon input of a one-ended hyperbolic group G, generators of a quasi-convex subgroup H, and knowledge of whether H is lonely in G, will decide whether H is associated to a splitting. Furthermore, if such a splitting exists then the algorithm will output this splitting.*

PROOF. Firstly, begin running the algorithm from Proposition II.4.6, which will terminate if and only if H is associated to a splitting.

Concurrently we run the following. If H is not lonely, then simultaneously search for an element $g \in \mathrm{Comm}_G(H) \setminus H$ and a finite index subgroup H' of H. If we find the former then continue, and if we find the latter then replace H with H' and then continue. At least one of these will terminate, and this ensures that $H \neq \mathrm{Comm}_G(H)$. If H is lonely, then just continue. We now run the algorithm presented in Proposition II.4.8. By Proposition II.1.26, exactly one of these two procedures will terminate, and if the first algorithm terminates then it will output a presentation of a splitting over a subgroup commensurable with H. $\square$

Corollary II.4.11. *There is an algorithm which takes in as input a one-ended hyperbolic group G and generators of a quasi-convex, residually finite subgroup H. This algorithm will then decide if H is associated to a splitting, and will output such a splitting if one exists.*

It is conjectured that the problem of deciding whether a given hyperbolic group has a finite quotient is undecidable, and in fact this problem is known to be equivalent to the well-known conjecture that there exists a hyperbolic group which is not residually finite [**20**]. Assuming this conjecture, it would be undecidable whether a given quasi-convex subgroup is lonely. This means that the hypothesis in Theorem II.4.10 is likely necessary unless we place further restrictions on H, such as requiring H be residually finite.

CHAPTER III

Big tiles in acylindrically hyperbolic groups

For some minutes Alice stood without speaking, looking out in all directions over the country—and a most curious country it was. There were a number of tiny little brooks running straight across it from side to side, and the ground between was divided up into squares by a number of little green hedges, that reached from brook to brook. "I declare it's marked out just like a large chess-board!" Alice said at last.

— L. Carroll, *Through the Looking Glass*

This chapter is concerned with a long-standing open problem in group theory, which asks whether every group is *monotileable*. We present progress on this question by proving that every acylindrically hyperbolic group is monotileable. Our methods extend those of Akhmedov, who proved a similar result for hyperbolic groups [**2**]. This chapter is based on joint work with L. Mineh [**92**].

III.1. Preliminaries

III.1.1. Monotileable groups. We first introduce some basic definitions and properties related to tiles and monotileable groups.

Definition III.1.1 (Monotileable group). Let G be a group. A subset $T \subset G$ is called a *tile* if G can be expressed as a disjoint union of left translates of T. In other words, T is a tile if there exists a subset $C \subset G$ such that

$$G = \bigsqcup_{g \in C} gT.$$

The set C is called the *centre* of tiling. The group G is called *monotileable* if every finite subset of G is contained in a finite tile.

In [**29**], Chou proved that many classes of groups are monotileable. Most of these were also independently recognised by Weiss [**136**].

Theorem III.1.2 (Chou). *The class of monotileable groups contains all of the following.*

1. *Finite groups.*
2. *Abelian groups.*
3. *Monotileable-by-monotileable groups.*
4. *Elementary amenable groups.*

(5) *Fully-residually*[1] *monotileable groups.*

In particular, residually finite groups are monotileable.

III.1.2. Acylindrical actions. Next, we state some basic definitions related to acylindrically hyperbolic groups.

Definition III.1.3. Let G be a group acting by isometries on a geodesic hyperbolic space X. The action of G is called *acylindrical* if for every $\varepsilon > 0$ there are constants $R, N \geq 0$ such that the set

$$\{g \in G \mid \mathrm{d}(x, gx) < \varepsilon, \mathrm{d}(y, gy) < \varepsilon\}$$

has at most N elements whenever $x, y \in X$ are such that $\mathrm{d}(x, y) \geq R$. An element $g \in G$ is called *elliptic* if it acts on X with bounded orbits. An element $g \in G$ is called *loxodromic* if the map $\mathbf{Z} \to X$ given by $n \mapsto g^n x$ is a quasi-isometric embedding for any $x \in X$.

Note that a loxodromic element $g \in G$ fixes exactly two points of $\partial_\infty X$, which we write as g^∞ and $g^{-\infty}$. Two loxodromic elements $g, h \in G$ are said to be *independent* if the sets $\{g^\infty, g^{-\infty}\}$ and $\{h^\infty, h^{-\infty}\}$ are disjoint.

Acylindrical actions and their individual isometries have been classified. If G acts acylindrically and the space X is hyperbolic, then every element is either elliptic or loxodromic [**17**, Lem. 2.2]. The actions are known to fall into the trichotomy described below.

Theorem III.1.4 ([**107**, Thm. 1.1]). *Let G be a group acting acylindrically on a geodesic hyperbolic space. Then either*

(1) *G has bounded orbits,*
(2) *G is virtually cyclic and contains a loxodromic element, or*
(3) *G contains infinitely many independent loxodromic elements.*

An action satisfying (3) in Theorem III.1.4 is called *non-elementary.* We now have the following core definition.

Definition III.1.5 (Acylindrically hyperbolic group). A group G is said to be *acylindrically hyperbolic* if it admits a non-elementary acylindrical action on a hyperbolic geodesic metric space.

We will use the following characterisation of acylindrical hyperbolicity.

Theorem III.1.6 ([**107**, Thm. 1.2]). *Let G be an acylindrically hyperbolic group. Then there exists a (possibly infinite) generating set S of G such that the Cayley graph $X = \mathrm{Cay}(G, S)$ is hyperbolic and the action of G upon X is non-elementary and acylindrical .*

[1] Given a group property Φ, recall that a group is said to be *fully-residually* Φ if for any finite subset $\{g_1, \dots, g_n\} \subset G$, there exists a homomorphism $f : G \to Q$, where Q is some group satisfying Φ, such that $f(g_i) \neq 1$ for each $i = 1, \dots, n$. It is an easy exercise to show that residually finite groups are fully-residually finite.

We will need the following fact about existence of loxodromic elements with particular endpoints.

THEOREM III.1.7 ([**73**, Thms. 2.9, 4.5]). *Let G be a group acting acylindrically and non-elementarily on a hyperbolic space X. Then for any pair of disjoint nonempty open sets $U, V \subset \Lambda G$, there is a loxodromic element $g \in G$ with $g^{\infty} \in U$ and $g^{-\infty} \in V$.*

III.2. Fixed points at infinity

Throughout this section we take X to be a δ-hyperbolic geodesic metric space and G a group with an acylindrical, non-elementary action on X by isometries. The goal of this section is to study fixed subsets of $\partial_\infty X$ under the induced action of G by homeomorphisms. We introduce the following notation. Given $g \in G$, we write

$$\operatorname{fix}(g) = \{p \in \partial_\infty X : g \cdot p = p\}.$$

The following two lemmas are well-known. We present simple proofs for the sake of completeness.

LEMMA III.2.1. *Let $g, h \in G$ be loxodromic elements. Then* $\operatorname{fix}(g)$ *and* $\operatorname{fix}(h)$ *are either disjoint or equal.*

PROOF. Suppose that g and h share at least one fixed point $p \in \partial_\infty X$. By Theorem III.1.4, the subgroup $\langle g, h\rangle$ is either virtually cyclic or contains infinitely many independent loxodromic elements. In the former case, g and h are commensurate, and so $\operatorname{fix}(g) = \operatorname{fix}(h)$. In the latter case, $\langle g, h\rangle$ contains a loxodromic element not fixing p, which is a contradiction. □

LEMMA III.2.2. *Let $p, q \in \partial_\infty X$ be distinct. Then the stabiliser of $\{p, q\}$ is either virtually cyclic and contains a loxodromic element, or is finite.*

PROOF. Denote by $N \leqslant G$ the stabiliser of $\{p, q\}$. Passing to an index two subgroup if necessary, we may assume that N fixes both p and q. By Proposition I.5.7, let ℓ be a bi-infinite quasi-geodesic with endpoints p and q. As N fixes p and q, $s \cdot \ell$ is quasi-geodesic lying a uniform finite Hausdorff distance from ℓ for any $s \in N$.

By Theorem III.1.4, we have that N either acts with bounded orbits, is virtually cyclic containing a loxodromic, or contains independent loxodromic elements. The final case may not occur, as N would contain a loxodromic element not fixing either p or q. Suppose, then, that N is not virtually cyclic containing a loxodromic. Then the orbits of N are bounded in X. It follows that the diameter of the orbit Nx is uniformly bounded for any $x \in \ell$. Picking $x, y \in \ell$ sufficiently far apart, the definition of acylindricity implies that N is finite. □

Since the action of G is non-elementary, $\partial_\infty X$ contains more than two points. The following is then a straightforward consequence of the above lemma.

Corollary III.2.3. *The pointwise stabiliser* $N = \{g \in G : \mathrm{fix}(g) = \partial_\infty X\}$ *of the boundary is a finite normal subgroup of* G.

The following statement will also be helpful.

Lemma III.2.4. *Let* $g \in G$ *be loxodromic and let* $h \in G$ *be such that* $h \cdot g^\infty \neq g^{-\infty}$. *There exists* $N \in \mathbf{N}$ *such that the element* hg^n *is also loxodromic for all* $n \geq N$. *In particular,* hg^n *has infinite order.*

Proof. Recall that $\widehat{X} = X \cup \partial_\infty X$ denotes the completion of X with its boundary (see § I.5.6). Let U and V be open neighbourhoods in $\widehat{X}$ of g^∞ and $g^{-\infty}$ respectively such that $\overline{U} \cap \overline{V} = \emptyset$, $\overline{U} \cup \overline{V} \neq \widehat{X}$, and $h\overline{U} \cap \overline{V} = \emptyset$.

We have that hU and V are open neighbourhoods of $h \cdot g^\infty$ and $g^{-\infty}$ with $\overline{hU} \cap \overline{V} = \emptyset$ and $\overline{hU} \cup \overline{V} \neq \widehat{X}$. Moreover, by [**73**, Prop. 4.4] there is some $N \in \mathbf{N}$ such that $g^n(\widehat{X} \setminus V)$ is contained in U for all $n \geq N$ (see the definition of *contractive G-completion* in [**73**, §2]). Therefore, $hg^n(X \setminus V)$ is contained in hU. Now applying [**73**, Lem. 2.6, Cor. 4.5], we have that hg^n is loxodromic for all $n \geq N$. ☐

Lemma III.2.5. *For any* $K \geq 0$ *there is a constant* $M = M(\lambda, c, K) \geq 0$ *such that the following is true.*

Let $x \in X$ *and let* $g \in G$ *be such that* $\mathrm{d}(x, gx) \leq K$. *Suppose that* $\mathrm{fix}(g)$ *is non-empty and let* $p \in \mathrm{fix}(g)$. *If* ℓ *is a* (λ, c)*-quasi-geodesic ray in* X *emanating from* x *and converging to* p *then for all* $y \in \ell$ *we have that* $\mathrm{d}(y, gy) \leq M$.

Proof. Using Lemma I.5.10, fix $\delta' \geq 0$ such that any ideal (λ, c)-quasi-geodesic triangle is δ'-slim. Let $\mu = \mu(\lambda, c) \geq 0$ be the Morse constant for (λ, c)-quasi-geodesics as in Lemma I.5.9. Consider the (λ, c)-quasi-geodesic triangle $T = \ell \cup g\ell \cup [x, gx]$, where $[x, gx]$ is some choice of geodesic path connecting x to gx. We have that T is δ'-slim. Fix $y \in \ell$ and write $L = \mathrm{d}(y, x)$. If $L \leq \delta' + K$, then choosing $M = 2\delta' + 3K$ would give that

$$\mathrm{d}(y, gy) \leq \mathrm{d}(y, x) + \mathrm{d}(x, gx) + \mathrm{d}(gx, gy) \leq 2L + K = M$$

and we are done. Therefore we suppose otherwise, that $L > \delta' + K$. By the δ'-slimness of T there exists some $z \in \ell \cup [x, gx]$ such that $\mathrm{d}(gy, z) \leq \delta'$. By our assumption on L and the fact that $\mathrm{d}(x, gx) \leq K$, it must be $z \in \ell$. We will show that $\mathrm{d}(y, z)$ is uniformly bounded above, which implies the lemma.

Parameterise ℓ as a (λ, c)-quasi-isometric embedding $\ell : [0, \infty) \to X$. Choose $a, b \geq 0$ such that $\ell(a) = y$, $\ell(b) = z$. We assume without loss of generality that $a \leq b$. The other case is similar. Let

$q = [x, z]$ be a geodesic connecting x to z. Then y is contained in the μ-neighbourhood of q. Choose $y' \in q$ such that $\mathrm{d}(y, y') \le \mu$. We have that

$$\begin{aligned}\mathrm{d}(z, x) &= \mathrm{d}(x, y') + \mathrm{d}(y', z)\\ &\ge \mathrm{d}(x, y) - \mu + \mathrm{d}(y, z) - \mu\\ &= \mathrm{d}(y, z) + L - 2\mu.\end{aligned}$$

We also have that

$$\begin{aligned}\mathrm{d}(z, x) &\le \mathrm{d}(z, gy) + \mathrm{d}(gy, gx) + \mathrm{d}(gx, x)\\ &\le \delta' + L + K.\end{aligned}$$

Combining the above, we deduce that

$$\mathrm{d}(y, z) \le K + \delta' + 2\mu,$$

so that lemma follows from setting $M = \max\{2\delta' + 3K, K + 2\delta' + 2\mu\}$. □

For the remainder of the section we will assume that the action of G is such that $\Lambda G = \partial_\infty X$.

Lemma III.2.6. *Given $g \in G$, either $\mathrm{fix}(g) = \partial_\infty X$ or $\mathrm{fix}(g)$ is nowhere dense.*

Proof. Suppose that $\mathrm{fix}(g) \neq \partial_\infty X$ and that it is not nowhere dense. Since $\mathrm{fix}(g)$ is closed, there is an open subset $U \subset \mathrm{fix}(g)$. Let $V \subset \partial_\infty X$ be an open set disjoint from $\mathrm{fix}(g)$. By Theorem III.1.7, there is loxodromic a z with $z^\infty \in U$ and $z^{-\infty} \in V$. Now z^g is a loxodromic with $(z^g)^\infty = g \cdot z^\infty = z^\infty$ but $(z^g)^{-\infty} = g \cdot z^{-\infty} \neq z^{-\infty}$, contradicting Lemma III.2.1. □

Lemma III.2.7. *Fix $x \in X$ and let $B \subset G$ be a subset with $\mathrm{diam}(Bx) < \infty$. Then there is $N \in \mathbf{N}$ such that for every $p \in \partial_\infty X$, there exists a neighbourhood U of p such that*

$$|\{b \in B : \mathrm{fix}(b) \cap U \neq \emptyset\}| \le N.$$

In particular, if $\mathrm{fix}(b)$ is discrete for every $b \in B$, then so is $\bigcup_{b \in B} \mathrm{fix}(b)$.

Proof. Write $K = \mathrm{diam}(Bx) < \infty$. Let $p \in \partial_\infty X$. By Proposition I.5.7, there exists a $(1, 20\delta)$-quasi-geodesic ray ℓ emanating from x and converging to p. By Lemma I.5.9, there is $\delta' \ge 0$ for which $(1, 20\delta)$-quasi-geodesic triangles in $X \cup \partial_\infty X$ are δ'-slim. Let $M = M(1, 20\delta, K) \ge 0$ be as in Lemma III.2.5. Let $R, N > 0$ be the constants provided by acylindricity, applied with $\varepsilon = M + 2\delta'$. Let $U \subset \partial_\infty X$ be a neighbourhood of p such that for any point $q \in U$, a $(1, 20\delta)$-quasi-geodesic ray emanating from x converging on q has an initial segment of length $2R$ that lies in a δ'-neighbourhood of an initial segment of ℓ of the same length.

Suppose that there are n distinct elements $b_1, \dots, b_n \in B$ with $q_1 \dots, q_n \in U$, where $q_i \in \mathrm{fix}(b_i)$. Let ℓ_i be $(1, 20\delta)$-quasi-geodesic rays from x to q_i for each i, and again observe that by Lemma III.2.5,

$d(z, bz) \leq M$ for any $z \in \ell_i$ and $b \in B$. Take $y \in \ell$ with $R < d(x, y) < 2R$, so that for each i there is a point $z_i \in \ell_i$ with $d(y, z_i) < \delta'$. Then for each i, we have

$$d(y, b_i y) \leq d(y, z_i) + d(z_i, b_i z_i) + d(b_i z_i, b_i y) < M + 2\delta'.$$

Of course, $d(x, b_i x) \leq K < M$, so acylindricity implies that $n \leq N$. The lemma follows. □

LEMMA III.2.8. *Fix $x \in X$ and let $B \subset G$ be a subset with* $\operatorname{diam}(Bx) < \infty$. *If* $\operatorname{fix}(b) \neq \partial_\infty X$ *for all $b \in B$, then* $\bigcup_{b \in B} \operatorname{fix}(b)$ *is nowhere dense.*

PROOF. Write $C = \bigcup_{b \in B} \operatorname{fix}(b)$. If the closure $\overline{C}$ contains an open set V, there is $p \in V \cap C$. Let U be the neighbourhood of p provided by Lemma III.2.7, so that U only meets finitely many of the sets $\operatorname{fix}(b)$. By Lemma III.2.6, each $\operatorname{fix}(b)$ is nowhere dense. Therefore $U \cap C$ is the union of finitely many nowhere dense subsets, and hence is nowhere dense. But then $U \cap V$ is an open set contained in the closure of $U \cap C$, a contradiction. □

The following corollary is immediate from Lemma III.2.7. We do not use it here, but include it as it may be useful for applications.

COROLLARY III.2.9. *There is $N \geq 0$ such that if $g \in G$ is elliptic with* $\operatorname{fix}(g) \neq \emptyset$, *then g has order at most N.*

III.3. Swingers and tiles

Let G be a countable group, and let S denote a (possibly infinite) generating set. We will denote by $|\cdot| = |\cdot|_S$ the word length of an element of G with respect to S, and $d = d_S$ for the associated edge-path metric on $\operatorname{Cay}(G, S)$. Let us introduce the following technical terminology.

DEFINITION III.3.1 (Swingers). Given $r > 0$, we say that an element $z \in G$ is an *r-swinger with respect to S* if z is loxodromic in $\operatorname{Cay}(G, S)$, and for every $b \in G$ with $1 \leq |b| \leq r$ we have

$$\left|z^{mi} b z^{mj}\right| > |z^m|$$

for each $i, j \in \{1, -1\}$, $m \geq 1$.

The nomenclature above is motivated by the dynamics of such elements. One can imagine that conjugating a swinger element z by such short elements b has the effect of 'swinging' z around, causing it to point in a totally different direction from both z and z^{-1}. The following lemma illustrates this intuition.

LEMMA III.3.2. *Let $z \in G$ be an r-swinger with respect to S. Then for all $b \in G$ with $1 \leq |b| \leq r$, the subgroup $\langle z, z^b \rangle$ is not virtually cyclic.*

Proof. Suppose otherwise, that $H = \langle z, z^b \rangle$ is virtually cyclic. Then both $\langle z \rangle$ and $\langle z^b \rangle$ have finite index in H. Thus the intersection $\langle z \rangle \cap \langle z^b \rangle$ has finite index in H, and in particular is non-trivial. Thus there are $n_1, n_2 \neq 0$ such that $(z^{n_1})^b = z^{n_2}$. We treat the case that $n_1 \geq n_2 \geq 1$, for the other cases are similar. Since z is an r-swinger, we have that for any $m \geq 1$,

$$\left| z^{m(n_1-n_2)} \right| = \left| z^{-n_2 m} z^{n_1 m} \right| = \left| (z^{n_1 m})^b z^{n_1 m} \right| \geq |z^{n_1 m} b z^{n_1 m}| - r > |z^{n_1 m}| - r.$$

As z is loxodromic, the map $n \mapsto z^n$ is a (λ, c)-quasi-isometry for some $\lambda \geq 1$ and $c \geq 0$. The above inequality thus yields

$$\lambda m(n_1 - n_2) + c > \tfrac{1}{\lambda} m n_1 - c - r$$

for all $m \geq 1$. However, since $n_1 - n_2 < n_1$, this inequality can only hold for finitely many m: a contradiction. ☐

Definition III.3.3. Let G be a group and let $S \subset G$ be a generating set. We say that $\mathrm{Cay}(G, S)$ *admits swingers* if G contains an r-swinger with respect to S for each $r > 0$.

We will show that if a group G acts acylindrically on a hyperbolic Cayley graph that admits swingers, then G is monotileable. The idea behind this argument is essentially due to Akhmedov and appears in [**2**]. Recall that a subset Y of a metric space is *r-separated* if $\mathrm{d}(x, y) \geq r$ for any distinct $x, y \in Y$. We begin with the following lemma.

Lemma III.3.4. *Let G be a group with generating set S and suppose that* $\mathrm{Cay}(G, S)$ *is δ-hyperbolic. For each $r > 0$, there is $R = R(r, \delta) > 0$ such that if $z \in G$ is an r-swinger with $|z| \geq R$, the subset*

$$D_{z,r} = \{x \in G : |xz| < |x| + r \quad \text{or} \quad |xz^{-1}| < |x| + r\}$$

is r-separated.

Proof. Let $x, y \in D_{z,r}$ be distinct elements so that there are $i, j \in \{1, -1\}$ such that $\left|xz^i\right| < |x| + r$ and $\left|yz^j\right| < |y| + r$. Pick geodesics $\alpha = [1, xz^i], \beta = [1, yz^j]$, and $\gamma = [xz^i, yz^j]$. Suppose for a contradiction that $\mathrm{d}(x, y) \leq r$, so that we have $1 \leq \left|y^{-1}x\right| \leq r$. Then since z is an r-swinger,

$$|\gamma| = \mathrm{d}(xz^i, yz^j) = \left|z^{-j} y^{-1} x z^i\right| > |z|$$

Therefore there is a point p of γ with

$$\mathrm{d}(xz^i, p) \geq \frac{1}{4}|z| \quad \text{and} \quad \mathrm{d}(yz^j, p) \geq \frac{1}{4}|z| \tag{3}$$

By the δ-slimness of the triangle $\gamma \cup [x, xz^i] \cup [x, yz^j]$, there is a point q of $[x, xz^i] \cup [x, yz^j]$ such that $\mathrm{d}(p, q) \leq \delta$. If q lies on $[x, yz^j]$ then

$$\begin{aligned}
\mathrm{d}(q, x) &= \mathrm{d}(x, yz^j) - \mathrm{d}(q, yz^j) \\
&\leq |z| + r - (\mathrm{d}(p, yz^j) - \mathrm{d}(p, q)) \\
&\leq \frac{3}{4}|z| + r + \delta
\end{aligned}$$

where the last inequality is an application of (3). It follows that

$$\mathrm{d}(p,x) \leq \mathrm{d}(p,q) + \mathrm{d}(q,x) \leq \frac{3}{4}|z| + r + 2\delta. \tag{4}$$

A similar argument gives the same upper bound when q lies on $[x, xz^i]$. Now by slimness of the triangle $\alpha \cup \beta \cup \gamma$, the point p lies in a δ-neighbourhood of $\alpha \cup \beta$. We may assume p lies in the δ-neighbourhood of a point t of α, with the other case similar. Since α is a geodesic representing xz, we have $|xz^i| = \mathrm{d}(1,t) + \mathrm{d}(t, xz^i)$. Moreover, by (4) and the definition of t, we have that

$$\mathrm{d}(x,t) \leq \mathrm{d}(x,p) + \mathrm{d}(p,t) \leq \frac{3}{4}|z| + r + 3\delta.$$

Then combining the above with the triangle inequality,

$$|xz^i| \geq \mathrm{d}(1,x) + \mathrm{d}(x, xz^i) - 2\,\mathrm{d}(x,t) \geq |x| + \tfrac{3}{2}|z| - 2r - 6\delta. \tag{5}$$

Since $x \in D_{z,r}$, the inequality $|xz^i| < |x| + r$ holds. Therefore (5) implies that

$$|z| < 2r + 4\delta.$$

Now choosing $R = 2r + 4\delta$ gives us the desired contradiction. □

In order to proceed we will also need the following basic statement about tiling sets with two elements.

LEMMA III.3.5. *Let G be a group, and A be a set upon which G acts on the right freely. Let $B \subset A$, and suppose there is $g \in G$ such that $Bg \subset B$. If the order of g is even or infinite, then B is a disjoint union of subsets of the form $\{b, bg\}$, where $b \in B$.*

PROOF. It is clear that there is a subsets $C_1, C_2 \subset B$ such that

$$B = \bigsqcup_{c \in C_1} \{cg^n : n \geq 0\} \sqcup \bigsqcup_{c \in C_2} \{cg^n : n \in \mathbf{Z}\}.$$

This follows, for example, by considering the partition of A induced by the $\langle g \rangle$-orbits and restricting this partition to B. Now, fix some arbitrary $c \in B$. Since the order of g is even or infinite and the action is free, we see that

$$\{cg^n : n \geq 0\} = \bigsqcup_{n \geq 0} \{c, cg\} g^{2n}, \quad \text{and} \quad \{cg^n : n \in \mathbf{Z}\} = \bigsqcup_{n \in \mathbf{Z}} \{c, cg\} g^{2n}.$$

The claim now follows. □

We are ready to prove the following important statement.

PROPOSITION III.3.6. *Let G be a countable group with (possibly infinite) generating set S. Suppose $\mathrm{Cay}(G,S)$ is δ-hyperbolic and admits swingers. If the action of G on $\mathrm{Cay}(G,S)$ is acylindrical, then for any finite subset $F \subset G$, there exists $z \in G$ such that $F \cup \{z\}$ is a tile for G. In particular, G is monotileable.*

PROOF. By translating F in advance, we may assume that $1 \in F$. If $|F| = 1$ then F is already a tile of G. Thus, we assume that $|F| > 1$, so that there is $v \in F$ with $v \neq 1$. Write $M = \max\{|g| : g \in F \cup F^{-1}\}$ and let $r = 4M + 1$. Let $R = R(r, \delta) > 0$ be the constant provided by Lemma III.3.4.

Now, by assumption there is an r-swinger $y \in G$. Note that no nontrivial $b \in G$ with $|b| \leq r$ sends one endpoint of y in $\partial_\infty X$ to another. Indeed, otherwise y and y^b would either share a single fixed point in $\partial_\infty X$, contradicting Lemma III.2.1, or generate a virtually cyclic subgroup, contradicting Lemma III.3.2. Let $N \in \mathbf{N}$ be the number provided by Lemma III.2.4 such that $v^{-1}y^n$ has infinite order for all $n \geq N$. We take $z = y^n$ with $n \geq N$ large enough to ensure $|z| \geq R$.

Let

$$C = C_{z,r} = \{x \in G : |xz^{-1}| < |x| + r\}.$$

By Lemma III.3.4, the set $D = D_{z,r}$ is r-separated. As $C \subset D$, the set C is also r-separated. We observe the following.

CLAIM III.3.7. *For all $s \in C$, we have $sv^{-1}z \in C$.*

PROOF. For the first statement, note that $v \neq 1$ and $|v| < r$. Since $s \in C \subset D$ and D is r-separated, $sv^{-1} \notin D$. By definition of D, we have $|sv^{-1}z| \geq |sv^{-1}| + r$. Then

$$|(sv^{-1}z)z^{-1}| = |sv^{-1}| \leq |sv^{-1}z| - r < |sv^{-1}z| + r.$$

Thus it follows from the definition of C that $sv^{-1}z \in C$.. □

Claim III.3.7 tells us that $Cv^{-1}z \subset C$, and we know $v^{-1}z$ has infinite order by our choice of z. Hence by Lemma III.3.5, there exists a subset $C' \subset C$ such that C decomposes as the disjoint union

$$C = \bigsqcup_{s \in C'} \{s, sv^{-1}z\}. \tag{6}$$

We now begin tiling our group. Our candidate tile will be $F \cup \{z\}$. For ease of notation we write $T = F \cup \{z\}$. Let

$$A = \bigcup_{s \in C'} sv^{-1}T.$$

Observe that for any $s \in C$, $\{s, sv^{-1}z\} \subset sv^{-1}T$, since $\{v, z\} \subset T$. Hence (6) shows $C \subset A$. We will show that the above union is in fact a disjoint union.

CLAIM III.3.8. *For any distinct $s, s' \in C'$ we have $sv^{-1}T \cap s'v^{-1}T = \emptyset$.*

PROOF. Suppose that $sv^{-1}T \cap s'v^{-1}T$ is non-empty. As $s \neq s'$, one of the following must hold:

(1) $sv^{-1}F \cap s'v^{-1}F \neq \emptyset$, or
(2) $sv^{-1}z \in s'v^{-1}F$, or
(3) $s'v^{-1}z \in sv^{-1}F$.

If the first case holds, then $\mathrm{d}(s, s') \leq 4M < r$, contradicting the fact that C is r-separated. Suppose the second case holds, so that $sv^{-1}z \in s'v^{-1}F$. Similarly to the first case, this implies that $\mathrm{d}(s', sv^{-1}z) < r$. Observing that $s' \in C, sv^{-1}z \in C$ by Claim III.3.7, and that C is r-separated, it must be the case that $sv^{-1}z = s'$. But this contradicts the fact that $\{s, sv^{-1}z\}$ and $\{s', s'v^{-1}z\}$ are disjoint as in (6). The final case may be dealt with identically. $\square$

In other words, we have that the translates of T by $\{sv^{-1} : s \in C'\}$ are pairwise disjoint and cover A. We now proceed to tile $G \setminus A$ by translates of T in a naïve way, picking remaining elements of minimal word length and covering them with the images of z in T. We show that the tiles we obtain in this fashion are disjoint from both A and from one another.

Claim III.3.9. *If $b \notin A$, then $bz^{-1}T \cap A = \emptyset$. Furthermore, if $c \notin A \sqcup bz^{-1}T$ and $|b| \leq |c|$, then $bz^{-1}T \cap cz^{-1}T = \emptyset$.*

Proof. To begin, note that for any $x \notin A$, we have $x \notin C$. It follows from the definition of C that

$$\left|xz^{-1}\right| \geq |x| + r. \tag{7}$$

An application of (7) gives that

$$\left|xz^{-1}z\right| = |x| \leq \left|xz^{-1}\right| - r < \left|xz^{-1}\right| + r.$$

Thus, $xz^{-1} \in D$ for any $x \notin A$. In particular, $bz^{-1} \in D$ and $cz^{-1} \in D$ for b and c as in the statement of the claim.

Let $b \notin A$ be an arbitrary element and suppose that $bz^{-1}T \cap A$ is non-empty. By the definition of A we have that $bz^{-1}T \cap sv^{-1}T \neq \emptyset$ for some $s \in C'$. By construction, we have that $b \notin A \supset sv^{-1}T$, and so either

(1) $bz^{-1}F \cap sv^{-1}F \neq \emptyset$, or
(2) $sv^{-1}z \in bz^{-1}F$.

In both cases we will deduce a contradiction. Note that by Claim III.3.7 we have that $sv^{-1}z \in C \subset D$. Recall that $bz^{-1} \in D$ since $b \notin A$. In the first case above, note that $\mathrm{d}(sv^{-1}, bz^{-1}) \leq 2M$, and so we may deduce that $\mathrm{d}(s, bz^{-1}) \leq 3M < r$. But this contradicts the fact that D is r-separated. In the second case, we see that $\mathrm{d}(sv^{-1}z, bz^{-1}) < r$. Similarly, this contradicts the fact that D is r-separated. Therefore $bz^{-1}T$ and A are disjoint.

Now suppose there is an element $c \notin A \sqcup bz^{-1}T$ such that $bz^{-1}T \cap cz^{-1}T$ is not empty. Since $b \neq c$ one of the following three cases must hold:

(1) $b \in cz^{-1}F$, or
(2) $c \in bz^{-1}F$, or
(3) $bz^{-1}F \cap cz^{-1}F \neq \emptyset$.

For the first case we have that $b = cz^{-1}t$ for some $t \in F$. Then

$$|b| = |cz^{-1}t| \geq |cz^{-1}| - |t| > |c| + r - r = |c| \geq |b|,$$

where the middle inequality follows from (7) and the fact that $|t| < r$ since $t \in F$. The second case immediately contradicts the construction of c, since $c \notin bz^{-1}T \supset bz^{-1}F$. In the third case, recall that $bz^{-1} \in D$ and $cz^{-1} \in D$. However, by assumption we have that $\mathrm{d}(bz^{-1}, cz^{-1}) \leq 2M < r$, which contradicts the fact that D is r-separated. □

Write $A_0 = G$ and $A_1 = A$. We will proceed with tiling the complement $G \setminus A$ by induction. Let $i \geq 1$ and suppose we have constructed a set A_i. Let $n \in \mathbf{N}$ be the minimal natural number for which there is $b \notin A_i$ with $|b| = n$. We claim there is a set A_{i+1}, formed of a disjoint union of A_i and pairwise disjoint translates of T, satisfying

$$\min\{|g| : g \notin A_{i+1}\} > n.$$

Indeed, since G is countable we may enumerate the elements of $G \setminus A_i$ with word length n. Let $\{g_1, g_2, \dots\}$ be such an enumeration and write $A_i^{(1)} = A_i$ and $n_1 = 1$. Now for any $j > 1$, let n_j be the smallest natural number such that $g_{n_j} \notin \bigcup_{k<j} g_{n_k} z^{-1}T$. Claim III.3.9 tells us that $A_i^{(j-1)}$ and $g_{n_j} z^{-1}T$ are disjoint, and also by induction that $g_{n_l} z^{-1}T$, $g_{n_k} z^{-1}T$ are pairwise disjoint for all $1 \leq k < l \leq j$. Let

$$A_i^{(j)} = A_i^{(j-1)} \sqcup g_{n_j} z^{-1}T.$$

Setting $A_{i+1} = \bigcup_{j\in\mathbf{N}} A_i^{(j)}$, we readily see that $g_k \in A_{i+1}$ for all $k \in \mathbf{N}$. Hence A_{i+1} satisfies the desired criterion. Now by construction, the sets $\{A_i : i \in \mathbf{N}\}$ form a nested family of subsets that exhaust G. Since each A_i is comprised of disjoint translates of $T = F \cup \{z\}$, the set $F \cup \{z\}$ is a tile for G. □

Remark III.3.10. We note that in the proof of Proposition III.3.6, we only use the boundedness of the set F with respect to d_S, rather than the finiteness. Combined with the results of the next section, one sees that if G is acylindrically hyperbolic with no nontrivial finite normal subgroup, then for any bounded subset $F \subset G$, there is $z \in G$ such that $F \cup \{z\}$ is a tile for G. It may be of interest, for example then, that any subset of a hyperbolically embedded subgroup of a torsion-free group G can be extended to a tile of G by adding a single element.

III.4. Finding swingers in acylindrically hyperbolic groups

In this section we take G to be an acylindrically hyperbolic group. As such, there is a generating set S of G such that the Cayley graph $X = \mathrm{Cay}(G, S)$ is δ-hyperbolic and the action of G on X by left translation is acylindrical, non-elementary, and cobounded. Throughout this section, $\mathrm{d} = \mathrm{d}_X$ will denote the metric on G induced from X. Given $g \in G$, we will write $|g| := \mathrm{d}(g, 1)$.

Lemma III.4.1. *Suppose that $B \subset G$ is a bounded subset of G such that no element of B fixes $\partial_\infty X$ pointwise. Then there is a loxodromic element $z \in G$ and an open set U containing $\mathrm{fix}(z)$ such that $U \cap \mathrm{fix}(z^b) = \emptyset$ for all $b \in B$.*

Proof. Write $B' = B \cup B^2$, where $B^2 = \{b^2 : b \in B\}$. and note that $\mathrm{diam}_X(B') < \infty$. By Lemma III.2.8, the set $C = \bigcup_{b \in B'} \mathrm{fix}(b)$ is not dense in $\partial_\infty X$. Hence there is an open set $V \subset \partial_\infty X$ disjoint from C. Applying Theorem III.1.7, there is a loxodromic element $z \in G$ with $\mathrm{fix}(z) \subset V$. We will first show that no element of B setwise fixes $\mathrm{fix}(z)$.

Let $b \in B$ and suppose that $z^\infty = (z^b)^\infty$. Now $(z^b)^\infty$ is exactly the translate $b \cdot z^\infty$, and so b fixes z^∞. Thus, $z^\infty \in C$, which contradicts the fact that $z^\infty \in V$. Similarly, if $z^{-\infty} = (z^b)^{-\infty}$ then $z^{-\infty} \in C$, which is again a contradiction.

It remains to show that $(z^b)^\infty \neq z^{-\infty}$ or $(z^b)^{-\infty} \neq z^\infty$. If either of these conditions fails, then z^b and z share the same fixed points by Lemma III.2.1. Thus b permutes the set $\{z^\infty, z^{-\infty}\}$, and so b^2 fixes it. But then $\{z^\infty, z^{-\infty}\} \subset C$, contradicting the construction of z.

Since B is bounded, we have that $\mathrm{diam}(\{z^b : b \in B\}) < \infty$. Moreover, $\mathrm{fix}(z^b)$ is discrete for every $b \in B$, so by Lemma III.2.7 we have that

$$A = \bigcup_{b \in B} \mathrm{fix}(z^b)$$

is discrete. In particular, since $\mathrm{fix}(z)$ is discrete and disjoint from A, there is some neighbourhood U of $\mathrm{fix}(z)$ which is disjoint from A. $\square$

We show that the dynamical condition obtained above can be used to find a swinger. The following lemma is standard.

Lemma III.4.2. *For every $\lambda \geq 1$, $c \geq 0$, and $s \geq 0$, there exists $t = t(\lambda, c, s) \geq 0$ such that the following holds. Let $x, y \in G$ be loxodromic elements such that $\langle x^\infty \cdot y^\infty \rangle_1 < s$, and the inclusions of $\langle x \rangle$ and $\langle y \rangle$ into G are (λ, c)-quasi-isometric embeddings. Then*

$$\mathrm{d}(x^m, y^n) > |x^m| + |y^n| - t.$$

for any $n, m \geq 1$.

Proof. Since $\langle x^\infty \cdot y^\infty \rangle_1 < s$, we have by definition that

$$\liminf_{i,j \to \infty} \langle x^i \cdot y^j \rangle_1 < s.$$

Let $m, n \geq 1$. Choose $M, N \geq 1$ such that $m \leq M$, $n \leq N$, and $\langle x^M \cdot y^N \rangle_1 < s$. Let $p = [1, x^M]$, $q = [1, y^N]$ be geodesics. By the Morse lemma (I.5.9), we have that there exists $\mu = \mu(\lambda, c) > 0$ such that x^m lies in the μ-neighbourhood of p and y^n lies in the μ-neighbourhood of q. Fix a and b be points on p and q such that $\mathrm{d}(x^m, a) \leq \mu$ and $\mathrm{d}(y^n, b) \leq \mu$. Since p and q are geodesics and

$\langle x^M \cdot y^N \rangle_1 < s$, we have that $\langle a \cdot b \rangle_1 < s$ also. Thus, by the definition of the Gromov product, we see that

$$\mathrm{d}(a, b) > |a| + |b| - 2s.$$

In particular, we deduce that

$$\mathrm{d}(x^m, y^n) > |x^m| + |y^n| - 2s - 4\mu.$$

Setting $t = 2s + 4\mu$, we are done. $\square$

LEMMA III.4.3. *Let $B \subset G$ is a bounded subset of G such that no element of B fixes $\partial_\infty X$ pointwise. Let $z \in G$ and $U \subset \partial_\infty X$ be as in the conclusion of Lemma III.4.1. Then there is $M \in \mathbf{N}$ such that for any $m \geq M$, $i, j \in \{\pm 1\}$, $b \in B$ we have that*

$$\left|z^{im} b z^{jm}\right| > |z^m|.$$

PROOF. We only consider the case where $i = j = 1$, as the others follow from identical reasoning.

Recall from § I.5.6 the definition of the basic open sets $U(p, r) \subset \partial_\infty X$, where $p \in \partial_\infty X$ and $r \geq 0$. Let $s > 0$ be the supremum of all s such that the basic open sets $U(z^\infty, s)$ are contained in U. Then for all $b \in B$, we have that $\langle z^\infty \cdot (z^b)^{-\infty} \rangle_1 < s$. As B is bounded, there exist uniform constants $\lambda \geq 1$, $c \geq 0$ such that for all $b \in B \cup \{1\}$, the inclusions of $\langle z^b \rangle$ into G are (λ, c)-quasi-isometric embeddings. Let $t = t(\lambda, c, s)$ be as in Lemma III.4.2, then for every $m \geq 0$, $b \in B$ we have that

$$\mathrm{d}(bz^{-m}b^{-1}, z^m) > \left|bz^{-m}b^{-1}\right| + |z^m| - t. \tag{8}$$

Let $R = \max\{|b| : b \in B \cup B^{-1}\} < \infty$. Now, fix some arbitrary $b \in B$. We have that $\mathrm{d}(z^b, bz) = \mathrm{d}(1, b) \leq R$. Hence, we see that

$$\mathrm{d}(bz^m, z^{-m}) \geq \mathrm{d}(bz^{-m}b^{-1}, z^m) - R.$$

Combining this with (8), we obtain

$$|z^m b z^m| = \mathrm{d}(bz^m, z^{-m}) > |bz^m| + |z^m| - t - 2R.$$

We also have by the triangle inequality that $|bz^m| \geq |z^m| - R$. Hence, we conclude that $|z^m b z^m| = \mathrm{d}(bz^m, z^{-m}) \geq 2|z^m| - t - 3R$. Rearranging, we write

$$|z^m b z^m| \geq |z^m| + (|z^m| - t - 3R). \tag{9}$$

As z is a loxodromic element, $|z^m| \to \infty$ as $m \to \infty$. That is, there is $M \in \mathbf{N}$ such that $|z^m| > t+3R$ for all $m \geq M$. Hence by (9) we have $|z^m b z^m| > |z^m|$. Since $b \in B$ was arbitrary and t, R and M do not depend on the choice of b, the result follows. $\square$

PROPOSITION III.4.4. *Suppose that G contains no nontrivial finite normal subgroups. Then $X = \mathrm{Cay}(G, S)$ admits swingers.*

Proof. Let $r > 0$ and $B = \{b \in G : 1 \leq |b| \leq r\}$. We show that G contains an r-swinger $z \in G$. Together with the assumption that G contains no finite normal subgroup, Corollary III.2.3 tells us that no nontrivial elements of G pointwise fix $\partial_\infty X$. Thus we may apply Lemma III.4.1 to find a loxodromic element $y \in G$ and an open subset U containing $\mathrm{fix}(y)$ such that $U \cap \mathrm{fix}(y^b) = \emptyset$ for all $b \in B$. Now Lemma III.4.3 tells us that there is $M \in \mathbf{N}$ such that $\left|y^{im} b y^{jm}\right| > |y^m|$ for all $m \geq M$, $i, j \in \{1, -1\}$, and $b \in B$. It follows that $z = y^M$ is an r-swinger with respect to S, completing the proof. □

III.5. Main results

We may now prove the main results of this chapter.

Theorem III.5.1. *Acylindrically hyperbolic groups are monotileable.*

Proof. We begin by showing that a countable acylindrically hyperbolic group G is monotileable. By [**33**, Thm. 6.14], G contains a maximal finite normal subgroup $K \lhd G$. The quotient G/K is again acylindrically hyperbolic [**100**, Lem. 3.9], and contains no nontrivial finite normal subgroups. Now combining Propositions III.3.6 and III.4.4 shows that G/K is monotileable. As monotileability is stable under extensions by finite groups, G is monotileable.

Now let G be an arbitrary acylindrically hyperbolic group. Let $F \subset G$ be a finite subset. Since G is acylindrically hyperbolic, it contains at least two independent loxodromic elements $g, h \in G$. Then by Theorem III.1.4, $H = \langle F, g, h \rangle$ is a countable acylindrically hyperbolic subgroup of G containing F. Therefore F extends to a finite tile of H. As the cosets of H are disjoint and cover G, it follows that F is contained in a finite tile of G. □

Corollary III.5.2. *One-relator groups are monotileable.*

Proof. Let G be a group with one-relator presentation $\langle S \ ; \ r \rangle$. If S contains three or more elements then G is acylindrically hyperbolic [**100**, Cor. 2.6], whence it is monotileable by Theorem III.5.1. Otherwise, combining [**100**, Prop. 4.21] and [**22**, Thm. 3.2], G is either acylindrically hyperbolic, a generalised Baumslag–Solitar group, or a mapping torus of an injective endomorphism of a finitely generated free group.

In the first case, we are done by Theorem III.5.1 again. In the second case, we observe that generalised Baumslag–Solitar groups are free-by-metabelian by a result of P. Kropholler [**84**, Thm. C, Cor. 2]. Both free groups and metabelian groups are monotileable, so free-by-metabelian groups are too. In the final case, such groups are known to be residually finite by a result of Borisov and Sapir [**13**, Thm. 1.2]. As residually finite groups are monotileable, the last case is covered also. □

Corollary III.5.3. *Two-dimensional Artin groups are monotileable.*

Proof. Let A_Γ be a two-dimensional Artin group. Recall that an Artin group is called *reducible* if Γ is a join of subgraphs Γ_1 and Γ_2, and any edge $e \in E\Gamma \setminus (E\Gamma_1 \cup E\Gamma_2)$ has label 2, and *irreducible* otherwise. If A_Γ is reducible, then $A_\Gamma \cong A_{\Gamma_1} \times A_{\Gamma_2}$. Being monotileable is closed under products, so we may restrict our attention to irreducible Artin groups.

If Γ has a single vertex, then A_Γ is a cyclic group. If Γ has two vertices, then it is called a *dihedral Artin group*, and it is known that A_Γ is an extension of a cyclic group by a free product of cyclic groups [**30**, § 2]. Since being monotileable is closed under extensions and free products, A_Γ is monotileable. Finally, if Γ has more than three vertices, then A_Γ is acylindrically hyperbolic [**133**, Thm. A], whence Theorem III.5.1 applies. $\square$

CHAPTER IV

Groups quasi-isometric to planar graphs

What you call Space is really nothing but a great Plane. I am in Space, and look down upon the insides of the things of which you only see the outsides. You could leave this Plane yourself, if you could but summon the necessary volition.

— E. A. Abbot, *Flatland*

The third and final problem we tackle in this thesis presents a coarse characterisation of those finitely generated groups which are *virtually planar*. That is, those containing finite-index subgroups which admit planar Cayley graphs. We characterise them as exactly those groups which are quasi-isometric to some planar graph. We do not assume that our group is finitely presented, and so the main technical feat of this chapter is to show that such a group is *accessible*, in the sense of Wall. This is achieved through a careful study of quasi-actions on planar graphs. This chapter is based on the article [**89**] of the present author.

IV.1. Preliminaries

IV.1.1. Accessibility. We begin this chapter by introducing *accessibility*, which will play a central role throughout. The following definition is essentially due to Wall [**135**], but translated into the language of Bass–Serre theory.

Definition IV.1.1 (Accessibility for groups). Let G be a finitely generated group. We say that G is *accessible* if G splits as a graph of groups $G(\Gamma) = (G_x, \varphi_e)$ such that every edge group G_e is finite, and every vertex group G_v has at most one end.

This is equivalent to asking that if one splits G over a finite subgroup and passes to the vertex groups iteratively, then this process terminates after a bounded number of steps. As discussed in § I.4, every finitely presented group and every torsion-free finitely generated group is accessible. However, there do exist finitely generated inaccessible groups [**47**].

There is also a graph-theoretical definition of accessibility, due to Thomassen and Woess [**129**]. Given a connected, locally finite graph X, ends $\omega_1, \omega_2 \in \Omega X$ and a finite subset $F \subset E(X)$, we say that F *separates* ω_1 and ω_2 if any bi-infinite path between these ends must cross some $e \in F$. If $K \subset \Gamma$ is a compact subgraph and U is a connected component of $\Gamma \setminus K$, we say that an end

$\omega \in \Omega(\Gamma)$ *lies in* U if every simple ray in ω has infinite intersection with U. We now have the following key definition.

Definition IV.1.2 (Accessibility for graphs). Let X be a connected, locally finite graph. We say that X is *accessible* if there exists $k \geq 1$ such that for any pair of distinct ends $\omega_1, \omega_2 \in \Omega X$, we have that ω_1, ω_2 can be separated by the removal of at most k edges.

It is clear that accessibility in the above sense is a quasi-isometry invariant amongst bounded-valence, connected graphs. Moreover, it was shown by Thomassen–Woess that a finitely generated group is accessible if and only if its Cayley graphs are accessible [**129**, Thm. 1.1].

IV.1.2. Cuts and disjoint paths. We next state a classical result which will play a key role throughout this chapter, due to Menger. We first need to introduce some notation.

Definition IV.1.3. Let Γ be a connected graph, and let $x, y \in V(\Gamma) \cup \Omega(\Gamma)$ be distinct. Define the *vertex separation* of x and y, denoted $\mathrm{vs}(x, y)$, as

$$\mathrm{vs}(x, y) = \inf\{|S| : S \subset V(\Gamma) \setminus \{x, y\},\ x,\ y \text{ lie in distinct components of } \Gamma \setminus S\}.$$

Similarly, define their *edge separation*, denoted $\mathrm{es}(x, y)$, as

$$\mathrm{es}(x, y) = \inf\{|F| : F \subset E(\Gamma),\ x,\ y \text{ lie in distinct components of } \Gamma \setminus F\}.$$

We now define the *minimal vertex end-cut size of* Γ as

$$\mathrm{vs}(\Gamma) = \min\{\mathrm{vs}(\omega_1, \omega_2) : \omega_1, \omega_2 \in \Omega(\Gamma),\ \omega_1 \neq \omega_2\}.$$

We similarly define the *minimal edge end-cut size of* Γ as

$$\mathrm{es}(\Gamma) = \min\{\mathrm{es}(\omega_1, \omega_2) : \omega_1, \omega_2 \in \Omega(\Gamma),\ \omega_1 \neq \omega_2\}.$$

If Γ has at most one end, we adopt the convention that $\mathrm{es}(\Gamma) = \mathrm{vs}(\Gamma) = \infty$.

The following is immediate, but useful to note.

Proposition IV.1.4. *Let* Γ *be a connected, bounded valence graph where every vertex has degree at most* $d \geq 0$. *Then* $\mathrm{vs}(\Gamma) \leq \mathrm{es}(\Gamma) \leq d \cdot \mathrm{vs}(\Gamma)$.

We now state the following important result. This can be seen as a precursor to the famous 'max-flow min-cut' theorem, applying to unweighted graphs.

Theorem IV.1.5 (Menger). *Let* Γ *be a connected, locally finite graph, and let* $x, y \in V(\Gamma) \sqcup \Omega(\Gamma)$. *Fix* $N \geq 1$. *Then* $\mathrm{es}(x, y) \geq N$ *if and only if there exists* N *pairwise edge-disjoint paths connecting* x *to* y. *Similarly, we also have that* $\mathrm{vs}(x, y) \geq N$ *if and only if there exists* N *pairwise internally vertex-disjoint paths connecting* x *to* y.

This was first proven for finite graphs by MENGER in 1927 [**98**]. It was later extended to locally finite, infinite graphs by HALIN in 1974 [**72**]. This extension follows fairly straightforwardly from a standard compactness argument. Note that one cannot drop the locally finite hypothesis as counterexamples exist. For example, consider a bi-infinite path L plus one extra vertex v, and add an edge between v and every vertex in L. The resulting graph Γ is two-ended and $\mathrm{vs}(\Gamma) = 2$, but Γ does not contain two disjoint bi-infinite paths.

COROLLARY IV.1.6. *Let Γ be a connected, locally finite graph, and suppose that* $\mathrm{vs}(\Gamma) \geq N$ *for some $N \geq 1$. Then between any two ends $\omega_1, \omega_2 \in \Omega(\Gamma)$ there exists a collection of N pairwise disjoint bi-infinite rays connecting ω_1 to ω_2.*

IV.1.3. Quasi-isometries between graphs. We now spend some time proving some preliminary lemmas relating specifically to quasi-isometries between graphs.

PROPOSITION IV.1.7. *Let Γ, Π be connected graphs, and $\psi : \Gamma \to \Pi$ a (not necessarily continuous) quasi-isometric embedding. Then there exists a subgraph $\Lambda \subset \Pi$ and a surjective, continuous quasi-isometry $\varphi : \Gamma \twoheadrightarrow \Lambda$. The inclusion $\Lambda \hookrightarrow \Pi$ is also a quasi-isometric embedding. Moreover, if Γ is bounded valence then so is Λ.*

REMARK IV.1.8. Note that the map $\varphi : \Gamma \twoheadrightarrow \Lambda$ is a quasi-isometry with respect to intrinsic metric d_Λ on Λ, not the restriction of the ambient metric d_Π where the statement is entirely trivial. The difference is subtle but important.

PROOF. We construct φ as follows. For every $v \in V(\Gamma)$, define $\varphi(v) = \psi(v)$. For every edge $e \in E(\Gamma)$, with endpoints a, b, set $\varphi(e)$ to a geodesic path between $\psi(a)$ and $\psi(b)$. Let $\Lambda = \varphi(\Gamma)$. We have that φ is thus a continuous, surjective map $\varphi : \Gamma \twoheadrightarrow \Lambda$. We claim that this is a quasi-isometry. Indeed, note that for all $x, y \in V(\Gamma)$ we have by construction that

$$\mathrm{d}_\Lambda(\varphi(x), \psi(y)) = \mathrm{d}_\Pi(\varphi(x), \psi(y)).$$

It follows that $\varphi : \Gamma \twoheadrightarrow \Lambda$ is a quasi-isometric embedding since ψ is, and thus φ is a quasi-isometry. It is also clear from this construction that the inclusion $\Lambda \hookrightarrow \Pi$ is also a quasi-isometric embedding.

Suppose now that Γ is bounded valence. We claim that Λ is also bounded valence. Let $d \geq 0$ be such that every vertex in Γ has degree at most d. Fix $\lambda \geq 1$ such that φ is a (λ, λ)-quasi-isometry. Fix $v \in V\Lambda$, and let $u \in \varphi^{-1}(v)$. Let $v_1, v_2 \ldots$ be the neighbours of v in Λ, and for each v_i choose some point $u_i \in \varphi^{-1}(v_i)$. Note that each u_i lies inside a bounded neighbourhood N of u. Since Γ is bounded valence, N intersects a bounded number of edges, say M. Recall that φ maps edges to geodesics between the images of their endpoints. This implies that if e is a (closed) edge in Γ, then $\varphi(e)$ intersects at most $2\lambda + 1$ distinct vertices in Λ. Thus, the image N in Λ can contain at most $k = M(2\lambda + 1)$ distinct vertices. Thus, v has at most k neighbours in Λ. Since v was arbitrary and all our bounds were uniform, it follows that Λ is bounded valence. $\square$

A similar construction to that given in Proposition IV.1.7 also gives the following.

PROPOSITION IV.1.9. *Let Γ, Π be connected graphs and $\psi : \Gamma \to \Pi$ a quasi-isometry. Then there exists a continuous quasi-inverse $\varphi : \Pi \to \Gamma$.*

Continuity of quasi-isometries is helpful for us, as it ensures that suitable restrictions of quasi-isometries are also quasi-isometries onto their image.

PROPOSITION IV.1.10. *Let Γ, Π be bounded valence, connected graphs. Let $\psi : \Pi \to \Gamma$ be a continuous quasi-isometry. Let $\Lambda \subset \Pi$ be a connected subgraph such that the inclusion map $\iota : \Lambda \hookrightarrow \Pi$ is a quasi-isometric embedding. Then the restriction $\psi|_\Lambda : \Lambda \twoheadrightarrow \psi(\Lambda)$ is a quasi-isometry.*

PROOF. Let $x, y \in \Lambda$, and let p be a geodesic in Λ connecting x to y of length n. Then $\psi(p)$ contains at most at cn vertices, where $c > 0$ is some constant depending only on the quasi-isometry constants of ψ and the maximum degree of Π. As ψ is continuous, we deduce that $\psi(p)$ contains a path of between $\psi(x)$ and $\psi(y)$ of bounded length. This gives the upper bound. For the lower bound, simply note that $\mathrm{d}_{\psi(\Lambda)}(\psi(x), \psi(y)) \geq \mathrm{d}_\Gamma(\psi(x), \psi(y))$, and apply the fact that ι and ψ are quasi-isometric embeddings. $\square$

REMARK IV.1.11. Again, we stress that in the statement of Proposition IV.1.10 it is important to note that we consider $\psi(\Lambda)$ with its own intrinsic path metric, and not the metric induced by the ambient graph Π. This is why we require the continuity of ψ, as it is important for us that the image $\psi(\Lambda)$ be a connected graph.

We now study how cuts and quasi-isometries interact.

LEMMA IV.1.12. *Let $\psi : \Gamma_1 \to \Gamma_2$ be a quasi-isometry between connected, locally finite graphs. Then there exists $r, R \geq 0$, depending only on the error constants of ψ, such that the following holds:*

Let $S \subset V(\Gamma_1)$, and let $x, y \in [\Gamma_1 \setminus B_{\Gamma_1}(S; R)] \cup \Omega(\Gamma_1)$ lie in distinct components of $\Gamma_1 \setminus S$. Then $\psi(x)$ and $\psi(y)$ lie in distinct components of $\Gamma_2 \setminus B_{\Gamma_2}(\psi(S); r)$.

In simple terms, this lemma tells us that if S separates two vertices/ends x and y (which aren't too close to S), then a uniform neighbourhood $\psi(S)$ separates $\psi(x)$ and $\psi(y)$. This is just a coarse version of the trivial statement that isometries send 'separating sets to separating sets'.

PROOF. Let φ be some choice of quasi-inverse. Fix $\lambda > 1$ sufficiently large such that both ψ and φ are (λ, λ)-quasi-isometries and φ is a λ-quasi-inverse to ψ. Note that such a $\lambda \geq 1$ depending only on the original error constants for ψ clearly exists. Let $r = 100\lambda^5$, and $R = \lambda(r + \lambda) + 1$. Let $x, y \in [\Gamma_1 \setminus B_{\Gamma_1}(S; R)] \cup \Omega(\Gamma_1)$ lie in distinct components of $\Gamma \setminus S$. To ease notation, let $S' = B_{\Gamma_2}(\psi(S); r)$, $x' = \psi(x)$, $y' = \psi(y)$.

Suppose that x', y' lie in the same component of $\Gamma_2 \setminus S'$. Then there is some path p connecting x' to y'. Then the λ^2-neighbourhood of $\varphi(p)$ contains a path p' connecting x to y. Such a path must intersect S. Routine calculation then reveals that p must therefore intersect the r-neighbourhood of $\psi(S)$. Since p was arbitrary, it follows that S' must separate x' from y', unless one of x' or y' lie in S'. However, by our choice of R, we also have that x' and y' necessarily lie outside of S'. The lemma follows. □

We now apply Lemma IV.1.12 to graphs equipped with quasi-actions. The following will be helpful. It essentially says that, in the presence of a cobounded quasi-action, if we delete a finite set S of vertices in our graph, and the remaining graph contains a finite connected component Z, then Z cannot be 'too big' compared to S. More precisely, we have the following statement.

Lemma IV.1.13. *Let Γ be a connected, locally finite graph equipped with a cobounded quasi-action by a group G. Then there exists $C \geq 1$ such that the following holds. If $S \subset V(\Gamma)$ is a finite set of vertices, and $Z \subset \Gamma$ is a finite connected component of $\Gamma \setminus S$, then*

$$\mathrm{d}_\Gamma(z, S) < C \,\mathrm{diam}_\Gamma(S) + C,$$

for all $z \in Z$.

Proof. The rough idea is that we choose a point z in Z which lies a maximal distance from S. If z is too far away from Z, then we have room to apply Lemma IV.1.12 and use the quasi-action to 'quasi-translate' a copy of S inside of Z near the point z. The image of z under this quasi-isometry, denoted y, will then threaten to lie inside Z but be further from S than z was, presenting us with a contradiction to the maximality the point z. We now fill in the details of this sketch below.

Fix $\lambda \geq 1$ such that the quasi-action of G on Γ is a λ-quasi-action, and B such that this quasi-action is B-cobounded. Fix $r, R \geq 0$ as in Lemma IV.1.12. Let S be a finite set of vertices and Z a finite connected component of $\Gamma \setminus S$. Since Z is finite, we may choose $z \in Z$ as to maximise $D := \mathrm{d}_\Gamma(z, S)$. Assume without loss of generality that $D > R$. Fix $s \in S$, and choose $g \in G$ such that $\mathrm{d}_\Gamma(z, \varphi_g(s)) \leq B$. Let ω be any end in Γ, and let $\xi = g\omega$. Let $y = \varphi_g(z)$. We have that $B_\Gamma(\varphi_g(S); r)$ separates y from ξ. To ease notation, write $S' = B_\Gamma(\varphi_g(S); r)$. Since D is assumed to be large, we have that S' is contained entirely within Z. In fact, S' lies outside the $\mathrm{diam}_\Gamma(S)$-neighbourhood of S, so every vertex in S' lies in the same connected component of $\Gamma \setminus S$, and so there is exactly one infinite connected component of $\Gamma \setminus S'$. Thus, for S' to separate y from ξ we must have that y lies in Z, and S' separates y from S.

Since z was chosen to lie a maximal distance from S, we see that $\mathrm{d}_\Gamma(y, S) \leq \mathrm{d}_\Gamma(z, S)$. Every path connecting y to S must pass through S', so we deduce that

$$\begin{aligned}\mathrm{d}_\Gamma(y, S) &\geq \mathrm{d}_\Gamma(y, S') + \inf_{s \in S} \mathrm{d}_\Gamma(s, S') \\ &\geq [\tfrac{1}{\lambda} D - \lambda - r] + [D - B - \lambda \,\mathrm{diam}_\Gamma(S) - \lambda - r].\end{aligned}$$

Combining and simplifying the above, we deduce that

$$D \leq \lambda^2 \operatorname{diam}_\Gamma(S) + \lambda(2\lambda + 2r + B).$$

Thus, by setting $C = \lambda(2\lambda + 2r + B)$ we are done. □

If we assume our graph has more than one end then we can say even more.

LEMMA IV.1.14. *Let Γ be a connected, locally finite graph equipped with a cobounded quasi-action by a group G. Suppose further that Γ has more than one end. Then there exists $C \geq 1$ such that the following holds. If $Z \subset \Gamma$ is a finite (not necessarily connected) subgraph such that $\Gamma \setminus Z$ is connected, then*

$$\mathrm{d}_\Gamma(z, \Gamma \setminus Z) < C,$$

for all $z \in Z$.

PROOF. Fix $\lambda > 1$ such that the quasi-action of G on Γ is a λ-quasi-action, and B such that this quasi-action is B-cobounded. Fix $r \geq 0$ as in Lemma IV.1.12.

Since Γ is multi-ended, there exists a finite connected subgraph $K \subset \Gamma$ such that $\Gamma \setminus K$ has at least two infinite components. By Lemma IV.1.12 we have that $N_g := B_\Gamma(\varphi_g(K); r)$ separates at least two infinite components of Γ, for all $g \in G$. Now, let Z be a finite subgraph of Γ such that $\Gamma \setminus Z$ is connected. Suppose that there exists some $z \in Z$ such that

$$\mathrm{d}_\Gamma(z, \Gamma \setminus Z) \geq \lambda \operatorname{diam}_\Gamma(K) + \lambda + 2r + B.$$

Then there exists some $g \in G$ such that N_g lies entirely within Z. If r_1, r_2 are any two infinite rays in $\Gamma \setminus N_g$, then clearly both rays will leave Z at some point. In particular, since $\Gamma \setminus Z$ is connected we have that r_1, r_2 lie in the same component of $\Gamma \setminus N_g$. Thus, there is exactly one infinite connected component of $\Gamma \setminus N_g$, a contradiction. The lemma therefore follows by setting $C = \lambda \operatorname{diam}_\Gamma(K) + \lambda + 2r + B$. □

Finally we note application of Lemma IV.1.13. Firstly, we will state two definitions. The first is standard across graph theory.

DEFINITION IV.1.15 (N-connected graph). Let Γ be a graph, and $N > 0$ a natural number. We say that Γ is *N-connected* if Γ is connected, and the removal of any collection of at most N vertices does not disconnect Γ.

The second definition is not standard, but helpful in our coarse setting.

DEFINITION IV.1.16 (Almost 2-connected graph). Let Γ be a infinite, connected, locally finite graph. We say that Γ is *almost 2-connected* if there exists a unique maximal 2-connected infinite subgraph $\Gamma_0 \subset \Gamma$ and the inclusion map $\Gamma_0 \hookrightarrow \Gamma$ is a quasi-isometry. The subgraph Γ_0 is called the *2-connected core* of Γ.

Note that every maximal 2-connected subgraph is necessarily isometrically embedded. This means that checking the quasi-isometry condition in the above definition amounts to just checking that the inclusion is coarsely surjective. Intuitively, an almost 2-connected graph Γ is obtained from a 2-connected graph Γ_0 by 'attaching' a selection of boundedly small finite graphs at vertices.

LEMMA IV.1.17. *Let* Γ *be a infinite, connected, locally finite graph equipped with a cobounded quasi-action. Suppose* $\mathrm{vs}(\Gamma) > 1$. *Then* Γ *is almost 2-connected.*

PROOF. Since $\mathrm{vs}(\Gamma) > 1$, no cut vertex in Γ separates ends. It follows from Lemma IV.1.13 that there exists some uniform constant $C \geq 0$ such that there is exactly one infinite, maximal, 2-connected subgraph Λ of Γ, and every other such subgraph has diameter at most C. The inclusion map $\Lambda \hookrightarrow \Gamma$ is clearly coarsely surjective, and thus a quasi-isometry. $\square$

IV.1.4. Planar graphs. Recall that a *topological embedding* between topological spaces X, Y is a continuous injection $f : X \hookrightarrow Y$ which restricts to a homeomorphism between X and $f(X)$ (with the subspace topology). We adopt the following convention.

DEFINITION IV.1.18 (Planar graph). Let Γ be a graph. We say that Γ is *planar* if there exists topological embedding $\Gamma \hookrightarrow \mathbf{S}^2$. Such an embedding is called a *drawing*.

The above leaves room for strange embeddings. An embedding $\Gamma \hookrightarrow \mathbf{S}^2$ is called *pointed* if for any 1-way infinite simple ray r in Γ, the images of the vertices in r accumulate at exactly one point in $\mathbf{S}^2$. The closure of the image of a planar graph in $\mathbf{S}^2$ can be viewed as an embedding of a particular compactification of Γ. In this chapter, we will generally restrict ourselves to the case where this closure is precisely the Freudenthal compactification of Γ. This can generally be ensured thanks to the following, due to RICHTER–THOMASSEN [**115**].

PROPOSITION IV.1.19 ([**115**, Lem. 12]). *Let* Γ *be a 2-connected locally finite planar graph. Then the Freudenthal compactification* $|\Gamma|_{\mathrm{Fr}}$ *embeds into* $\mathbf{S}^2$.

This result can be pushed further, and the 2-connected assumption may be dropped. The following argument was suggested by B. RICHTER.[1]

PROPOSITION IV.1.20. *Let* Γ *be a connected, locally finite, planar graph. Then the Freudenthal compactification* $|\Gamma|_{\mathrm{Fr}}$ *embeds into* $\mathbf{S}^2$.

PROOF. Fix an embedding $\vartheta : \Gamma \hookrightarrow \mathbf{S}^2$. We will 'thicken-up' Γ to construct a 2-connected, locally finite planar Π such that (some subdivision of) Γ is a subgraph of Π and this inclusion is a quasi-isometry. From there, the embedding of $|\Pi|_{\mathrm{Fr}}$ into $\mathbf{S}^2$ given by Proposition IV.1.19 will induce an embedding of $|\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$. Intuitively, our plan is to add a small cycle surrounding each cut

[1]Personal communication.

vertex. In order to make this sketch formal, we will construct Π as an ascending union of finite planar graphs, which is necessarily planar. Indeed, if an ascending union of finite planar graphs contained some subdivision of K_5 or $K_{3,3}$ then certainly some finite graph in this chain would have to, which is contradicts KURATOWSKI's theorem.

We first assume that Γ contains no cut edges. That is, for every $e \in E(\Gamma)$ we have that $\Gamma \setminus e$ is connected. Choose some root $v_0 \in V(\Gamma)$. We call $u \in V(\Gamma)$ a *marked vertex* if it is a cut vertex of Γ. Given $r \geq 0$ we construct Π_i as follows. Consider the closed r-neighbourhood of v_0 in Γ, which we will denote $N_r = B_\Gamma(v_0; r)$. This is a finite subgraph of Γ. Let N_r' denote the subdivision of N_r, where we subdivide every edge into three edges. Suppose $u \in VN_r$ is a marked vertex of Γ, and $\mathrm{d}_\Gamma(v_0, u) < r$. Let u' be the image of u in N_r'. Let $u_1, \ldots, u_n$ be the neighbours of u' in VN_r'. Since Γ is locally finite, n is finite. Let e_i be the (necessarily unique) edge of N_r' with endpoints u', u_i. The drawing ϑ restricts to a drawing of N_r, and thus of N_r' since these graphs are homeomorphic. This induces a cyclic ordering to $e_1, \ldots, e_n$, particularly their clockwise ordering about u'. We assume this is exactly how they are ordered. We add a path of length 2 joining every u_i to u_{i+1}, with indices taken modulo n. We repeat this for every such u', and call the resulting graph Π_r. By construction Π_r is a finite planar graph. Repeating this construction for every r, it is clear that we have an ascending chain

$$\Pi_0 \subset \Pi_1 \subset \Pi_2 \subset \cdots .$$

Let $\Pi = \bigcup_r \Pi_r$. Let Γ' denote the subdivision of Γ, where each edge is divided into three edges. Clearly Γ is homeomorphic to Γ', and the natural map $\Gamma \to \Gamma'$ is a $(3, 0)$-quasi-isometry. We have a natural inclusion $\Gamma' \hookrightarrow \Pi$ which is an isometry onto its image (indeed, the new paths added to Π) create no new shortcuts. Every vertex of Π is adjacent to a vertex of Γ', so this map is a quasi-isometry. Finally, note that Π is certainly 2-connected. For suppose Π contained a cut vertex w. This vertex certainly must lie in Γ', as every $u \in V(\Pi) \setminus V(\Gamma)'$ has degree 2 and lies on a simple cycle so cannot be a cut vertex. Since Γ contains no cut edges this implies that w is induced by a cut vertex of Γ. But then by construction if u_1, u_2 are neighbours of w in Π then there is a path connecting them which avoids w. Thus Π is 2-connected. Applying Proposition IV.1.19 we conclude that the Freudenthal compactification of any connected, locally finite, planar graph *without cut edges* embeds into $\mathbf{S}^2$.

Finally, we deal with the case where Γ has cut edges. We replace Γ with Γ' where Γ' is obtained from Γ by 'doubling' every edge. That is, $V(\Gamma') = V(\Gamma)$, and if u, v in Γ are connected by k edges in Γ then they are connected by $2k$ edges in Γ. The inclusion $\Gamma \hookrightarrow \Gamma'$ is certainly a quasi-isometry, and Γ' is clearly a locally finite. To see that Γ' is planar, note that doubling edges clearly preserves planarity in finite graphs. We now repeat a similar construction to above, and write Γ' as an ascending union of finite planar graphs, where the r-th element of this chain is obtained by doubling the edges in the r-ball about some root vertex v_0 in Γ. By the previous case, the Freudenthal compactification of Γ'

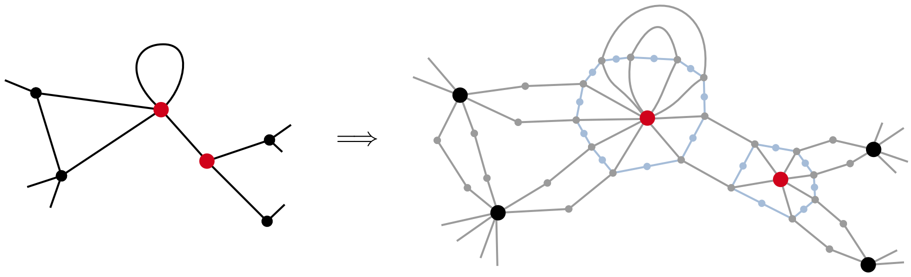

FIGURE IV.1. Constructing a 2-connected planar super-graph. The marked vertices are depicted in red.

is planar, and so the same can be said about Γ. A cartoon of this full construction can be found in Figure IV.1. $\square$

In light of Proposition IV.1.20, we introduce the following terminology for the sake of brevity.

DEFINITION IV.1.21 (Good drawing). Let Γ be a connected, locally finite, planar graph. Then a *good drawing* of Γ is an embedding $\vartheta : |\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$.

Some planar graphs admit even nicer drawings in the plane, in the following sense. We say that a planar graph Γ admits a *vertex accumulation point-free drawing*, or *VAP-free drawing*, if there exists a topological embedding $\vartheta : \Gamma \hookrightarrow \mathbf{R}^2$ such that $\vartheta(V(\Gamma))$ is a discrete subset of $\mathbf{R}^2$. The following is a standard fact and follows easily from Proposition IV.1.20.

COROLLARY IV.1.22. *Let Γ be a connected, locally finite, one-ended, planar graph. Then Γ admits a VAP-free drawing.*

REMARK IV.1.23. Not every planar graph admits a VAP-free drawing. For example, the Cayley graph of the one-relator group

$$\mathbf{Z}^2 * \mathbf{Z} = \langle a, b, c \ ; \ [a, b] \rangle$$

is planar (it is a 'tree of flats') but admits no such drawing. Proving this is an easy exercise in the Jordan curve theorem.

An important feature of planar graphs is their *faces*. Informally, it is clear what we mean by a 'face'. However, this word could be referring to several different but closely related concepts, especially in the realm of infinite planar graphs. Is a 'face' a component of $\mathbf{S}^2 \setminus \vartheta(|\Gamma|_{\mathrm{Fr}})$, or actually a particular form of subgraph of Γ? For the sake of clarity, we now set up some notation to help us discuss faces without ambiguity.

DEFINITION IV.1.24 (Faces and facial subgraphs). Let Γ be a connected, locally finite, planar graph with a fixed good drawing ϑ. The connected components of $\mathbf{S}^2 \setminus \vartheta(|\Gamma|_{\mathrm{Fr}})$ are referred to as the *faces* of Γ, with respect to the drawing ϑ. Let $\mathcal{D}(\Gamma)$ denote the set of faces. Given $U \in \mathcal{D}(\Gamma)$, write

$$\mathcal{F}[U] := \vartheta^{-1}(\partial U) \setminus \Omega(\Gamma).$$

We call the subgraph $\mathcal{F}[U]$ the *facial subgraph of* Γ *bordering* U. Let $\mathcal{F}(\Gamma)$ denote the set of facial subgraphs of Γ.

$$\mathcal{F}^{\infty}(\Gamma) = \{f \in \mathcal{F}(\Gamma) : f \text{ is infinite}\}, \quad \mathcal{F}^{c}(\Gamma) = \{f \in \mathcal{F}(\Gamma) : f \text{ is compact}\}.$$

Remark IV.1.25. Note that $f \in \mathcal{F}^{\infty}(\Gamma)$ need not be connected in general. However, the closure of f in $|\Gamma|_{\mathrm{Fr}}$ will always be connected.

Remark IV.1.26. To be completely rigorous, we should really include mention of the drawing ϑ in our notation, as the sets $\mathcal{F}(\Gamma)$ etc. depend not just on Γ but also on ϑ. However, for our purposes this drawing will always be fixed in advance and so there is no risk of confusion.

There is a key benefit to working with 2-connected planar graphs, which is that their faces are always cycles, in some sense. More generally, we have the following.

Proposition IV.1.27 ([**115**, Prop. 3]). *Let K be a compact, 2-connected, locally connected subset of the sphere. Then the boundary of every component of $\mathbf{S}^2 \setminus K$ is a simple closed curve.*

In particular, if Γ is a 2-connected locally finite planar graph then the closure of every $f \in \mathcal{F}(\Gamma)$ in $|\Gamma|_{\mathrm{Fr}}$ is a simple closed curve. More precisely, if $f \in \mathcal{F}^{c}(\Gamma)$ then f is a simple loop, and if $f \in \mathcal{F}^{\infty}(\Gamma)$ then f is a disjoint union of bi-infinite lines.

IV.2. One-ended groups

IV.2.1. Discussion. Recall that the ultimate goal of this chapter is to prove that a finitely generated group which is quasi-isometric to a planar graph is virtually planar. In this section, we first deal with the one-ended case. This is fairly self-contained, and serves as a good demonstration of the techniques used in this chapter.

Our strategy is to show that such a group is necessarily quasi-isometric to a complete Riemannian plane, and thus conclude by Theorem I.4.1 stated in the introduction. Now, if we want to show that a (2-connected, locally finite) one-ended planar graph Γ is quasi-isometric to a complete Riemannian plane, then our natural instinct is glue 2-cells into the faces and extend the graph metric on the graph to a Riemannian metric on the resulting plane. The following two pathologies could arise, which will halt this plan in its tracks.

(1) There could be 'infinite face paths', so the resulting complex is not a plane.
(2) The finite faces of Γ could be arbitrarily big, which will stop the inclusion of our graph into the constructed Riemannian surface from being a quasi-isometry.

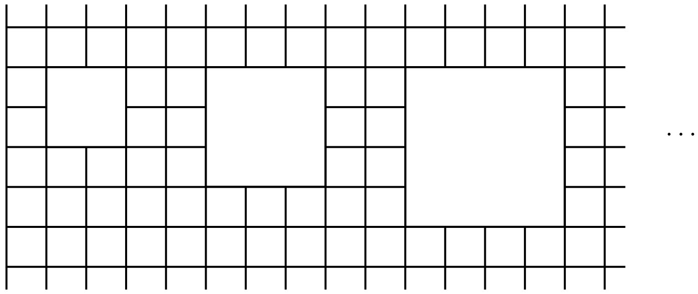

FIGURE IV.2. One-ended planar graph which is not quasi-isometric to a complete Riemannian plane. It is not immediately clear whether or not this graph is quasi-isometric to a finitely generated group.

Consider, for example, the one-ended planar graph depicted in Figure IV.2. This graph is obtained by taking a half-grid and removing bigger and bigger 'holes'. It is not immediately clear that this graph is not quasi-isometric to some finitely generated group.

Our strategy is thus to prove that neither of the two pathologies described above can occur in a one-ended planar graph which is quasi-isometric to a Cayley graph. We show this by studying the induced quasi-action on the planar graph, and using this to obtain some control over the local features of the graph.

IV.2.2. Quasi-actions on one-ended planar graphs. Throughout this section, let $X = \mathrm{Cay}(G,S)$ be a Cayley graph of a finitely generated group G. Let Γ be a connected planar graph, and let $\varphi : X \to \Gamma$ be a quasi-isometry with quasi-inverse $\psi : \Gamma \to X$. We will assume that Γ is bounded valence, and that both φ and ψ are continuous; indeed, Propositions IV.1.7, IV.1.9 demonstrate that these assumptions are completely inconsequential. We may also assume that Γ is 2-connected, by an application of Lemma IV.1.17. Thus, every $f \in \mathcal{F}(\Gamma)$ is either a simple cycle or a disjoint union of bi-infinite rays. The connected components of infinite faces shall be referred to as *bi-infinite face paths* in this section.

By Corollary IV.1.22, Γ admits a VAP-free drawing $\vartheta : \Gamma \hookrightarrow \mathbf{R}^2$. Fix $\lambda, B \geq 1$ be such that the induced quasi-action of G on Γ is a B-cobounded λ-quasi-action. Recall the following standard construction.

LEMMA IV.2.1. *Given any $v \in X$, there exists three geodesic rays γ_1, γ_2, γ_3 based at v such that*

$$\mathrm{d}_X(\gamma_i(n), \gamma_j(m)) \to \infty,$$

as $n, m \to \infty$, for any distinct $i, j = 1, 2, 3$.

We now apply the above and push the resulting rays through the quasi-isometry φ, obtaining a similar feature in Γ.

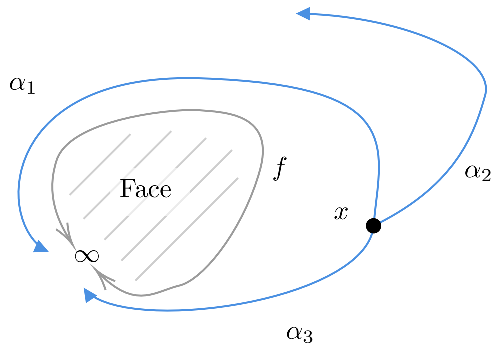

FIGURE IV.3. The ray α_2 is forced to diverge from f due to the Jordan curve theorem.

LEMMA IV.2.2. *There exists a constant $r \geq 0$ such that for every $x \in \Gamma$ there exists three distinct quasi-geodesic rays $\alpha_1, \alpha_2, \alpha_3 : [0, \infty) \to \Gamma$ based in $B_\Gamma(x; r)$. Moreover, we may assume that the α_i satisfy the following further property: for every $i \neq j$, $\mathrm{d}_\Gamma(z, \alpha_i) > \lambda^2$ for every $z \in \alpha_j$.*

PROOF. Let $y = \psi(x)$. Using Lemma IV.2.1 construct geodesic rays γ_1, γ_2, γ_3 based at y which pairwise diverge. Let $M > 0$ be large. Then there is some $N \geq 0$ such that for all $n, m \geq N$ we have that

$$\mathrm{d}_X(\gamma_i(n), \gamma_j(m)) > M.$$

By transitivity we may assume that this N does not depend on the choice of y or x.

Let $\alpha_i = \varphi \circ \gamma_i|_{[N,\infty)}$. If we choose M to be sufficiently large with respect to λ, it will be clear that the λ^2-neighbourhoods of the α_i are disjoint. The lemma follows. □

LEMMA IV.2.3. *Γ contains at most one bi-infinite face path.*

PROOF. Suppose there exists two distinct such bi-infinite face paths f_1 and f_2. Pick $x_i \in f_i$ and let ℓ be a geodesic between x_1 and x_2. It is easy to see that ℓ is a compact subset which separates Γ into at least two infinite connected components. This contradicts the fact that Γ is one-ended. □

LEMMA IV.2.4. *If Γ contains a bi-infinite face path f, then* $\mathrm{Haus}_\Gamma(f, \Gamma) = \infty$.

PROOF. Assume there is an infinite face f. Take a vertex x, and apply Lemma IV.2.2 to obtain three quasi-geodesic rays $\alpha_1, \alpha_2, \alpha_3 : [0, \infty) \to \Gamma$ which are based at x, and which pairwise diverge. Two of these rays, say α_1, α_2, will together form a Jordan curve (if we include the point at infinity) which separates α_3 from the infinite face f. Any path from α_3 to f must cross through α_1 and α_2. Since α_3 diverges from these two rays, it also diverges from f. See Figure IV.3 for a cartoon.

□

LEMMA IV.2.5. *There exists some uniform constant $r \geq 0$ such that every finite face cycle of Γ has length at most n.*

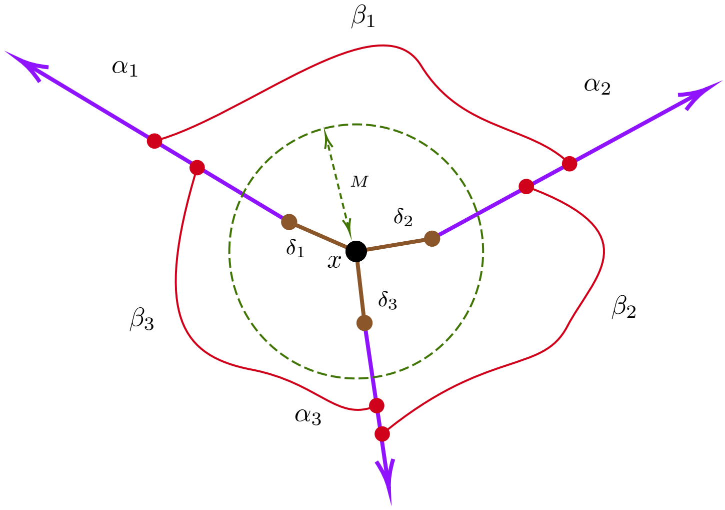

FIGURE IV.4. Our choice of the paths β_i traces out a Jordan curve separating x from infinity.

PROOF. If Γ has no infinite face then choose $x \in \Gamma$ arbitrarily. Otherwise, apply Lemma IV.2.4 and choose $x \in \Gamma$ such that x lies at least $M = 1000\lambda^{1000}Br$ from this infinite face, where r is the constant given in Lemma IV.2.2.

Apply Lemma IV.2.2 and obtain rays $\alpha_1, \alpha_2, \alpha_3 : [0, \infty) \to \Gamma$ based in $B_\Gamma(x; r)$ with disjoint λ^2-neighbourhoods. Let δ_i be a geodesic from α_i to x. Since there is no infinite face near x, we may connect α_i to α_{i+1} by a path β_i which stays inside the region bounded by $\alpha_i \cup \alpha_{i+1} \cup \delta_i \cup \delta_{i+1}$ not containing α_{i+2} (where indices are taken modulo 3). We can also assume that each β_i is disjoint from the M-neighbourhood of all of the δ_j. The subgraph $\alpha_1 \cup \alpha_2 \cup \alpha_3 \cup \beta_1 \cup \beta_2 \cup \beta_3$ contains a Jordan curve J which separates x from infinity, which is disjoint from the M-neighbourhood of x. See Figure IV.4 for a cartoon.

Suppose now we translate this figure somewhere using our quasi-action. Fix $g \in G$. To ease notation, let us denote the quasi-isometry φ_g with the following shorthand:

$$\varphi_g(x) = x'.$$

We have that $\alpha_i' \cap \alpha_j' = \emptyset$ for all $i \neq j$, and also that

$$\beta_i' \cap (\alpha_{i+2}' \cup \delta_{i+2}') = \emptyset \tag{10}$$

for each $i = 1, 2, 3$. Each $\alpha_i' \cup \delta_i'$ is still a quasi-geodesic ray based at x', heading towards the lone end of Γ. If we try to draw the β_i', then the equation (10) forces us to once again trace a Jordan curve separating x' from the end of Γ. In particular, we see that J' still separates x' from infinity. We have that J' is certainly disjoint from the B-neighbourhood of x'. Let S be the vertex set of J'. Then S has diameter in Γ which depends only on J and the quasi-action, say $D \geq 1$. We have that S separates x' from the end of Γ. Let Z be the finite connected component of $\Gamma \setminus S$ containing x'. Any face of Γ which intersects the B-neighbourhood of x' cannot 'cross' J', and so is contained in

Z. Applying Lemma IV.1.13, we see that Z can only contain boundedly many vertices. This puts a uniform upper bound on the lengths of such faces. Since $g \in G$ was chosen arbitrarily and the quasi-action is B-cobounded, the lemma follows. □

LEMMA IV.2.6. *The planar graph Γ contains no infinite facial subgraphs.*

PROOF. The method to prove this is very similar to the proof of Lemma IV.2.5, so we will only sketch the argument.

Suppose there were some $f \in \mathcal{F}^{\infty}(\Gamma)$. Then by Lemma IV.2.4 we can choose some $x \in \Gamma$ which is arbitrarily far from f. We apply the same construction as in the proof of Lemma IV.2.5 and obtain a Jordan curve J separating x from infinity which is disjoint from the M neighbourhood of x, where M is taken as in the proof of Lemma IV.2.5. As before, we may translate this Jordan curve anywhere using the quasi-action of G. The curve $\varphi_g(J)$ still separates $y = \varphi_g(x)$ from infinity in the plane, and lies disjoint from the B-neighbourhood of y. But then the B-neighbourhood of y cannot contain an infinite face. Since this quasi-action is B-cobounded and $g \in G$ was chosen arbitrarily, we are forced to conclude that there is no infinite face in Γ. □

THEOREM IV.2.7. *Let X be a one-ended Cayley graph graph. Suppose X is quasi-isometric to a planar graph. Then X is quasi-isometric to a complete Riemannian plane.*

PROOF. By Lemmas IV.2.5, IV.2.6, we may assume that Γ contains no infinite face paths, and that every face of Γ has bounded length. It is then clear that if we attach a 2-cell along each face path, then the resulting cell complex K is homeomorphic to $\mathbf{R}^2$. Subdivide the 2-cells of K and obtain a triangulation, so that the inclusion of Γ into the 1-skeleton K^1 is an isometric embedding. Since the 2-cells of K have boundedly small boundaries, we only need to subdivide each cell into a bounded number of triangles. Thus, the inclusion $\Gamma \hookrightarrow K^1$ is coarsely surjective and thus a quasi-isometry. We now extend the metric on K^1 to a metric on K, and obtain a piecewise-linear complete plane which is quasi-isometric to Γ. We then modify this metric locally into a smooth Riemannian metric on K. □

The corollary below follows immediately from combining Theorem IV.2.7 with Theorem I.4.1.

COROLLARY IV.2.8. *Let G be a finitely generated, one-ended group. Suppose G is quasi-isometric to a planar graph. Then G is a virtual surface group.*

IV.3. Coboundary diameters and cutting up graphs

IV.3.1. Inner and outer coboundary diameters. Let Γ be a graph and $\Lambda \subset \Gamma$ a subgraph. Define the *coboundary of* Λ, denoted $\delta\Lambda$, as the set of edges with one endpoint in Λ and the other in $\Gamma \setminus \Lambda$. We will often abuse notation and view $\delta\Lambda$ itself as a subspace of Γ.

Definition IV.3.1 (Coboundary diameters). Let Γ be a connected graph and $\Lambda \subset \Gamma$ a connected subgraph. Define the *inner coboundary diameter* of Λ, denoted $\|\Lambda\|_{\mathrm{in}}$, as

$$\|\Lambda\|_{\mathrm{in}} = \sup\{\mathrm{diam}_\Lambda(\delta U \cap \Lambda) : U \text{ is a connected component of } \Gamma \setminus \Lambda\}.$$

Similarly, we define the *outer coboundary diameter* of Λ as

$$\|\Lambda\|_{\mathrm{out}} = \sup\{\|U\|_{\mathrm{in}} : U \text{ is a connected component of } \Gamma \setminus \Lambda\}.$$

We say that Λ has *uniform coboundary* if both the inner and outer coboundary diameters of Λ are finite. If $\Lambda = \Gamma$, we adopt the convention that $\|\Lambda\|_{\mathrm{in}} = \|\Lambda\|_{\mathrm{out}} = 0$.

We may write $\|\Lambda\|_{\mathrm{in}}^\Gamma$ or $\|\Lambda\|_{\mathrm{out}}^\Gamma$ if the intended super-graph Γ is ever unclear.

Intuitively, we imagine the coboundary of Λ as a selection of tight cuts, separating Λ from the components of its complement. The inner diameter measures the size of these cuts from 'inside' Λ, whereas the outer diameter measures this diameter from the other side of this cut. Without placing further restrictions on Λ, there is no reason these values need to correlate. However, we can say something in the presence of an appropriate group action.

Definition IV.3.2 (Cocompactly stabilised subgraph). Let Γ be a connected graph and G a group acting on Γ. We say that a connected subgraph $\Lambda \subset \Gamma$ is *cocompactly stabilised* if the set-wise stabiliser $\mathrm{Stab}(\Lambda) \leq G$ acts cocompactly on Λ.

Lemma IV.3.3. *Let G be a group acting on a connected, bounded valence graph Γ. Let $\Lambda \subset \Gamma$ be a connected, cocompactly stabilised subgraph. If $\|\Lambda\|_{\mathrm{in}}$ is finite, then Λ has uniform coboundary.*

Proof. Since $\mathrm{Stab}(\Lambda)$ acts on Λ cocompactly and Γ has bounded valence, we see that H acts with finite quotient on the set of connected components of $\Gamma \setminus \Lambda$. If U is a connected component of $\Gamma \setminus \Lambda$, then we have that δU is a finite set of edges, and so $\|U\|_{\mathrm{in}}$ is finite. Since there are only finitely many orbits of these components, we deduce that $\|\Lambda\|_{\mathrm{out}}$ must be finite, and so Λ has uniform coboundary. □

Proposition IV.3.4. *Let Γ be a locally finite graph and $\Lambda \subset \Gamma$ a connected subgraph with $\|\Lambda\|_{\mathrm{in}} < \infty$. Then the natural map $\Omega(\Lambda) \to \Omega(\Gamma)$ is injective.*

Proof. Let γ_1, γ_2 be rays in Λ. We need to show that if γ_1, γ_2 approach distinct ends in Λ, then they approach distinct ends in Γ. Let $K \subset \Lambda$ be a compact subgraph such that infinite subpaths of γ_1, γ_2 lie in distinct components of $\Lambda \setminus K$. Let U_i denote the component of $\Lambda \setminus K$ containing the tail of γ_i. Any path in Γ connecting U_1 to U_2 must pass through some connected component U of $\Gamma \setminus \Lambda$, where δU intersects both U_1 and U_2. Only finitely many such U exist since $\|\Lambda\|_{\mathrm{in}}$ is finite, and moreover δU can only contain finitely many edges. Thus, let

$$K' = K \cup \bigcup_U \delta U,$$

where U ranges over the aforementioned components of $\Gamma \setminus \Lambda$. Clearly K' is compact, and the tails of γ_1, γ_2 lie in distinct components of $\Gamma \setminus K'$. $\square$

REMARK IV.3.5. With Γ, Λ as above, i.e. $\|\Lambda\|_{\mathrm{in}} < \infty$, it follows from Proposition IV.3.4 that the closure of Λ in the Freudenthal compactification $|\Gamma|_{\mathrm{Fr}}$ of Γ is naturally homeomorphic to $|\Lambda|_{\mathrm{Fr}}$. In particular, if Γ is planar and $\vartheta : |\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$ is a good drawing of Γ, then the restriction of ϑ to (the closure of) Λ is a good drawing of Λ.

PROPOSITION IV.3.6. *Let Γ be a connected, locally finite graph, and let $\Lambda \subset \Gamma$ be a connected subgraph with uniform coboundary. Then the inclusion $\Lambda \hookrightarrow \Gamma$ is a quasi-isometric embedding.*

PROOF. Take a geodesic in Γ connecting $x, y \in V\Lambda$. Any segment of this geodesic lying outside of Λ has length at most $2 + \|\Lambda\|_{\mathrm{out}}$, and we also know that there exists a path inside of Λ of length at most $\|\Lambda\|_{\mathrm{in}}$, with precisely the same endpoints. Thus, such a geodesic can be transformed into a path contained in Λ whose length is proportional to the length of the original geodesic. It follows immediately that the inclusion $\Lambda \hookrightarrow \Gamma$ is a quasi-isometric embedding. $\square$

IV.3.2. Nested cuts and the Boolean ring. Let X be a connected graph. An *edge-cut*, or just *cut*, of X is a subset $F \subset E(X)$ such that $X \setminus F$ is disconnected. We say that a cut F is *finite* if $|F| < \infty$. A finite cut F is said to be *tight* if $X \setminus F$ contains exactly two connected components. Recall that here we are just removing the interiors of edges, not their end vertices too. Given a subset $b \subset V(X)$, let b^* denote the complement $V(X) \setminus b$. Let δb, called the *coboundary of* b, denote the set of edges in X with exactly one endpoint in b. The term 'coboundary' is suggestive of the cohomology at play here, since if we take a the 0-cochain in $C^0(X; \mathbf{Z}_2)$ supported by U, then δU is precisely the support of coboundary of this cochain. Given any $b \subset V(X)$, clearly δb is a cut and $\delta b = \delta b^*$. Let

$$\mathscr{B}(X) = \{b \subset V(X) : \delta b \text{ is finite}\}.$$

Clearly $\mathscr{B}(X)$ is closed under the operations of union, intersection, and complementation. This makes $\mathscr{B}(X)$ into a Boolean ring. That is, a commutative ring with unity such that every element r satisfies idempotent $r^2 = r$. The multiplicative operation is intersection, while the additive operation is symmetric difference, denoted $\triangle$ in this thesis.

Given a group G acting on X, this induces an action of G upon $\mathscr{B}(X)$. Thus, we may view $\mathscr{B}(X)$ as a G-module. Let $\mathscr{B}_n X$ denote the subring of $\mathscr{B}(X)$ generated by elements A such that $|\delta b| < n$. We may abuse terminology and say that $b \in \mathscr{B}(X)$ is *tight* if its coboundary δb is a tight cut. Given $b_1, b_2 \in \mathscr{B}(X)$, we say that b_1 *crosses* b_2 if the intersections

$$b_1 \cap b_2, \quad b_1 \cap b_2^*, \quad b_1^* \cap b_2, \quad b_1^* \cap b_2^*$$

are all non-empty. If b_1 and b_2 do not cross, we say they are *nested*. We say that a subset $\mathcal{E} \subset \mathscr{B}(X)$ is *nested* if it is closed under taking complements and any two $b_1, b_2 \in \mathcal{E}$ are nested. We now state the following key theorem due to Dicks and Dunwoody [**39**, II.2.20].

Theorem IV.3.7 (Dicks–Dunwoody). *Let X be a connected graph and G a group acting on X. Then there is a sequence $\mathcal{E}_1 \subset \mathcal{E}_2 \subset \dots$ of G-invariant nested subsets of $\mathscr{B}(X)$ consisting of tight elements, such that $\mathcal{E}_n$ generates $\mathscr{B}_n X$ as a Boolean ring.*

The following characterisation of accessibility is helpful, due to Thomassen and Woess [**129**].

Theorem IV.3.8 ([**129**, Thm. 7.6]). *Let X be a connected, locally finite graph equipped with a cocompact action by a group G. Then X is accessible if and only if there exists $n \geq 1$ such that $\mathscr{B}(X) = \mathscr{B}_n(X)$. In other words, X is accessible if and only if there is a nested, G-invariant, G-finite generating set of $\mathscr{B}(X)$.*

Recall that the set of directed edges of a tree has a natural ordering. The structure of nested sets can be encoded similarly in a tree via the following theorem, originally due to Dunwoody [**44**].

Theorem IV.3.9 ([**44**, Thm. 2.1]). *Let $(\mathcal{E}, \leq)$ be a partially ordered set equipped with an involution $* : \mathcal{E} \to \mathcal{E}$ satisfying the following:*

(1) *For any $A, B \in \mathcal{E}$, at least one of $A \leq B$, $A \leq B^*$, $A^* \leq B$, $A^* \leq B^*$ hold.*
(2) *If $A \leq B$ for $A, B \in E$, then $B^* \leq A^*$.*
(3) *Given $A, B \in \mathcal{E}$, at most one of $A \leq B$, $A \leq B^*$ hold. Similarly, at most one of $A \leq B$, $A^* \leq B$ hold.*
(4) *If $A \leq B$ for $A, B \in \mathcal{E}$, then there are at most finitely many $C \in \mathcal{E}$ such that $A \leq C \leq B$.*

Then there exists a tree $T = T(\mathcal{E})$ such that the set of oriented edges $\vec{E}(T)$ can be naturally identified with $\mathcal{E}$, and the order $\leq$ on $\mathcal{E}$ is precisely the order determined by edge-paths in T.

Note that condition (4) above ensures that the resulting tree T is simplicial. Without it, we would instead obtain an $\mathbf{R}$-tree.

We will refrain from giving a proof of the above theorem, but it is worth at least stating where the vertices of T come from. Define a relation $\sim$ on $\mathcal{E}$ as follows. Given $A, B \in \mathcal{E}$, say that $A \ll B$ if $A \leq B$ and $A \leq C \leq B$ implies $A = C$ or $B = C$. We then define

$$A \sim B \text{ if } A = B \text{ or } A \ll B^*.$$

It is an exercise to check that $\sim$ is an equivalence relation on $\mathcal{E}$. The vertex set of T can then be taken to be the set of $\sim$-equivalence classes. Those directed edges of T which point 'into' a vertex v are then precisely those $A \in \mathcal{E}$ such that $A \in v$. We call T the *structure tree* of $\mathcal{E}$. For more details, see [**44**].

Now, if the partially ordered set $\mathcal{E}$ above is equipped with an appropriate action by a group G, then we obtain an action of G on the structure tree $T(\mathcal{E})$. In particular, this is an invitation to employ the toolbox of Bass–Serre theory (see § I.5.5 for definitions). The proof of the following observation is essentially due to THOMASSEN and WOESS; see [**129**, § 7].

THEOREM IV.3.10. *Let G be a finitely generated group acting freely and cocompactly on a locally finite, connected graph X. Let $\mathcal{E} \subset \mathscr{B}(X)$ be a nested, G-invariant, G-finite subset consisting of tight elements with structure tree $T = T(\mathcal{E})$. Then, given $v \in V(T)$, we have that the vertex stabiliser $G_v = \mathrm{Stab}(v)$ acts freely and compactly on a subgraph $X_v \subset X$ such that:*

(1) *Each X_v has uniform coboundary in X, and so we may canonically identify $\Omega(X_v)$ with a subset of $\Omega(X)$.*
(2) *If ω_1, ω_2 lie in $\Omega(X_v)$ then there is no $b \in \mathcal{E}$ which separates ω_1 and ω_2.*
(3) *If $\mathcal{E}$ is taken to be $\mathcal{E}_n$ as in Theorem IV.3.7, then we also have that $\mathrm{es}(X_v) \geq n$.*

PROOF. Let $\mathcal{E}$ be a nested G-invariant, G-finite subset consisting of tight elements, and let $T = T(\mathcal{E})$ be the structure tree for this set. So $\mathcal{E}$ can be naturally identified with $\vec{E}(T)$ and each $v \in V(T)$ can be identified with those $e \in \mathcal{E}$ which point 'into' v. Fix $v \in V(T)$. Given $b \in v$ and $m \geq 1$, let $R(m, b)$ be the subgraph of X induced by the set of vertices which lie a distance of at most m from b^*. We choose m sufficiently large so that $R(m, b)$ satisfies the following:

(1) $R(m, b)$ contains all geodesics in $X[b]$ containing endpoints of edges in δb,
(2) If there exists n pairwise edge-disjoint paths $p_1, \ldots, p_n$ in $X[b]$ with endpoints in δb, then n pairwise edge-disjoint paths $p_1', \ldots, p_n'$ in $R(m, b)$ such that p_i' has the same endpoints as p_i.

Recall from § I.5.3 that $X[b]$ denotes the subgraph of X induced by b. Such an m clearly exists since G acts on $\mathcal{E}$ with finitely many orbits. We then define

$$X_v = \bigcap_{b \in v} X[b^*] \cup \bigcup_{b \in v} (\delta b \cup R(m, b)).$$

We assume that X_v is an induced subgraph. If not, then add back the missing edges. Since each $b \in \mathcal{E}$ is tight, we deduce that X_v is connected. If $\mathcal{E} = \mathcal{E}_n$, then fact that $\mathrm{es}(X_v) \geq n$ essentially follows directly from property (2) above, together with MENGER's theorem (IV.1.5). □

We conclude this subsection with some helpful, miscellaneous results about tight elements of $\mathscr{B}(X)$.

The following was first observed by DUNWOODY in [**45**] in the case of minimal cuts, and subsequently extended by THOMASSEN and WOESS in [**129**] to tight cuts of bounded size.

PROPOSITION IV.3.11 ([**129**, Prop. 4.1]). *Let X be a (possibly locally infinite) connected graph, $e \in E(X)$, $k \geq 1$. Then there exists only finitely many tight $b \in \mathscr{B}(X)$ such that δb contains e and $|\delta b| \leq k$.*

Note that Proposition IV.3.11 immediately implies that $\mathcal{E}_n/G$ is finite for all $n \geq 1$. Moreover, the following corollary is immediate.

Corollary IV.3.12. *Let X be a connected, locally finite, cocompact graph. Then for every $n > 0$ there exists $m > 0$ such that for all tight $b \in \mathscr{B}(X)$, if $|\delta b| < n$ then* $\operatorname{diam}(\delta b) < m$.

We also note the following tricks for creating tight cuts.

Proposition IV.3.13. *Let X be a connected graph and let $b_0 \in \mathscr{B}(X)$ such that $X[b_0]$ is connected. Let U be a connected component of $X \setminus X[b_0]$, and let b_1 denote the set of vertices of U. Then b_1 is a tight element of $\mathscr{B}(X)$.*

Proof. Clearly $\delta b_1 \subset \delta b_0$, so in particular δb_1 is finite and thus $b_1 \in \mathscr{B}(X)$. By construction, $X[b_1]$ is connected. We need only observe that $X[b_1^*]$ is connected, but this is clear since $X[b_0]$ is connected and every edge in δb_1 abuts b_0. □

Proposition IV.3.14. *Let X be a connected graph. Let $b \in \mathscr{B}(X)$, $e_1, e_2 \in \delta b$. Suppose there exists paths $p \subset X[b]$, $q \subset X[b^*]$, both connecting an endpoint of e_1 to an endpoint of e_2. Then there exists a tight element $b' \in \mathscr{B}(X)$ such that $\delta b' \subset \delta b$ and $e_1, e_2 \in \delta b'$.*

Proof. Let $b_0 \subset b$ be the vertex set of the connected component of $X[b]$ which contains p. Now, let b' be the vertex set of the connected component of $X[b_0]$ which contains q. By Proposition IV.3.13, this is a tight element, and certainly $\delta b' \subset \delta b$ and $e_1, e_2 \in \delta b'$. □

The following fact is also useful and worthy of mention. This was first observed by Möller [**101**], and a short proof can be found in [**129**].

Proposition IV.3.15 ([**129**, Prop. 7.1]). *Let X be a connected, locally finite graph. Let $\mathcal{E}$ be a subset of $\mathscr{B}(X)$, and let R be the subring of $\mathscr{B}(X)$ generated by $\mathcal{E}$. If $\omega_1, \omega_2 \in \Omega(X)$ are separated by some $b \in R$, then there is some $b' \in \mathcal{E}$ which separates them too.*

IV.3.3. Cuts and quasi-isometries. We now study how the above concepts can be pushed through quasi-isometries.

Lemma IV.3.16. *Let $\psi : \Gamma_1 \to \Gamma_2$ be a quasi-isometry between connected, locally finite graphs. Let $b \in \mathscr{B}(\Gamma_1)$. Then there exists $b' \in \mathscr{B}(\Gamma_2)$ such that* $\operatorname{Haus}_{\Gamma_2}(\psi(b), b')$ *is uniformly bounded, and $\delta b'$ is contained in a bounded neighbourhood of $\psi(\delta b)$. Moreover if $\Gamma_1[b]$ is connected then so is $\Gamma_2[b']$.*

Proof. Let $\varphi : \Gamma_2 \to \Gamma_1$ be some choice of quasi-inverse to ψ. Let $\lambda \geq 1$ be such that ψ and φ are (λ, λ)-quasi-isometries, and φ is a λ-quasi-inverse to ψ.

Let $b' = B_{\Gamma_2}(\psi(b); R) \cap V(\Gamma)_2$ for some large $R > 0$, say $R = 100\lambda^5$. Clearly if $\Gamma_1[b]$ is connected then so is $\Gamma_2[b']$. Suppose u and v are adjacent vertices in Γ_2 such that $u \in b'$ and $v \notin b'$. Since R is sufficiently large compared to the quasi-isometry constants, we see that $\varphi(v) \notin b$, but certainly both $\varphi(v)$ and $\varphi(u)$ lie in a bounded neighbourhood of b. Thus, u and v lie in a bounded neighbourhood of $\psi(\delta b)$. Since u and v are arbitrary and Γ_2 is locally finite, we have that $\delta b'$ is finite and so $b' \in \mathscr{B}(\Gamma_2)$. □

The following lemma is a coarse version of Corollary IV.3.12.

Lemma IV.3.17. *Let X, Γ be bounded valence, connected graphs. Let X be cocompact, and suppose X and Γ are quasi-isometric. Then for every $n > 0$ there exists $m > 0$ such that for all tight $b \in \mathscr{B}(\Gamma)$, if $|\delta b| < n$ then* $\operatorname{diam}(\delta b) < m$.

Proof. Let $\varphi : X \to \Gamma$ be a quasi-isometry with quasi-inverse $\psi : \Gamma \to X$. We assume without loss of generality that these are continuous maps. As usual, fix $\lambda \geq 1$ which is larger than all quasi-isometry constants involved. Fix $n > 0$, and let $(b_i)_{i\geq 1}$ be a sequence of tight cuts in $\mathscr{B}(\Gamma)$ such that $|\delta b_i| < n$ for all i, and $\operatorname{diam}(\delta b_i) \to \infty$. We will find boundedly small tight cuts of arbitrarily large diameter in X, contradicting Corollary IV.3.12.

Since each δb_i contains at most n edges, it is clear that we can choose decompositions

$$\delta b_i = C_i \sqcup D_i,$$

for each i such that the infimal distance $\mathrm{d}_\Gamma(C_i, D_i) \to \infty$ as $i \to \infty$. Each b_i is tight, so let $U_i = \Gamma[b_i]$, $W_i = \Gamma[b_i^*]$ be the two connected components of $\Gamma \setminus \delta b_i$. Let p_i be a path through U_i connecting C_i to D_i, and let q_i be a path through W_i connecting C_i to D_i.

For each $i \geq 1$, using Lemma IV.3.16 let $b_i' \in \mathscr{B}(X)$ be such that $\operatorname{Haus}(\psi(b_i), b_i')$ is uniformly bounded, and $\delta b_i'$ is contained in a bounded neighbourhood of $\psi(\delta b_i)$. Since each b_i was tight, we may assume that $X[b_i']$ is connected for every $i \geq 1$. As X is bounded valence, we deduce that there exists some uniform $N > 0$ such that each b_i' satisfies $|\delta b_i'| < N$. We are almost done, but each b_i' may not be tight. We will now find some tight $b_i'' \in \mathscr{B}(X)$ such that $\delta b_i'' \subset \delta b_i'$, and $\operatorname{diam}(\delta b_i'') \to \infty$ as $i \to \infty$. Once this is achieved, we are done.

Let $r > 0$ be such that for every $i \geq 1$, $\delta b_i'$ is contained in the r-neighbourhood of $\psi(\delta b_i)$. Assume without loss of generality that for all $i \geq 1$, we have that

$$\mathrm{d}_\Gamma(C_i, D_i) > 3\lambda(r + \lambda).$$

Let $R = \lambda(r + \lambda) + 1$ In particular, there must exist a subpath p_i' of p_i such that p_i' lies outside of the R-neighbourhood of δb_i, but p_i' begins in the $2R$-neighbourhood of C_i, and ends in the $2R$-neighbourhood of D_i. Let $q_i = \psi(p_i')$. By our choice of R, q_i is disjoint from the r-neighbourhood of $\psi(\delta b_i)$, and thus disjoint from $\delta b_i'$. It follows that q_i is contained a single connected component of $X \setminus \delta b_i'$. Note that the endpoints of q_i are uniformly close to $\delta b_i'$ by construction, but they become arbitrarily far part as $i \to \infty$.

Let U_i denote the connected component of $X \setminus X[b_i']$ containing q_i, and let b_i'' be the vertex set of U_i. By Proposition IV.3.13, b_i'' is tight. Clearly $\delta b_i'' \subset \delta b_i'$, and $\operatorname{diam}(\delta b_i'') \to \infty$ as $i \to \infty$. The lemma follows. □

Finally, we record the following pigeonhole argument, which will allow us to push the increased cut-size created in Theorem IV.3.10 through a quasi-isometry.

LEMMA IV.3.18. *Let Γ_1, Γ_2 be bounded valence connected graphs, and let $\varphi : \Gamma_1 \to \Gamma_2$ be a continuous quasi-isometry. Let $m > 0$ and fix a subgraph $\Lambda \subset \Gamma_1$ such that $\operatorname{vs}(\Lambda) \geq m$. Then*

$$\operatorname{vs}(\varphi(\Lambda)) \geq Cm,$$

where $C = C(\varphi) > 0$ is some constant depending only on φ and Γ_1.

PROOF. Fix $\lambda \geq 1$ such that φ is a (λ, λ)-quasi-isometry. To simplify notation, let $\Lambda = \Lambda_1$ and write $\varphi(\Lambda) = \Lambda_2$. Suppose that there are k vertices $v_1, \ldots, v_k$ in Λ_2 whose removal separates distinct ends $\omega_1, \omega_2 \in \Omega(\Lambda_2)$. Let $\xi_i = \varphi^{-1}(\omega_i)$ for each $i = 1, 2$, recalling that quasi-isometries induce well-defined bijections on the corresponding sets of ends. Since $\operatorname{vs}(\Lambda) \geq m$ and the Γ_i are locally finite, we have by MENGER's theorem (IV.1.5) that there exists m pairwise disjoint bi-infinite paths $\alpha_1, \ldots, \alpha_m$ in Λ between ξ_1 and ξ_2.

Thus each $\alpha_i' := \varphi(\alpha_i)$ is a bi-infinite path between the ends ω_1 and ω_2. Each α_i' must pass through some v_j, so by the pigeonhole principle there exists some v_j such that at least m/k of the α_i' pass through v_j. By relabelling, we can assume without loss of generality that $\alpha_1', \ldots, \alpha_{m/k}'$ pass through v_1. Let $x \in \varphi^{-1}(v_1)$. As φ is a quasi-isometry, we have that $\alpha_1, \ldots \alpha_{m/k}$ must intersect the closed r-neighbourhood of x, for some $r = r(\varphi) > 0$ depending only on φ. Combining this observation with the assumption that Γ_1 is bounded valence together with the fact that the α_i are disjoint, we deduce that

$$\frac{m}{k} \leq |B_{\Gamma_1}(x; r)|.$$

The right-hand side is bounded above by some uniform constant since Γ_1 is bounded valence. The lemma follows. □

We conclude this section by stating the following result, which is just a convenient repackaging of some of the above.

THEOREM IV.3.19. *Fix $N > 0$, Let G be a finitely generated group acting freely and cocompactly on a connected, locally finite graph X. Let Γ be a connected, bounded valence graph, and let $\varphi : X \to \Gamma$ be a continuous quasi-isometry. Then G splits as a graph of groups $G(\Theta) = (G_x, \alpha_e)$ such that the following hold:*

(1) *Each edge group G_e is finite.*

(2) *For each vertex group* G_v *acts freely and cocompactly on a subgraph* $X_v \subset X$ *with uniform coboundary.*
(3) *The quasi-isometry* φ *restricts to a quasi-isometry* $\varphi_v : X_v \twoheadrightarrow \Gamma_v$*, where* $\Gamma_v = \varphi(X_v)$*.*
(4) *For each* $v \in V(\Theta)$*, we have that either:*
   (a) X_v *has at most one end, or*
   (b) X_v *is multi-ended, and* $\mathrm{vs}(\Gamma_v) > N$*.*

IV.4. (Relative) cohomological planarity

IV.4.1. The CHomP property. In [**46**] it is shown that if a group G acts freely and cocompactly on a locally finite connected simplicial two-complex K such that $H^1(K; \mathbf{Z}_2) = 0$, then the 1-skeleton of K is an accessible graph. We need a certain weakening of this condition, and introduce the following terminology.

Definition IV.4.1 (Cohomological planarity). Let K be a connected polyhedral complex. We say that K is ***cohom****ologically* ***p****lanar* (or just ***CHomP***) if the natural map

$$H^1_c(K; \mathbf{Z}_2) \to H^1(K; \mathbf{Z}_2)$$

is the zero map.

In plain(er) English, a polyhedral complex K is CHomP if and only if every compactly supported 1-cocycle is a coboundary. It is easy to verify that a (possibly non-compact) surface without boundary is planar if and only if it is CHomP. This motivates our choice of nomenclature.

This condition was first considered in the context of accessibility by Groves–Swarup in [**69**] (see also [**127**]). If K is a connected surface without boundary, then this condition is equivalent to K being planar. This motivates our name for the property.

Proposition IV.4.2. *Let* K*,* L *be connected polyhedral complexes, and* $\psi : K \to L$ *induce a surjection between first homology groups with* $\mathbf{Z}_2$*-coefficients. If* K *is CHomP then so is* L*.*

Proof. Since $\mathbf{Z}_2$ is a field, it follows from the universal coefficients theorem that ψ induces an injection $H^1(L; \mathbf{Z}_2) \hookrightarrow H^1(K; \mathbf{Z}_2)$. We have the following commutative square:

$$\begin{array}{ccc} H^1_c(L;\mathbf{Z}_2) & \longrightarrow & H^1_c(K;\mathbf{Z}_2) \\ \big\downarrow & & \big\downarrow{\scriptstyle 0} \\ H^1(L;\mathbf{Z}_2) & \hookrightarrow & H^1(K;\mathbf{Z}_2) \end{array}$$

It follows immediately that L is CHomP. ☐

The CHomP property is very relevant to the question of accessibility. In particular, the following theorem is proven in [**39**, § VI].

Theorem IV.4.3 (Dicks–Dunwoody). *Let K be a connected, 2-dimensional polyhedral complex which is locally finite away from the 0-skeleton, and equipped with a cocompact action $G \curvearrowright K$ such that the stabiliser of every edge is finite. Suppose further that K is CHomP. Then $\mathscr{B}K$ admits a G-invariant, nested, G-finite generating set.*

Recall that a 2-dimensional polyhedral complex is said to be *locally finite away from the 0-skeleton* if every 1-cell is a face of at most finitely many 2-cells.

Remark IV.4.4. Some remarks are in order. Firstly, we note that in [**39**], they do not make the assumption that K is CHomP, but instead make the stronger assumption that $H_1(K; \mathbf{Z}_2) = 0$. Their arguments, however, go through unchanged with our weaker hypothesis. Indeed, the hypothesis $H_1(K; \mathbf{Z}_2) = 0$ is only used to ensure that 'tracks' on K separate. But in fact, the proof only requires that **compact** tracks separate. The CHomP property also implies this (and in fact is equivalent to all compact tracks separating).

Secondly, in [**39**] it is assumed that K is simplicial. This is easily weakened to K being a polygonal complex by passing to the second barycentric subdivision K'' of K, as we then have a natural equivariant surjection $\mathscr{B}(K'') \twoheadrightarrow \mathscr{B}(K)$, since K is locally finite away from the 0-skeleton.

Thirdly, we would like to point out that the assumption that edge stabilisers are finite is also unnecessary. Indeed, it was noted by Dunwoody[2] that this assumption can be dropped essentially by applying Proposition IV.3.11 in its place. This idea is realised in the appendix of [**89**]. In this chapter we do not need to make use of this extension, as our hypotheses will ensure that edge stabilisers are always trivial.

It will be helpful to define a 'coarse' variant of the CHomP property for graphs. For this, we introduce the following terminology.

Definition IV.4.5. Let X be a connected graph, and $\varepsilon \geq 0$. The *ε-filling of X*, denoted $K_\varepsilon(X)$, is the 2-dimensional polyhedral complex obtained from X by attaching a 2-cell along every closed loop of length at most ε.

It is easy to see that if X be a locally finite graph and $G \curvearrowright X$ cocompactly. Then for all $\varepsilon \geq 0$, this action extends to a cocompact action $G \curvearrowright K_\varepsilon(X)$. The following terminology is also helpful.

Definition IV.4.6 (Coarsely simply connected). Let Γ be a connected graph and $\varepsilon > 0$. We say that Γ is *ε-coarsely simply connected* if $K_\varepsilon(\Gamma)$ is simply connected. If there exists $\varepsilon > 0$ such that Γ is ε-coarsely simply connected, then we may just call Γ *coarsely simply connected*.

It is a classical fact that a finitely generated group is finitely presented if and only if its Cayley graphs are coarsely simply connected. We now introduce the following.

[2]Personal communication.

Definition IV.4.7 (Coarsely CHomP). Let X be a connected graph. We say that X is *coarsely CHomP* if there exists $\varepsilon > 0$ such that $K_\varepsilon(X)$ is CHomP.

Combining Theorem IV.4.3 with Theorem IV.3.8, the following corollary is immediate, which we state now for later reference.

Corollary IV.4.8. *Let X be a connected, locally finite graph equipped with a cocompact group action $G \curvearrowright X$. Suppose X is coarsely CHomP. Then X is accessible.*

Note that simply connected complexes are clearly CHomP. In particular, the above implies Dunwoody's theorem that (almost) finitely presented groups are accessible [**46**].

IV.4.2. The relatively CHomP property and $\mathcal{H}$-elliptic cuts. The CHomP property is still too strong for our needs. For example, even a planar surface with non-empty boundary need not be CHomP. To remedy this, we will 'cone off' certain bad regions of our complex, in order consider cuts 'relative' to these nasty bits. The 'bad' regions should be thought of as behaving a bit like the boundary components of a planar surface.

Definition IV.4.9 (Systems of subgraphs). Let X be a graph equipped with an action $G \curvearrowright X$. Then a *system of subgraphs* is a set $\mathcal{H}$ of (not necessarily connected) subgraphs of X. We introduce the following terminology.

- If for all $Y \in \mathcal{H}$, $g \in G$ we have that $gY \in \mathcal{H}$, we say that $\mathcal{H}$ is *G-invariant.* This induces an action $G \curvearrowright \mathcal{H}$.
- If at most finitely many $Y \in \mathcal{H}$ intersect any bounded subset of X, we say that $\mathcal{H}$ is *locally finite.*

If K is a polyhedral complex, then a system of subgraphs of K is taken to mean a system of subgraphs of the 1-skeleton K^1.

Remark IV.4.10. Note that if $\mathcal{H}$ is a locally finite system of subgraphs and X is a locally finite graph, then replacing each $Y \in \mathcal{H}$ with a bounded tubular neighbourhood of itself does not affect this property, nor does replacing each Y with a subgraph of itself.

Definition IV.4.11 ($\mathcal{H}$-elliptic cuts). Let X be a connected graph, and $\mathcal{H}$ be a collection of subgraphs of X. We say that $b \in \mathscr{B}(X)$ is *$\mathcal{H}$-elliptic* if for all $Y \in \mathcal{H}$, either $b \cap Y$ or $b^* \cap Y$ is finite. Let $\mathscr{B}_\mathcal{H}X$ denote the set of $\mathcal{H}$-elliptic cuts.

If K is a polyhedral complex and $\mathcal{H}$ is a system of subgraphs of K, then write $\mathscr{B}_\mathcal{H}K := \mathscr{B}_\mathcal{H}K^1$.

Proposition IV.4.12. *Let X be a connected graph and $\mathcal{H}$ a peripheral system. Then $\mathscr{B}_\mathcal{H}(X)$ is indeed a subring of $\mathscr{B}(X)$.*

Proof. Let $b_1, b_2 \in \mathscr{B}_{\mathcal{H}}(X)$. We need to check that $b_1 \cap b_2$ and $b_1 \triangle b_2$ are in $\mathscr{B}_{\mathcal{H}}(X)$ (recall that '$\triangle$' denotes symmetric difference). Fix $H \in \mathcal{H}$. We now split into three cases.

Suppose first that $b_1 \cap H$ and $b_2 \cap H$ are finite. Then $(b_1 \cap b_2) \cap H$ is certainly finite and so $b_1 \cap b_2 \in \mathscr{B}_{\mathcal{H}}(X)$. Also,

$$(b_1 \triangle b_2) \cap H = b_1 \cap H \triangle b_2 \cap H,$$

which is finite.

Secondly, assume that $b_1 \cap H$ and $b_2^* \cap H$ are finite. Then $(b_1 \cap b_2) \cap H \subset b_1 \cap H$ is finite, and

$$(b_1 \triangle b_2)^* \cap H = (H \cap b_1^* \cap b_2^*) \cup (H \cap b_1 \cap b_2).$$

which is also finite.

Finally, suppose that $b_1^* \cap H$ and $b_2^* \cap H$ are finite. Then

$$H \cap (b_1 \cap b_2)^* = H \cap (b_1^* \cup b_2^*) = (b_1^* \cap H) \cup (b_2^* \cap H),$$

which is finite. Furthermore, $b_1^* \triangle b_2^* = b_1 \triangle b_2$, and so

$$(b_1 \triangle b_2) \cap H = (b_1^* \triangle b_2^*) \cap H = (b_1^* \cap H) \triangle (b_2^* \cap H),$$

which is finite.

In each case, we have shown that both $b_1 \cap b_2$ and $b_1 \triangle b_2$ lie in $\mathscr{B}_{\mathcal{H}}(X)$. It follows that $\mathscr{B}_{\mathcal{H}}(X)$ is a subring of $\mathscr{B}(X)$. $\square$

We will almost exclusively be interested in peripheral systems which satisfy the following condition, which we call *tameness*.

Definition IV.4.13 (Tame system of subgraphs). Let X be a connected graph, then a peripheral system $\mathcal{H}$ is called *tame* if for any $\mathcal{H}$-cut $b \in \mathscr{B}_{\mathcal{H}}(X)$ we have that only finitely many $Y \in \mathcal{H}$ intersect both b and b^*.

It is tempting to say that a G-invariant, locally finite system of subgraphs of a cocompact complex is necessarily tame. Sadly, this is not the case, as the following example shows.

Example IV.4.14. Consider the 3-regular tree $X = T_3$ with a distinguished end $\omega_0 \in \Omega(X)$, thus viewing it as a binary tree with 'infinite level sets'. Let $G \subset \mathrm{Aut}(X)$ be the subgroup which stabilises ω_0, and thus preserves the level sets. It is easy to see that G acts on X (edge-)transitively. Let $\mathcal{H}$ denote the set of level sets. Clearly $\mathcal{H}$ is G-invariant and locally finite. Moreover, it is not too hard to see that every $b \in \mathscr{B}(X)$ is an $\mathcal{H}$-cut, and so $\mathscr{B}_{\mathcal{H}}(X) = \mathscr{B}(X)$. However, $\mathcal{H}$ is not tame, as removing any edge from X separates vertices in infinitely many level sets.

Mercifully, however, we can achieve tameness by placing one more sensible condition on $\mathcal{H}$. This will be helpful to us later on.

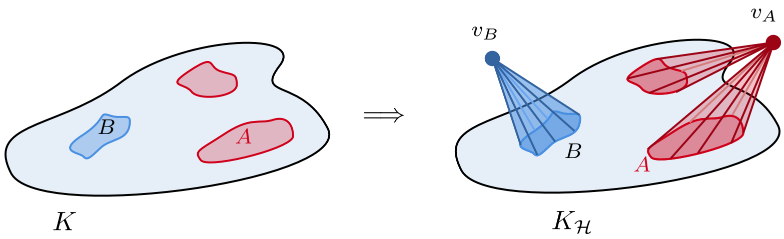

FIGURE IV.5. Constructing the coned-off complex $K_{\mathcal{H}}$. Here, $\mathcal{H} = \{A, B\}$.

PROPOSITION IV.4.15. *Let X be a connected graph. Let $\mathcal{H}$ be a locally finite system of subgraphs such that each connected component of every $Y \in \mathcal{H}$ has infinite 0-skeleton. Then $\mathcal{H}$ is tame.*

PROOF. Suppose $\mathcal{H}$ were not tame, so fix $b \in \mathscr{B}_{\mathcal{H}}(X)$ such that infinitely many $Y \in \mathcal{H}$ intersect both b and b^*. For any such Y, one of either b or b^* is finite. Since every component of Y is infinite, we must have that some edge in Y lies in δb. But then δb is finite, and so by the pigeonhole principle we see that some edge in δb is contained in infinitely many $Y \in \mathcal{H}$. This contradicts the fact that $\mathcal{H}$ is locally finite. $\square$

Tameness is important as it will allow us to reduce most questions about $\mathscr{B}_{\mathcal{H}}(K)$ to questions about $\mathscr{B}(K')$ for some other complex K', via the following construction.

DEFINITION IV.4.16 (Coned-off complex). Let K be a connected 2-dimensional polyhedral complex and $\mathcal{H}$ a system of subgraphs of K. We define the *coned-off complex*, denoted $K_{\mathcal{H}}$, as follows: For every $Y \in \mathcal{H}$, form the topological cone

$$CY = Y \times [0,1]/(x,1) \sim (y,1),$$

equipped with the obvious cell structure, with a single new vertex v_Y added, called the *cone vertex*. Attach CY to K by identifying the base of the cone $Y \times \{0\} \subset CY$ with $Y \subset K$.

See Figure IV.5 for a cartoon of this construction. Below we record a couple of easy observations.

PROPOSITION IV.4.17. *Let K be a connected, 2-dimensional, locally finite, polyhedral complex and locally finite system of subgraphs of K. Then we have the following:*

(1) *If $\mathcal{H}$ is locally finite then $K_{\mathcal{H}}$ is locally finite away from the 0-skeleton.*
(2) *If K is equipped with a G-action and $\mathcal{H}$ is G-invariant then this extends to an action on $K_{\mathcal{H}}$. If the action on K is cocompact and $\mathcal{H}$ is locally finite then the induced action upon K is also cocompact.*

The coned-off complex will, in general, be locally infinite, as the cone vertex v_A corresponding to $A \in \mathcal{H}$ will have degree exactly equal to the number of vertices in A. Its purpose is to provide a lens with which we may inspect the subring $\mathscr{B}_{\mathcal{H}}(K)$, via the following.

Proposition IV.4.18. *Let K be a connected 2-dimensional simplicial complex and $\mathcal{H}$ a system of subgraphs. Then the inclusion $K \hookrightarrow K_{\mathcal{H}}$ induces a ring homomorphism*

$$F : \mathscr{B}(K_{\mathcal{H}}) \to \mathscr{B}_{\mathcal{H}}(K).$$

If $\mathcal{H}$ is tame then F is surjective. If every $Y \in \mathcal{H}$ has infinite 0-skeleton then F is injective. Finally, if K is equipped with a G-action and $\mathcal{H}$ is G-invariant then F is G-equivariant.

Proof. Let $F : \mathscr{B}(K_{\mathcal{H}}) \to \mathscr{B}(K)$ be given by restriction. That is, $F(b) = b \cap V(\Gamma)$. It is routine to check that this is a (G-equivariant) homomorphism of Boolean rings.

We claim that the image of F is contained in $\mathscr{B}_{\mathcal{H}}(K)$. Firstly, for the sake of a contradiction that there exists some $b \in \mathscr{B}(K_{\mathcal{H}})$ and some $Y \in \mathcal{H}$ such that both $Y \cap F(b)$ and $Y \cap F(b^*)$ are infinite. Now, either $v_Y \in b$ or $v_Y \notin b$. In either case we see that δb is infinite, a contradiction. The claim follows.

Suppose now that $\mathcal{H}$ is tame. We claim that F is surjective. Fix $b \in \mathscr{B}_{\mathcal{H}}(K)$. We extend b to an element $b' \in \mathscr{B}(K_{\mathcal{H}})$ by including a cone vertex v_Y if and only if $b \cap Y$ is infinite. Note that $\delta b'$ is therefore obtained by adding to δb a collection of finitely many edges for each $Y \in \mathcal{H}$ such that both $b \cap Y$ and $b^* \cap Y$ are non-empty. It then follows from the tameness of $\mathcal{H}$ that $\delta b'$ is also finite. Clearly $F(b') = b$, and so F is surjective.

Finally, suppose every $Y \in \mathcal{H}$ has infinite 0-skeleton. We need only show that F has trivial kernel. Here, the 0-element is the empty set. Clearly any $b \in \mathscr{B}(K_{\mathcal{H}})$ which satisfies $F(b) = \emptyset$ cannot contain any cone vertices lest b must have infinite coboundary. Thus, any element b such that $F(b)$ is empty must itself be empty. $\square$

Remark IV.4.19. Note that we lose nothing by assuming that every $Y \in \mathcal{H}$ has infinite 0-skeleton. It is easy to check that if $\mathcal{H}' \subset \mathcal{H}$ is such that $\mathcal{H} \setminus \mathcal{H}'$ contains only finite subgraphs, then $\mathscr{B}_{\mathcal{H}}(K) = \mathscr{B}_{\mathcal{H}'}(K)$.

We now have the following key definition.

Definition IV.4.20 (Relatively CHomP). Let K be a connected 2-dimensional polyhedral complex and $\mathcal{H}$ a system of subgraphs of K^1. We say that K is *CHomP relative to* $\mathcal{H}$ if $K_{\mathcal{H}}$ is CHomP.

Similarly, if X is a connected graph and $\mathcal{H}$ is a system of subgraphs, then we say that X is *coarsely CHomP relative to* $\mathcal{H}$ if there exists $\varepsilon > 0$ such that $K_\varepsilon(X)$ is CHomP relative to $\mathcal{H}$.

Combining Propositions IV.4.17 and IV.4.18 with Theorem IV.4.3, we deduce the following.

Theorem IV.4.21 (Accessibility over $\mathcal{H}$-elliptic cuts). *Let K a connected, locally finite, 2-dimensional polyhedral complex, equipped with a proper, cocompact action by a group G. Let $\mathcal{H}$ be a tame, G-invariant, locally finite system of infinite subgraphs such that K is CHomP relative to $\mathcal{H}$. Then there exists a G-invariant, nested, G-finite generating set of $\mathscr{B}_{\mathcal{H}}(K)$.*

The following corollary is immediate.

Corollary IV.4.22. *Let X be a connected, locally finite graph equipped with a free, cocompact action $G \curvearrowright X$. Let $\mathcal{H}$ be a tame, G-invariant, locally finite system of infinite subgraphs such that X is coarsely CHomP relative to $\mathcal{H}$. Then there exists a G-invariant, nested, G-finite generating set of $\mathscr{B}_{\mathcal{H}}(X)$.*

Proof. Since X is coarsely CHomP relative to $\mathcal{H}$, then there exists $\varepsilon > 0$ such that $K = K_\varepsilon(X)$ is CHomP relative to $\mathcal{H}$. Since X is locally finite and $\mathcal{H}$ is a locally finite system of subgraphs, we have that $K_{\mathcal{H}}$ is locally finite away from the 0-skeleton. Thus, we have by Theorem IV.4.21 that $\mathscr{B}(K_{\mathcal{H}})$ has a G-invariant, nested, G-finite generating set. Applying Proposition IV.4.18, we obtain such a generating set for $\mathscr{B}_{\mathcal{H}}(X)$. Note finally that the image of a nested set under this isomorphism is clearly nested, so we are done. $\square$

There is a small subtlety which arises here. Note that the generating set of $\mathscr{B}(K_{\mathcal{H}})$ obtained via Theorem IV.4.3 consists of *tight* elements, but when we pass this generating set through the natural isomorphism given by Proposition IV.4.18, the corresponding generating set of $\mathscr{B}_{\mathcal{H}}(X)$ may contain elements which are not tight. Luckily this isn't a problem, thanks to the following trick.

Proposition IV.4.23. *Let X be a connected, locally finite graph equipped with a group action $G \curvearrowright X$. Suppose that $R \subseteq \mathscr{B}(X)$ is a G-invariant subring with a G-invariant, nested, G-finite generating set. Then there exists another G-invariant subring $S \subseteq \mathscr{B}(X)$ which contains R, such that S admits a G-invariant, nested, G-finite generating set consisting of tight elements.*

Proof. First, note that any $b \in \mathscr{B}(X)$ can be expressed as a disjoint union $b = c_1 \sqcup \ldots \sqcup c_k$, where each $c_i \in \mathscr{B}(X)$ is such that $X[c_i]$ is connected. Now, by an application of Proposition IV.3.13, every c_i is equal to the complement of a disjoint union of finitely many tight elements of $\mathscr{B}(X)$. It follows from this observation that every $b \in \mathscr{B}(X)$ lies in some subring of $\mathscr{B}(X)$ generated by a nested set of finitely many tight elements, each of which is such that their coboundary is a subset of δb.

Let $a_1, \ldots, a_n$ be orbit representatives of generating set of R. By the previous paragraph, for every $i = 1, \ldots, n$ let $A_i \subset \mathscr{B}(X)$ be a finite nested set, where each $c \in A_i$ is a tight element such that $\delta c \subset \delta a_i$, and a_i lies in the subring generated by A_i. Since set of translates of the a_i is a nested set, and every $c \in A_i$ satisfies $\delta c \subset \delta a_i$, we see that the G-orbit of the A_i is also nested. We may now take S to be the subring

$$S = \langle g \cdot c : g \in G,\ i \in \{1, \ldots, n\},\ c \in A_i \rangle.$$

This generating set is G-invariant, nested, G-finite, and consists of tight elements, and S clearly contains R. This proves the proposition. $\square$

REMARK IV.4.24. For the deformation space aficionados, Theorem IV.4.21 should be compared with [**71**, Prop. 2.24] of the GUIRARDEL–LEVITT paper on JSJ decompositions, which is a relative version of Dunwoody's accessibility theorem for (relatively) finitely presented groups. This result states that if a given group is finitely presented relative to a collection of peripheral subgroups, then the relative JSJ deformation space exists over any class of subgroups $\mathscr{A}$. We briefly describe how this result relates to our own.

Set $\mathscr{A}$ to the class of finite groups, and let X be a Cayley graph of a finitely generated group G. Let $\mathscr{H}$ be a set of peripheral subgroups, say finitely generated for the sake of simplicity. Take $\mathcal{H}$ to be the set of some connected tubular neighbourhoods of the cosets of the elements of $\mathscr{H}$, viewed as subgraphs of X. Then it is not too hard to see that $\mathscr{B}_{\mathcal{H}}X$ is a finitely generated G-module if and only if the aforementioned relative JSJ deformation space exists. Indeed, if $\mathcal{E}$ is a G-finite nested generating set of $\mathscr{B}_{\mathcal{H}}X$, then the structure tree of $\mathcal{E}$ is seen to be $(\mathscr{A},\mathscr{H})$-universally elliptic and so the relative JSJ deformation space over $\mathscr{A}$ exists. Conversely, if $\mathscr{B}_{\mathcal{H}}X$ is not finitely generated as a G-module then we can find reduced $(\mathscr{A},\mathscr{H})$-trees with arbitrarily large quotients, which will contradict relative accessibility in the sense of [**71**, Prop. 2.24].

It is also worth noting that if G is finitely presented relative to $\mathscr{H}$ then this implies that the Cayley graph of G is coarsely CHomP relative to $\mathcal{H}$. The converse is, of course, false.

We conclude this section by recording the following lemma relating to maps between coned-off complexes.

LEMMA IV.4.25. *Let K, L be locally finite, 2-dimensional, connected polyhedral complexes, and let $\mathcal{H}_1$, $\mathcal{H}_2$ be a system of subgraphs of K, L respectively. Let $\psi : K \to L$ be a continuous map. Let $f : \mathcal{H}_1 \to \mathcal{H}_2$ be a bijection such that $\psi(Y) \subset f(Y)$ for all $Y \in \mathcal{H}_1$. Then the map ψ extends to a map $\Psi : K_{\mathcal{H}_1} \to L_{\mathcal{H}_2}$. If ψ is a proper map*[3] *then so is Ψ.*

Suppose further that for any $Y \in \mathcal{H}_1$, the restriction $\psi|_Y : Y \to f(Y)$ induces a surjection of connected components. That is, for every $Y \in \mathcal{H}_2$ and every connected component $U \subset Y$, we have that $U \cap \psi(f^{-1}(Y))$ is non-empty. Then, if ψ induces a surjection between first homology groups (with any coefficients) then so does Ψ.

PROOF. The map $\psi|_Y \times \mathrm{id} : Y \times [0,1] \to f(Y) \times [0,1]$ descends to a map $\psi_Y : CY \to C[f(Y)]$. If ψ is proper, then ψ_Y is also clearly proper. We now define Ψ via the formula

$$\Psi(x) = \begin{cases} \psi(x) & \text{if } x \in K, \\ \psi_Y(x) & \text{if } x \in CY \setminus K, \text{ where } Y \in \mathcal{H}_1. \end{cases}$$

This is clearly continuous, and proper if ψ is proper.

Now, let us assume that for any $Y \in \mathcal{H}_1$, the restriction $\psi|_Y : Y \to f(Y)$ induces a surjection of connected components, and that ψ induces a surjection between the first homology groups. We

[3] Recall that a continuous map between topological spaces is said to be *proper* if the inverse image of any compact set is compact.

need to check that Ψ does too. In this proof we write $H_1(X)$ for the sake of notational simplicity. The argument goes through verbatim with any coefficients.

We will prove this using a Mayer–Vietoris sequence. Let $A = K_{\mathcal{H}_1} \setminus K$, i.e. A is the union of the interiors of the cones attached to K. Then a small open neighbourhood U of $K_{\mathcal{H}_1} \setminus A$ is homotopy equivalent to K. The intersection $U \cap A$ is homotopy equivalent to the disjoint union $\bigsqcup_{Y \in \mathcal{H}_1} Y$. Similarly, let $B = L_{\mathcal{H}_2} \setminus L$, and let W be a small open neighbourhood of L in $L_{\mathcal{H}_2}$. As before, W is homotopy equivalent to L and $W \cap B$ is homotopy equivalent to $\bigsqcup_{Z \in \mathcal{H}_2} Z$. Note that $\Psi(A) \subset B$, and we may clearly choose U, W such that $\Psi(U) \subset W$. Constructing the Mayer–Vietoris sequence for these decompositions, we arrive at the following diagram with exact rows.

$$
\begin{tikzcd}
\cdots \arrow[r] & H_1(A) \oplus H_1(K) \arrow[r] \arrow[d, twoheadrightarrow] & H_1(K_{\mathcal{H}_1}) \arrow[r] \arrow[d] & H_0(\bigsqcup_{\mathcal{H}_1} Y) \arrow[r] \arrow[d, twoheadrightarrow] & H_0(A) \oplus H_0(K) \arrow[r] \arrow[d, "\cong"] & \cdots \\
\cdots \arrow[r] & H_1(B) \oplus H_1(L) \arrow[r] & H_1(L_{\mathcal{H}_2}) \arrow[r] & H_0(\bigsqcup_{\mathcal{H}_2} Z) \arrow[r] & H_0(B) \oplus H_0(L) \arrow[r] & \cdots
\end{tikzcd}
$$

The first arrow from the left is onto, since ψ induces a surjection on first homology by assumption and also both A, B are contractible. The third arrow is a surjection by our hypotheses. The fourth arrow is an isomorphism since K and L are connected and the bijection $a : \mathcal{H}_1 \to \mathcal{H}_2$ induces a bijection between the connected components of A and B. It follows from the four lemma that the second vertical arrow from the left is a surjection, which is precisely what we set out to show. $\square$

IV.4.3. Quasi-isometries and the CHomP property. We now discuss how all of the above discussion can be pushed through quasi-isometries. Throughout this subsection, we fix the following notation.

(1) Let Γ_1, Γ_2 be connected, locally finite graphs, and for $i = 1, 2$ let $\mathcal{H}_i$ be a locally finite system of subgraphs of Γ_i.
(2) Let $\lambda \geq 1$. Let $\psi : \Gamma_1 \to \Gamma_2$ be a (λ, λ)-quasi-isometry between connected, locally finite graphs, with λ-quasi-inverse $\varphi : \Gamma_2 \to \Gamma_1$.
(3) Let $a : \mathcal{H}_1 \to \mathcal{H}_2$ be a bijection such that $\psi(Y) \subset a(Y)$ for all $Y \in \mathcal{H}_1$.
(4) Assume further that the restriction $\psi|_Y : Y \to a(Y)$ induces a surjection of connected components. That is, for every $Y \in \mathcal{H}_2$ and every connected component $U \subset Y$, we have that $U \cap \psi(a^{-1}(Y))$ is non-empty.
(5) Finally, assume there is $r \geq 0$ such that $\mathrm{Haus}(a(Y), \psi(Y)) \leq r$ for all $Y \in \mathcal{H}_1$

We have the following small result, which enables us to push $\mathcal{H}$-elliptic cuts through quasi-isometries.

Proposition IV.4.26. *Let $\omega_1, \omega_2 \in \Omega(\Gamma_1)$ be distinct. If there exists $b_1 \in \mathscr{B}_{\mathcal{H}_1}(\Gamma_1)$ which separates ω_1 and ω_2, then there also exists $b_2 \in \mathscr{B}_{\mathcal{H}_2}(\Gamma_2)$ which separates $\psi(\omega_1)$ and $\psi(\omega_2)$.*

Proof. Let $b_2 = B_{\Gamma_2}(\psi(b_1); R)$ for some large R, say $R = 100\lambda^5$. Suppose u and v are adjacent vertices in Γ_2 such that $u \in b_2$ and $v \notin b_2$. Since R is sufficiently large compared to the quasi-isometry constants, we see that $\varphi(v) \notin b_1$, but certainly both $\varphi(v)$ and $\varphi(u)$ lie in a bounded neighbourhood of b_1. Thus, u and v lie in a bounded neighbourhood of $\psi(\delta b_1)$. Since u and v are arbitrary and Γ_2 is locally finite, we have that δb_2 is finite and $b_2 \in \mathscr{B}(\Gamma_2)$. It then follows quickly from applying (5) above, combined with the facts that our graphs are locally finite and $\mathrm{Haus}(b_2, \psi(b_1)) < R < \infty$, that actually $b_2 \in \mathscr{B}_{\mathcal{H}_2}(\Gamma_2)$. $\square$

We now show that if Γ_1 is coarsely CHomP relative to $\mathcal{H}_1$ then Γ_2 is coarsely CHomP relative to $\mathcal{H}_2$.

Lemma IV.4.27. *For all sufficiently large $\varepsilon > 0$ the following holds. For every closed combinatorial loop ℓ in Γ_1, we have that ℓ and $\varphi \circ \psi \circ \ell$ are freely homotopic in $K_\varepsilon(\Gamma_1)$,*

Proof. Triangulate S^1 with n vertices $v_1, \dots, v_n$ such that $\ell(z)$ is a vertex if and only if z is a vertex. Let e_i be the 1-cell joining v_i to v_{i+1}, where indices are taken modulo n. To ease notation, let $\ell' = \varphi \circ \psi \circ \ell$.

We have that $\ell'(e_i)$ is a path of bounded length, and that $\mathrm{d}_\Gamma(\ell(v_i), \ell'(v_i))$ is also bounded. Let γ_i be a geodesic joining $\ell(v_i)$ to $\ell'(v_i)$. For each i we have that the concatenation

$$\alpha_i = \ell(e_i)\gamma_{i+1}\ell'(e_i)^{-1}\gamma_i^{-1}$$

is a loop of bounded length, say at most ε. Thus each α_i is null-homotopic in $K = K_\varepsilon(\Gamma_1)$. It is immediate from this observation that the map $S^1 \sqcup S^1 \to K$ sending each circle respectively to ℓ and ℓ' extends to a map $A \to K$, where A is the annulus and $S^1 \sqcup S^1$ is identified with ∂A. It follows that ℓ and ℓ' are freely homotopic. $\square$

In other words, the above lemma tells us that if $\varepsilon > 0$ is taken to be sufficiently large then the map $\varphi \circ \psi$ induces the identity map on the set of free homotopy classes in $K_\varepsilon(\Gamma_1)$. In [**19**, Prop. I.8.24] it is shown that finite presentability is a quasi-isometry invariant amongst finitely generated groups. The actual meat of the proof is devoted to proving the following.

Proposition IV.4.28. *For every $\varepsilon > 0$ there exists $\eta > 0$ such that the following holds. Let ℓ be a closed combinatorial loop, and suppose that $\varphi \circ \ell$ is null-homotopic in $K_\varepsilon(\Gamma_1)$. Then ℓ is null-homotopic in $K_\eta(\Gamma_2)$. In particular, if Γ_2 is coarsely simply connected then so is Γ_1.*

Applying this, we can prove the following.

Proposition IV.4.29. *For every sufficiently large $\varepsilon > 0$ there exists $\eta > 0$ such that $\psi : \Gamma_1 \to \Gamma_2$ extends to a continuous map $\hat{\psi} : K_\varepsilon(\Gamma_1) \to K_\eta(\Gamma_2)$. Moreover, we may assume that the induced map*

$$\hat{\psi}_* : H_1(K_\varepsilon(\Gamma_1); \mathbf{Z}_2) \to H_1(K_\eta(\Gamma_2); \mathbf{Z}_2)$$

on the first homology is a surjection.

Proof. Fix $\varepsilon > 0$ sufficiently large so that the hypothesis of Lemma IV.4.27 is satisfied, and choose $\eta > 0$ as in Proposition IV.4.28. Let ℓ be a null-homotopic loop in $K_1 := K_\varepsilon(\Gamma_1)$. We need only show that $\ell' = \psi \circ \ell$ is null-homotopic in $K_2 := K_\eta(\Gamma_2)$. Let $\ell'' = \varphi \circ \ell'$. Then ℓ'' and ℓ are freely homotopic in K_1, and so ℓ'' is null-homotopic. Applying Proposition IV.4.28 we deduce that ℓ' is null-homotopic. We now extend ψ to a map $\hat{\psi} : K_\varepsilon(\Gamma_1) \to K_\eta(\Gamma_2)$ by mapping every 2-cell σ with boundary loop ℓ to a homotopy disc bounded by $\psi \circ \ell$.

We now argue that the induced map on the first homology can be taken to be surjective. We have that the following diagram commutes:

$$\begin{tikzcd} H_1(\Gamma_1; \mathbf{Z}_2) \arrow[r, "\psi_*"] \arrow[d, two heads, "p"'] & H_1(\Gamma_2; \mathbf{Z}_2) \arrow[d, two heads, "q"] \\ H_1(K_1; \mathbf{Z}_2) \arrow[r, "\hat{\psi}_*"'] & H_1(K_2; \mathbf{Z}_2) \end{tikzcd}$$

Here, p, q are the (necessarily surjective) maps on homology induced by obvious inclusions. Recall that φ is continuous, and so induces a map φ_* on homology between Γ_2 and Γ_1. Lemma IV.4.27 implies that if z is a cycle in Γ_2, then

$$q(z + \psi_* \circ \varphi_*(z)) = 0 \implies q(z) = \hat{\psi}_* \circ p \circ \varphi_*(z).$$

Since q is surjective, we are done. $\square$

We can now prove the main theorem of this section.

Theorem IV.4.30. *If Γ_1 is coarsely CHomP relative to $\mathcal{H}_1$ then Γ_2 is coarsely CHomP relative to $\mathcal{H}_2$.*

Proof. Let $\varepsilon > 0$ be sufficiently large, and choose $\eta > 0$ as in Proposition IV.4.29. To ease notation, write $K = K_\varepsilon(\Gamma_1)$, $L = K_\eta(\Gamma_2)$. We thus have a continuous map $\hat{\psi} : K \to L$ which restricts to ψ on the 1-skeleton, and induces a surjection on first homology. Since ψ was a quasi-isometry of locally finite graphs, it is clear that $\hat{\psi}$ is a proper map. By Lemma IV.4.25, we have that $\hat{\psi}$ induces a map $\Psi : K_{\mathcal{H}_1} \to L_{\mathcal{H}_2}$, and Ψ induces a surjection between first homology groups. By Proposition IV.4.2 we see that $L_{\mathcal{H}_2}$ is CHomP, and thus Γ_2 is coarsely CHomP relative to $\mathcal{H}_2$. $\square$

IV.4.4. CHomP-ness of planar graphs. We now wish to record the fact that locally finite, 2-connected planar graphs with 'small faces' are coarsely CHomP relative to their infinite faces. The main result of this section is Theorem IV.4.35.

Let K be a locally finite, 2-connected, planar, polyhedral complex. Then the 1-skeleton K^1 is a locally finite planar graph and thus admits a good drawing, in the sense of Definition IV.1.21. With this in mind, an embedding $K \hookrightarrow \mathbf{S}^2$ is said to be a *good embedding* if it restricts to a good drawing of the 1-skeleton.[4]

[4]It is not immediately clear whether every such complex has a good embedding, but we will not need this. Every complex we consider will automatically have this property.

If we assume that K has a fixed good embedding, we may now define the set of *facial subgraphs* of K as

$$\mathcal{F}(K) = \{f \in \mathcal{F}(K^1) : f \text{ is not the boundary of a 2-cell in } K\}.$$

As we did with planar graphs, we define $\mathcal{F}^\infty(K)$ and $\mathcal{F}^{\mathrm{c}}(K)$ as the subsets of $\mathcal{F}(K)$ containing precisely the infinite and finite facial subgraphs, respectively.

Our immediate goal is to show that any K as above with a good embedding is CHomP relative to $\mathcal{F}(K)$. We first note the following very easy fact, the proof of which is an exercise in point-set topology.

LEMMA IV.4.31. *Let X be a topological space, and $A \subset X$ a subspace. Then the embedding $A \hookrightarrow X$ extends to a continuous injection of cones $CA \hookrightarrow CX$.*[5]

We may apply this to prove the following.

LEMMA IV.4.32. *Let K be a locally finite, planar, 2-connected polyhedral complex equipped with a good embedding $K \hookrightarrow \mathbf{S}^2$. Then the coned-off complex $K_{\mathcal{F}}$, where $\mathcal{F} = \mathcal{F}(K)$, admits a continuous injection into the 2-sphere $\mathbf{S}^2$.*

PROOF. Let $\vartheta : K \hookrightarrow \mathbf{S}^2$ denote the given good embedding of K. By Proposition IV.1.27, we see that for every $f \in \mathcal{F}(K)$ there exists an open disk D_f in the complement of $\mathbf{S}^2 \setminus \vartheta(K)$, such that ∂D_f is a simple closed curve and $\vartheta(f)$ is contained in ∂D_f. We may assume that $f \neq f'$ implies D_f and $D_{f'}$ are disjoint. Note that the cone of the circle is precisely the closed disk. By Lemma IV.4.31, there is a continuous injection from the cone Cf into D_f. Repeating this for every $f \in \mathcal{F}K$, we may extend the embedding $K \hookrightarrow \mathbf{S}^2$ to a continuous injection of $K_{\mathcal{F}}$ into $\mathbf{S}^2$, where $\mathcal{F} = \mathcal{F}K$. □

Next, we have the following technical lemma.

LEMMA IV.4.33. *Let K be a connected 2-dimensional polyhedral complex. Suppose K admits a continuous injection into $\mathbf{S}^2$, and every 1-cell is a face of exactly two 2-cells. Then K is CHomP.*

PROOF. Assume without loss of generality that K is simplicial, by passing to the second barycentric subdivision. Let $c : E(K) \to \mathbf{Z}_2$ be a compactly supported cocycle. For every simplex σ, we have that either exactly 0 or 2 its faces lie in the support of c. If two of the faces of σ do indeed lie in this support, say $e_1, e_2 \subset \sigma$, connect the midpoints of the e_i together with a straight-line segment through σ. Repeat this for every such σ. The resulting figure will be a union of finitely many simple closed curves $C = \ell_1 \cup \ldots \cup \ell_k$ in K, disjoint from the 0-skeleton and contained in only finitely many 2-simplices. Consider the image $C' = \vartheta(C)$. Since C is compact, $\vartheta|_C$ is a topological embedding. In particular, C' is a disjoint union of k simple closed curves in $\mathbf{S}^2$, and each of these curves is separating by the Jordan curve theorem.. Consider $\mathbf{S}^2 \setminus C'$, and 2-colour the components of this black

[5]In general, this will not be an embedding. The canonical example is that $C[0,1)$ does not embed into $C[0,1]$.

and white so that two components whose closures intersect are distinct colours. Note that every component of $\mathbf{S}^2 \setminus C'$ intersects $\vartheta(V(K))$ non-trivially. We now define a 0-cochain $a : V(K) \to \mathbf{Z}_2$, where $a(v) = 1$ if and only if $\vartheta(v)$ lies in a white component of $\mathbf{S}^2 \setminus C'$. We have that $c = \delta a$. In particular, every compactly supported 1-cocycle is a coboundary, and so K is CHomP. □

Piecing the above together, we deduce the following.

LEMMA IV.4.34. *Let K be a locally finite, 2-connected, planar polyhedral complex equipped with a good embedding. Then K is CHomP relative to $\mathcal{F}(K)$.*

We can now prove the following.

THEOREM IV.4.35. *Let Γ be a locally finite, 2-connected planar graph with a fixed drawing. Suppose that there exists $n > 0$ such that every $f \in \mathcal{F}^c(\Gamma)$ is a cycle of length at most n. Then Γ is coarsely CHomP relative to $\mathcal{F}^\infty(\Gamma)$.*

PROOF. To ease notation, write $\mathcal{F} = \mathcal{F}^\infty(\Gamma)$. Fix $\varepsilon > 0$ such that every finite facial subgraph of Γ is a cycle of length at most ε. Let $K = K_\varepsilon(\Gamma)$, so every $f \in \mathcal{F}^c(\Gamma)$ is null-homotopic in K. Next, form a second complex L by attaching a 2-cell along every $f \in \mathcal{F}^c(\Gamma)$. Note that since Γ is 2-connected, every such face is a simple closed curve, and moreover we see that the resulting complex L is itself planar, and satisfies $\mathcal{F}(L) = \mathcal{F}^\infty(\Gamma) = \mathcal{F}$. Finally, it is clear by construction that L the natural embedding of L into $\mathbf{S}^2$ is a good embedding. By Lemma IV.4.34, we deduce that L is CHomP relative to $\mathcal{F}$.

Next, we deduce that K is also CHomP relative to $\mathcal{F}$. As every finite face is null-homotopic in K, the identity map $\Gamma \to \Gamma$ extends to a continuous map $\psi : L \to K$. As K and L are locally finite 2-complexes, it is easy to see that this map can be taken to be proper. Note that since ψ restricts to an homeomorphism between 1-skeletons, it induces a surjection between first homology groups. By Lemma IV.4.25, ψ induces a proper map $\Psi : L_{\mathcal{F}} \to K_{\mathcal{F}}$ between coned-off complexes, which itself induces a surjection between first homology groups. By Proposition IV.4.2, we have that $K_{\mathcal{F}}$ is CHomP, as required. □

IV.5. Friendly-faced subgraphs of planar graphs

We now give some consideration to the relationship between the faces of a planar graph and the faces of its subgraphs. Now would be a good time for the reader to review our notation and terminology pertaining to the faces of a planar graph (see Definition IV.1.24).

IV.5.1. Discussion. The faces of a subgraph of a planar graph can be wildly different from the faces of the super-graph. This has the potential to cause issues for us further down the line, so the goal of this section is to try and claw back some control here. With this in mind, we have the following definition.

Definition IV.5.1 (Friendly-faced subgraphs). Let Γ be a connected, locally finite planar graph with fixed good drawing $\vartheta : |\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$. We say that a connected subgraph $\Lambda \subset \Gamma$ is *friendly-faced* if the following are satisfied:

(1) The good drawing ϑ of Γ restricts to a good drawing of Λ.
(2) There exists $r > 0$ such that for every $U_1 \in \mathcal{D}(\Lambda)$ (with the aforementioned induced good drawing) there exists $U_2 \in \mathcal{D}(\Gamma)$ such that $U_2 \subset U_1$ and
$$f_1 \subset B_\Gamma(f_2; r),$$
where $f_i = \mathcal{F}[U_i]$ for $i = 1, 2$.

Intuitively, a subgraph is friendly-faced if we can approximate its own facial subgraphs with those of the super-graph. Condition (1) is just to ensure that the induced drawing of the subgraph is sensible enough for us to work with, and is equivalent to asking that the inclusion $\Lambda \hookrightarrow \Gamma$ induces an inclusion $\Omega(\Lambda) \hookrightarrow \Omega(\Gamma)$ on the sets of ends, and thus an embedding $|\Lambda|_{\mathrm{Fr}} \hookrightarrow |\Gamma|_{\mathrm{Fr}}$ of Freudenthal compactifications.

Example IV.5.2. We illustrate this definition with some examples. Consider the planar graph Γ_1 and highlighted subgraph Λ_1 sketched in Figure IV.6. We have that Λ_1 is a friendly-faced subgraph of Γ_1, as for every facial subgraph of Λ_1 one can find a facial subgraph of Γ_1 such that the first is contained in a bounded neighbourhood of the latter. In particular, every facial subgraph of Λ_1, except for one, is also a facial subgraph of Γ_1. The 'inner' face is the exception, but a subgraph of the infinite 'inner' face of Γ_1 can serve as a good approximation.

In the second example in Figure IV.7, we have that Λ_2 is not a friendly-faced subgraph of Γ_2. In particular, the infinite 'inner' facial subgraph of Λ_2 is not contained in a bounded neighbourhood of any facial subgraph of Γ_2. Note that Λ_2 is obtained from Γ_2 by deleting a single edge. In this sense, being 'friendly-faced' is an easy property to break.

Another non-example is given in Figure IV.2. Viewing the graph depicted there as a subgraph of the two-dimensional grid in the obvious way, we have that it is not a friendly-faced subgraph of the grid.

An easy but important example of a friendly-faced subgraph is the *2-connected core* of an *almost 2-connected graph* (recall Definition IV.1.16).

Proposition IV.5.3. *Let Γ be a locally finite, almost 2-connected planar graph with fixed good drawing. Let $\Gamma_0 \subset \Gamma$ be its 2-connected core. Then Γ_0 is a friendly-faced subgraph of Γ.*

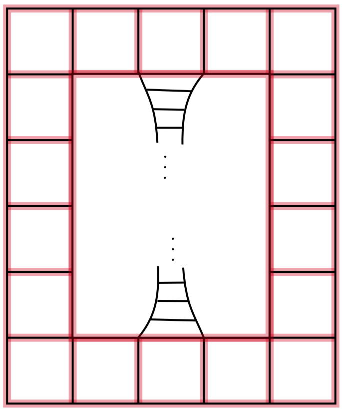

FIGURE IV.6. The subgraph Λ_1 is highlighted in red.

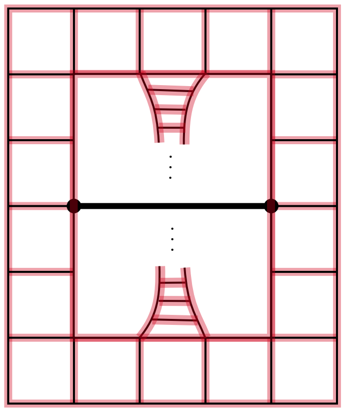

FIGURE IV.7. The subgraph Λ_2 is highlighted in red.

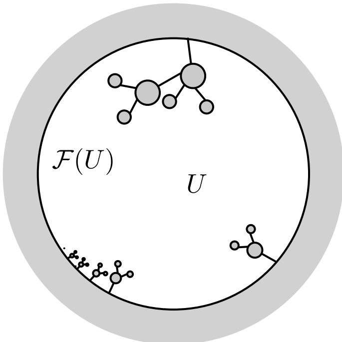

FIGURE IV.8. The simple body and adjoined cacti of a facial subgraph of an almost 2-connected planar graph. Each adjoined cactus has bounded diameter.

PROOF. Left as an exercise to the reader. □

REMARK IV.5.4. It is possible to give quite a concrete description of the facial subgraphs of an almost 2-connected, locally finite planar graph Γ.

Let Γ_0 be its 2-connected core. The closure of every facial subgraph of Γ_0 in the Freudenthal compactification is a simple closed curve by Proposition IV.1.27. Let $U \in \mathcal{D}(\Gamma)$, and let $U_0 \in \mathcal{D}(\Gamma_0)$ be the unique element such that $U_0 \supset U$. Assume that $\operatorname{diam}_\Gamma(f)$ is sufficiently large (possibly infinite). Then $f = \mathcal{F}[U]$ is obtained from $f_0 = \mathcal{F}[U_0]$ by attaching boundedly small cacti graphs to f_0 at cut vertices. We may refer to f_0 as the *simple body* of f, and the components of $f \setminus f_0$ as the *adjoined cacti* of f. See Figure IV.8 for cartoon.

We also note the following lemma, which sheds some light on how the friendly-faced property interacts with chains on subgraphs.

LEMMA IV.5.5. *Let Γ be a locally finite, connected planar graph. Let $\Lambda \subset \Pi \subset \Gamma$ be connected subgraphs, such that*

(1) *Π has uniform coboundary in Γ,*

(2) Λ *is a friendly-faced subgraph of* Γ.

Then Λ *is also a friendly-faced subgraph of* Π.

PROOF. Since Π has uniform coboundary in Γ, it is clear that ϑ restricts to a good drawing of Π. Note also that Π is quasi-isometrically embedded into Γ, by Proposition IV.3.6.

Let $U \in \mathcal{D}(\Lambda)$. Using the fact that Λ is a friendly-faced subgraph of Γ, let $W \in \mathcal{D}$ be such that $W \subset U$ and

$$\mathcal{F}[U_0] \subset B_\Gamma(\mathcal{F}[W]; \varepsilon),$$

for some uniform $\varepsilon > 0$. Let $V \in \mathcal{D}(\Pi)$ be the unique face such that $U \supset V \supset W$.

Let $x \in \mathcal{F}[U]$. Then there exists a path of length at most ε in Γ connecting x to some point in $\mathcal{F}[W]$. But clearly this path must intersect $\mathcal{F}[V]$. Since Π is quasi-isometrically embedded in Γ, there exists a path of bounded length in Π connecting x to $\mathcal{F}[V]$. Since x was arbitrary, it follows immediately that Λ is a friendly-faced subgraph of Γ. $\square$

In the presence of an appropriate group action, it is often possible to 'thicken up' a subgraph into a friendly-faced subgraph. In particular, it is possible to prove the following.

PROPOSITION IV.5.6. *Let* Γ *be a connected, locally finite, cocompact, planar graph. Let* $\Lambda \subset \Gamma$ *be a 2-connected, cocompactly stabilised subgraph with uniform coboundary. Then there exists some* $\delta > 0$ *such that* $B_\Gamma(\Lambda; \delta)$ *is a friendly-faced subgraph of* Γ.

We will need to prove a stronger, more coarse version of this. To this end, the rest of this section is devoted to proving the following.

THEOREM IV.5.7 (Neighbourhoods with friendly faces). *Let* X *be a connected, locally finite, cocompact graph and* $Y \subset X$ *a connected, cocompactly stabilised subgraph with uniform coboundary. Let* Γ *be a connected, bounded valence, planar graph, and* $\varphi : X \twoheadrightarrow \Gamma$ *a continuous, surjective quasi-isometry. Suppose further that* $\varphi(Y)$ *is almost 2-connected. Then there exists* $\delta > 0$ *such that*

$$\varphi(B_X(Y; \delta))$$

is a friendly-faced subgraph of Γ.

Note that if we set $X = \Gamma$ and take φ to be the identity map, then we recover Proposition IV.5.6.

The reader short on time may proceed to the next section, if they so wish. The rest of § IV.5 is devoted to proving Theorem IV.5.7, and nothing beyond this point in this section will make a second appearance.

IV.5.2. Coboundary diameters and quasi-isometries. The first step towards proving Theorem IV.5.7 is to study how coboundary diameters (see Definition IV.3.1) interact with quasi-isometries. In particular, we will show the following.

THEOREM IV.5.8 (Neighbourhoods with uniform coboundary). *Let X be a connected, locally finite, cocompact graph and $Y \subset X$ a connected, cocompactly stabilised subgraph with uniform coboundary. Let Γ be a connected, bounded valence graph and $\varphi : X \twoheadrightarrow \Gamma$ a continuous, surjective quasi-isometry. Then there exists $\delta > 0$ such that*

$$\varphi(B_X(Y;\delta))$$

has uniform coboundary in Γ.

We will prove this through a series of lemmas.

LEMMA IV.5.9. *Let Γ be a connected, bounded valence graph and $\Lambda \subset \Pi \subset \Gamma$ be connected subgraphs. If both $\|\Lambda\|_{\mathrm{in}}^{\Gamma}$ and $\mathrm{Haus}_\Gamma(\Lambda, \Pi)$ are finite then $\|\Pi\|_{\mathrm{in}}$ is finite too.*

PROOF. Write $k = \|\Lambda\|_{\mathrm{in}}^{\Gamma}$, $r = \mathrm{Haus}_\Gamma(\Lambda, \Pi)$. Fix a connected component U of $\Gamma \setminus \Pi$. Then there exists a unique component V of $\Gamma \setminus \Lambda$ such that $U \subset V$. Let $x, y \in \delta U \cap \Pi$. Choose $x', y' \in \Lambda$ such that

$$\mathrm{d}_\Gamma(x,x') = \mathrm{d}_\Gamma(x,\Lambda), \quad \mathrm{d}_\Gamma(y,y')\, \mathrm{d}_\Gamma(x,\Lambda).$$

In particular, x', y' lie adjacent in δV and a geodesic in connecting x to x' or y to y' is internally disjoint from Λ. Such a geodesic is contained in $B_\Gamma(\delta V; r)$ and so has bounded length, say length at most $L \geq 0$, since δV contains boundedly many vertices and Γ is bounded valence. Note that L depends only on r, Γ and Λ.

There is a path through Λ of length at most k connecting x' to y'. It follows that there is a path through Π of length at most $k + 2L$ connecting x to y. Since x, y and U were arbitrary, it follows that $\|\Pi\|_{\mathrm{in}} \leq k + 2L$. □

LEMMA IV.5.10. *Let X, Γ be connected bounded valence graphs, and let $\varphi : X \twoheadrightarrow \Gamma$ be a surjective quasi-isometry. Let $Y \subset X$ be a subgraph. Then for all $\varepsilon > 0$ there exists $\delta > 0$ such that*

$$\varphi(B_X(Y;\delta)) \supset B_\Gamma(\varphi(Y);\varepsilon).$$

PROOF. Choose $x \in B_\Gamma(\varphi(Y);\varepsilon)$. Since φ is surjective, let $z \in \varphi^{-1}(x)$. We have that $\mathrm{d}_X(z, Y)$ is uniformly bounded, and so the lemma follows. □

LEMMA IV.5.11. *Let X, Γ be connected bounded valence graphs, and let $\varphi : X \to \Gamma$ be a continuous quasi-isometry. Let $Y \subset X$ a subgraph such that $\|Y\|_{\mathrm{in}}$ is finite. Then there exists $\varepsilon_0 > 0$ such that for all $\varepsilon > \varepsilon_0$ we have that $\|B_\Gamma(\varphi(Y);\varepsilon)\|_{\mathrm{in}}$ is finite.*

PROOF. Let $\psi : \Gamma \to X$ be a choice of continuous quasi-inverse to φ. Fix $\lambda \geq 1$ so that φ and ψ are (λ, λ)-quasi-isometries and λ-quasi-inverses to each other. Let $\varepsilon_0 = \lambda^2 + 2\lambda$ and fix $\varepsilon > \varepsilon_0$. To ease notation, write $N = B_\Gamma(\varphi(Y); \varepsilon)$.

Let U be some connected component of $\Gamma \setminus N$. Since $\varepsilon > \lambda^2$, we have that $\psi(U)$ will be contained entirely in some connected component of $X \setminus Y$. Call this connected component V. Let $u, v \in \delta U \cap N$, and let γ be a path connecting u to v, such that $\gamma \cap N = \{u, v\}$. We have that $\psi(\gamma)$ is a path connecting $\psi(u)$ to $\psi(v)$, contained entirely within V.

Let p_1 be a geodesic of length ε from u to $\varphi(Y)$, say terminating at $a \in \varphi(Y)$. Similarly let p_2 be a geodesic of length ε joining y to some $b \in \varphi(Y)$. We have that $\psi(a)$, $\psi(b)$ lie at most a distance λ from Y, so let $a', b' \in Y$ lie a minimal distance from $\psi(a)$, $\psi(b)$ respectively. Let q_1 be a geodesic joining $\psi(a)$ to a' and q_2 a geodesic joining b' to $\psi(b)$. Let

$$q = q_1 \psi(p_1) \psi(\gamma) \psi(p_2) q_2$$

be the concatenation of these paths. This is a path which starts and ends in Y. Since $\varepsilon > \lambda^2$, we have that $\psi(\gamma)$ is disjoint from Y, and so $\psi(u)$ and $\psi(v)$ lie in the same connected component of $X \setminus Y$. There will exist some subpath l_1 of $q_1\psi(p_1)$ which connects $\psi(u)$ to some $c \in Y$, which is otherwise disjoint from Y. Similarly, there will exist some subpath l_2 of $q_2\psi(p_2)$ connecting $\psi(v)$ to some $d \in Y$ which is otherwise disjoint from Y. In particular, since $\|Y\|_{\text{in}}$ is finite, say $\|Y\|_{\text{in}} = k$, there is some path l_3 contained in Y of length at most k connecting c to d. Note that the endpoints of l_3 lie at most a distance λ from $\psi(p_1)$, $\psi(p_2)$.

Now, let $l = \varphi(l_3)$. This is a path of length at most $\lambda k + \lambda$, whose endpoints lie at a distance of at most $\lambda^2 + 2\lambda < \varepsilon$ from p_1 and p_2. Join l to p_1 and p_2 with geodesics f_1, f_2. Note that f_1 and f_2 are contained within N. Thus, the union $p_1 \cup p_2 \cup f_1 \cup f_2 \cup l$ contains a path q connecting u to v, and is contained entirely within N. Finally, note that q contains at most $4\varepsilon + \lambda k + \lambda$ edges, and so since U, u, v were arbitrary it follows that $\|N\|_{\text{in}}$ is finite. $\square$

LEMMA IV.5.12. *Let X, Γ be connected bounded valence graphs, and let $\varphi : X \to \Gamma$ be a continuous quasi-isometry. Let G act upon X, and let $Y \subset X$ be a connected, cocompactly stabilised subgraph. Suppose that $\|Y\|_{\text{in}}$ and $\|\varphi(Y)\|_{\text{in}}$ are finite. Then $\varphi(Y)$ has uniform coboundary.*

PROOF. Let $\psi : \Gamma \to X$ be a choice of continuous quasi-inverse to φ. Fix $\lambda \geq 1$ so that φ and ψ are (λ, λ)-quasi-isometries and λ-quasi-inverses to each other. Fix a connected component U of $\Gamma \setminus \varphi(Y)$. Let $v, u \in \delta U \cap U$. We need to find a path of bounded length through U connecting v to u. Clearly there is some simple path in U connecting them together, so denote this path p. We will replace p with a path p' of bounded length.

Fix $\varepsilon > 2\lambda^3 + \lambda^2$. Decompose p into a composition of paths

$$p = p_0 q_1 p_1 \dots q_n p_n,$$

where each p_i is contained in the closed ε-neighbourhood of $\varphi(Y)$, and each q_i is contained in the complement of the open ε-neighbourhood of Y. See Figure IV.9 for a cartoon of this decomposition. Let x_i, y_i denote the endpoints of q_i. We will first find a short path in U connecting x_i to y_i.

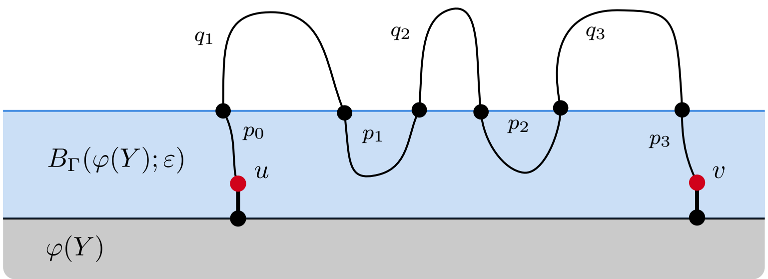

FIGURE IV.9. Decomposing the path p into the p_i and q_i.

Fix i, and to ease notation let $x = x_i$, $y = y_i$, $q = q_i$. We have that

$$\mathrm{d}_\Gamma(x, \varphi(Y)) = \mathrm{d}_\Gamma(y, \varphi(Y)) = \varepsilon.$$

Let $q' = \psi(q)$, which is a path connecting $x' := \psi(x)$ to $y' := \psi(y)$. We have that q' is a contained in the complement of the ε'-neighbourhood of Y, where

$$\varepsilon' = \tfrac{1}{\lambda}\varepsilon - \lambda > 2\lambda^2 > 0.$$

We also have that $x', y' \in B_X(Y; r)$, where $r = \lambda\varepsilon + \lambda$. Choose $a, b \in B_X(Y; \varepsilon' + 1)$ such that $\mathrm{d}_X(a, x')$, $\mathrm{d}_X(a, x')$ are minimal. In particular, these distances are at most r.

Since Y is cocompactly stabilised, so is $N := B_X(Y; \varepsilon')$. Thus, N has uniform coboundary. The vertices $a, b \in \delta N$ lie in the same connected component of $X \setminus N$, and so there is some path in $X \setminus N$ connecting them of bounded length, say k, where this bound depends only on ε' and Y. Let l_2 be such a path of minimal length. Let l_1 be a geodesic connecting x' to a and l_3 a geodesic connecting b to y'. The concatenation $l = l_1 l_2 l_3$ is a path of length at most $2r + k$ connecting x' to y', and is disjoint from N. Let $l' = \varphi(l)$. This is a path of length at most $\lambda(2r + k) + \lambda$. Moreover, we have that

$$\inf_{z \in l} \mathrm{d}_\Gamma(z, \varphi(Y)) \geq \tfrac{1}{\lambda}\varepsilon' - \lambda > \tfrac{1}{\lambda}(2\lambda^2) - \lambda > 0.$$

The endpoints of l' are $x'' := \varphi \circ \psi(x)$ and $y'' := \varphi \circ \psi(y)$, which lie a distance at most λ from x and y respectively. Join x to x'' by a geodesic of length at most λ, and similarly join y to y''. These geodesics are disjoint from $\varphi(Y)$ since $\varepsilon > \lambda$. Adjoin these geodesics to the start and end of l' and form a new path l'' connecting x to y of length at most $\lambda(2r + k) + 2\lambda$. Importantly, l'' has a uniformly bounded length depending only on Y and the quasi-isometries. and is disjoint from $\varphi(Y)$. Thus, we can replace q with this path l''.

Returning to our original set-up, thanks to the above we may now assume that each path q_i is a path of uniformly bounded length. Since Γ is bounded valence and $|\delta U|$ is bounded, $B_\Gamma(\delta U; \varepsilon)$ contains at most boundedly many vertices, say M. Since each q_i starts at a unique vertex in

$B_\Gamma(\delta U; \varepsilon)$, we see that $n \leq M$. Also, the union of the p_i contains at most M vertices. It follows that

$$\mathrm{d}_U(u, v) \leq M(\lambda(2r + k) + 2\lambda) + M,$$

This bound depends only on X, Y, Γ, and the quasi-isometries φ, ψ. Since U, u, and v were arbitrary it follows that $\|\varphi(Y)\|_{\mathrm{out}}$ is finite, and thus $\varphi(Y)$ has uniform coboundary. □

Piecing the above together we can deduce the main result of this subsection, which we restate for the convenience of the reader.

Theorem IV.5.8 (Neighbourhoods with uniform coboundary). *Let X be a connected, locally finite, cocompact graph and $Y \subset X$ a connected, cocompactly stabilised subgraph with uniform coboundary. Let Γ be a connected, bounded valence graph and $\varphi : X \twoheadrightarrow \Gamma$ a continuous, surjective quasi-isometry. Then there exists $\delta > 0$ such that*

$$\varphi(B_X(Y; \delta))$$

has uniform coboundary in Γ.

Proof. By Lemma IV.5.11 we may choose $\varepsilon > 0$ so that $\|B_\Gamma(\varphi(Y); \varepsilon)\|_{\mathrm{in}}$ is finite. Apply Lemma IV.5.10 and choose $\delta > 0$ such that $Y' := \varphi(B_X(Y; \delta)) \supset B_\Gamma(\varphi(Y); \varepsilon)$. By Lemma IV.5.9 we have that $\|Y'\|_{\mathrm{in}}$ is finite. Then by Lemma IV.5.12 we have that Y' has uniform coboundary, which is exactly what we needed to show. □

IV.5.3. Neighbourhoods with friendly faces. We now work towards a proof of Theorem IV.5.7. Most of the heavy lifting will be done by the next lemma.

Lemma IV.5.13. *Let Γ be a bounded valence, connected, planar graph with fixed good drawing ϑ. Let $\Lambda \subset \Gamma$ be an almost 2-connected subgraph with uniform coboundary. Then there exists $r > 0$ such that if $\Pi \subset \Gamma$ is a connected subgraph satisfying*

$$B_\Gamma(\Lambda; r) \subset \Pi \quad \textit{and} \quad \mathrm{Haus}_\Gamma(\Lambda, \Pi) < \infty,$$

then Π is a friendly-faced subgraph of Γ.

Proof. Since Λ has uniform coboundary, we know by Proposition IV.3.4 that the inclusions $\Lambda \hookrightarrow \Gamma$, $\Pi \hookrightarrow \Gamma$ induce injections on the sets of ends, and so ϑ restricts to a good drawing of both Λ and Π. Thus, it makes complete sense for us to speak of the faces of these subgraphs.

Let $r = \|\Lambda\|^{\Gamma}_{\mathrm{out}}$, and let $\Pi \subset \Gamma$ be such that

$$B_\Gamma(\Lambda; r) \subset \Pi \quad \text{and} \quad \mathrm{Haus}_\Gamma(\Lambda, \Pi) < \infty.$$

Let $U_1 \in \mathcal{D}(\Pi)$, and let U_0 be the unique element in $\mathcal{D}(\Lambda)$ such that $U_1 \subset U_0$.

Since Λ is almost 2-connected, we have that $\mathcal{F}[U_0]$ is formed of a *simple body* and *adjoined cacti* (see Remark IV.5.4 for a description of the faces of an almost 2-connected planar graph).

Claim IV.5.14. *Let $x, y \in \mathcal{F}[U_1] \cap \Lambda$. Then there is $U_2 \in \mathcal{D}(\Gamma)$ such that $U_2 \subset U_1$ and $x, y \in \mathcal{F}[U_2]$.*

Proof. Clearly $x, y \in \mathcal{F}[U_0]$. Suppose for the sake of a contradiction that such a $U_2 \in \mathcal{D}(\Gamma)$ does not exist. Then there must exist a simple path p in Γ such that the initial and terminal vertices of p lie on f_0, and p is otherwise drawn entirely within U_0, and also such that $\vartheta(x)$, $\vartheta(y)$ do not lie in the closure of a common component of $U_0 \setminus \vartheta(p)$. See Figure IV.10 for a cartoon.

Let u be the second vertex to appear along p, and v be the penultimate. In particular, u and v lie are adjacent to Λ, and are contained in the same connected component C of $\Gamma \setminus \Lambda$. Note that $\vartheta(C) \subset U$. Since $\|\Lambda\|^{\Gamma}_{\text{out}} = r$, we have that $\mathrm{d}_C(x, y) \leq r$. Thus, there exists a path of bounded length through C connecting x to y. In particular, we may assume that the path p we constructed above has length at most $r + 2$. This path will be contained in Π. But this contradicts the assumption that x and y both lie on the facial subgraph bordering U_1, where $U_1 \in \mathcal{D}(\Pi)$ was such that $U_1 \subset U_0$. Thus, the claim follows. □

We can extend the above claim from two vertices to arbitrarily many, via the following technical statement.

Claim IV.5.15. *There exists some uniform $N > 0$ depending only on Λ such that the following holds. Let S be a set of vertices in $\mathcal{F}[U_0]$ with $|S| > N$. Suppose S satisfies the following:*

(†) *For any two $x, y \in S$ there exists $U \in \mathcal{D}(\Gamma)$ such that $U \subset U_0$ and $x, y \in \mathcal{F}[U]$.*

Then there exists $U_2 \in \mathcal{D}(\Gamma)$ such that $U_2 \subset U_0$ and $S \subset \mathcal{F}[U_2]$.

Proof. Assume $N > 3$. Suppose to the contrary that this were not true. Then it is easy to see that there must exist a tripod T embedded in $|\Gamma|_{\mathrm{Fr}}$ with the following properties:

(1) The leaves of T lie in $\mathcal{F}[U_0]$, and T otherwise maps into U_0 under ϑ,
(2) $U_0 \setminus \vartheta(T)$ contains exactly three connected components V_1, V_2, V_3,
(3) There exists three distinct $x_1, x_2, x_3 \in S$ such that
$$\vartheta(x_i) \in (\overline{V_i} \cap \overline{V_{i+1}}) \setminus \overline{V_{i+2}},$$
where indices are taken modulo 3.

See Figure IV.11 for a cartoon.

Now, let us proceed to add more elements from S to this picture while preserving (†). In particular, it is not hard to see that any $y \in S$ much satisfy one of the following properties:

(1) Either y lies on the simple body of $\mathcal{F}[U_0]$, and coincides precisely with a leaf of T, or
(2) The vertex y lies on an adjoined cactus C of $\mathcal{F}[U_0]$, and C also contains a leaf of T.

If neither of these happens, then there must exist some x_i such that y and x_i do not lie on a common subface of U_0, and thus we contradict (†).

Finally, note that an adjoined cactus can only contain boundedly many vertices, since Γ is bounded valence, say at most m vertices. Setting $N > 3m$, the claim follows. ☐

We now show that Π is friendly-faced. Note that since Λ has uniform coboundary, Γ is bounded valence, and $\mathrm{Haus}_\Gamma(\Pi, \Lambda)$ is finite, we see that each connected component of $\Pi \setminus \Lambda$ has at most boundedly many vertices. Thus, if $\mathrm{diam}_\Pi(\mathcal{F}[U_1])$ is sufficiently large then we must have that

$$\mathcal{F}[U_1] \subset B_\Pi(\mathcal{F}[U_1] \cap \Lambda; m),$$

where $m > 0$ is some uniform constant. We also assume that $\mathcal{F}[U_1]$ is large enough so that $\mathcal{F}[U_1] \cap \Lambda$ contains more than N vertices, where $N > 0$ is as in Claim IV.5.15. Now, by Claims IV.5.14 and IV.5.15, there exists some $U_2 \in \mathcal{D}(\Gamma)$ such that $U_2 \subset U_1$, and

$$\mathcal{F}[U_1] \cap \Lambda \subset \mathcal{F}[U_2].$$

We thus deduce that $\mathcal{F}[U_1]$ lies inside of a bounded neighbourhood of $\mathcal{F}[U_2]$ in Γ. Since U_1 was an arbitrary face of Λ with $\mathcal{F}[U_1]$ sufficiently large, it follows that Π is friendly-faced. ☐

We are now able to prove Theorem IV.5.7, which we restate for the reader's convenience.

Theorem IV.5.7 (Neighbourhoods with friendly faces). *Let X be a connected, locally finite, cocompact graph and $Y \subset X$ a connected, cocompactly stabilised subgraph with uniform coboundary. Let Γ be a connected, bounded valence, planar graph, and $\varphi : X \twoheadrightarrow \Gamma$ a continuous, surjective quasi-isometry. Suppose further that $\varphi(Y)$ is almost 2-connected. Then there exists $\delta > 0$ such that*

$$\varphi(B_X(Y;\delta))$$

is a friendly-faced subgraph of Γ.

Proof. By Lemmas IV.3.3 and IV.5.9 we have that $B_X(Y;\varepsilon)$ has uniform coboundary for every $\varepsilon > 0$. By Theorem IV.5.8, there exists some $\delta' > 0$ such that $\varphi(B_X(Y;\delta'))$ has uniform coboundary in Γ. Now, combining Lemma IV.5.10 with Lemmas IV.5.13, we deduce that there is some $\delta > \delta'$ such that $\varphi(B_X(Y;\delta))$ is friendly-faced. ☐

IV.6. Quasi-actions on planar graphs

The next technical challenge is to study quasi-actions on planar graphs and their subgraphs. We open with a discussion of what is to come.

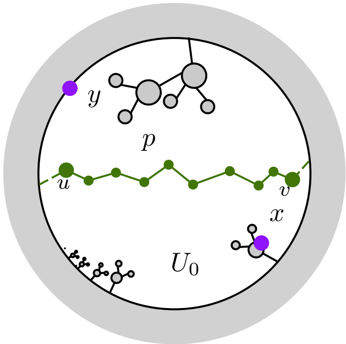

FIGURE IV.10. There exists a path p in Γ passing through U_0 which separates x from y in this face.

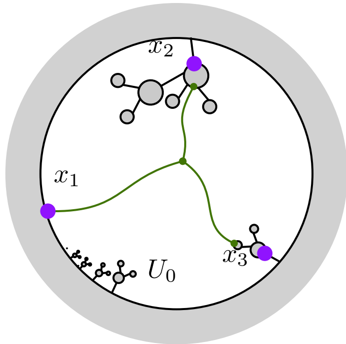

FIGURE IV.11. If x_1, x_2, x_3 pairwise lie on subfaces of U_0, but share no common face, then a tripod as above must exist in $|\Gamma|_{\mathrm{Fr}}$.

IV.6.1. Initial discussion. The following definition will be useful for brevity.

DEFINITION IV.6.1 (Quasi-planar tuple). A *quasi-planar tuple* is a sextuple

$$\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi),$$

satisfying the following:

(1) X is a connected, locally finite, multi-ended graph,
(2) G is a group acting freely and cocompactly on X,
(3) Γ is a connected, bounded valence planar graph with good drawing $\vartheta : |\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$,
(4) The map $\varphi : X \twoheadrightarrow \Gamma$ is a continuous surjective quasi-isometry with continuous quasi-inverse $\psi : \Gamma \to X$.

By Propositions IV.1.7, IV.1.9, we lose nothing by assuming that Γ is bounded valence, or that the quasi-isometries are continuous. The reason we want φ to be surjective is so we can apply Lemma IV.5.10 later in this section.

Note that there is an induced quasi-action of G upon Γ (see Definition I.5.2). The following definition sets out criteria for this quasi-action to be 'nice'.

DEFINITION IV.6.2 (Well-behaved quasi-planar tuple). Let $\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi)$ be a quasi-planar tuple. We then define two *good behaviour* properties as follows:

(GB1) There exists $m, n > 0$ such that the following holds. Let $f \in \mathcal{F}(\Gamma)$ satisfy $\mathrm{diam}_\Gamma(f) > n$, then for all $g \in G$ there exists $f' \in \mathcal{F}(\Gamma)$ such that

$$\mathrm{Haus}_\Gamma(f', \varphi_g(f)) < m.$$

(GB2) Γ is 2-connected.

We say that $\mathfrak{Q}$ is *well-behaved* if both of the above are satisfied.

The property (GB1) should be interpreted as a coarse analogue to the planar embedding theorem of WHITNEY, which states that the drawing of a 3-connected planar graph in $\mathbf{S}^2$ is unique, up to post-composition with homeomorphisms of $\mathbf{S}^2$. In particular, automorphisms of a 3-connected planar graph will take 'faces to faces'. With this in mind, (GB1) intuitively asserts that the drawing of Γ is 'coarsely preserved' by the induced quasi-action. The reason we ask for (GB2) is that it will help us apply results of § IV.4.4 later on.

REMARK IV.6.3. While (GB1) only asserts the existence of f' such that $\mathrm{Haus}_\Gamma(f', \varphi_g(f))$ is bounded for any given f and g, if $f \in \mathcal{F}^\infty(\Gamma)$ then we note that this f' is necessarily unique. This follows, for example, from the observation that at most one infinite facial subgraph accumulates in any given end. If this were not the case for some end ω, then ω would be a local cut point of $|\Gamma|_{\mathrm{Fr}}$, contradicting Proposition I.5.3.

Recall that $\mathcal{F}^{\mathrm{c}}(\Gamma)$ denotes the set of compact facial subgraphs of Γ. Note that (GB1) has the following consequence.

PROPOSITION IV.6.4. *Let $\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi)$ be a quasi-planar tuple, and suppose $\mathfrak{Q}$ satisfies (GB1). Then there exists $R > 0$ such that* $\mathrm{diam}_\Gamma(f) < R$ *for all $f \in \mathcal{F}^{\mathrm{c}}(\Gamma)$.*

PROOF. Suppose to the contrary that Γ contains a sequence $(f_n)_{n\geq 1}$, $f_n \in \mathcal{F}^{\mathrm{c}}(\Gamma)$ of compact facial subgraphs such that $\mathrm{diam}(f_n) \geq n$ for all $n \geq 1$. Let $v_0 \in \Gamma$ be arbitrary. Pick a vertex v_n on each f_n and let $g_n \in G$ be such that $\varphi_{g_n}(v_n)$ lies within a bounded distance of v_0 (recall the quasi-action of G on Γ is cobounded). By (GB1) we deduce that a bounded neighbourhood of v_0 must intersect compact facial subgraphs of arbitrarily large diameter. This contradicts the fact that Γ is locally finite. □

In § IV.7, we will see that good behavior is enough to deduce accessibility. The sole purpose of § IV.6 is to reduce the general problem to the well-behaved case. In particular, we will prove the following statement, which we state now for the convenience of the reader.

THEOREM IV.6.5 (Splittings with well-behaved vertex groups). *Let G be a finitely generated group, and suppose that G is quasi-isometric to some connected planar graph Γ with good drawing ϑ. Then G splits as a graph of groups $G(\Theta) = (G_x, \alpha_e)$ such that*

(1) *Each edge group G_e is finite,*
(2) *For each vertex group G_u, one of the following holds:*
   (a) *G_u has at most one end, or*
   (b) *There is a well-behaved quasi-planar tuple*

$$\mathfrak{Q}_u = (G_u, X_u, \Gamma_u, \vartheta_u, \varphi_u, \psi_u),$$

where $\Gamma_u \subset \Gamma$ and ϑ_u is the restriction of ϑ to the closure of Γ_u in $|\Gamma|_{\mathrm{Fr}}$.

Apart from this main theorem, no other results of § IV.6 will make a second appearance outside of this section. Thus, the reader who is short on time may now skip to § IV.7, if they so wish.

IV.6.2. Intersecting neighbourhoods of facial subgraphs. We first need to prove some technical lemmas which tell us that if Γ is a planar graph which is quasi-isometric to some cocompact graph X, then intersections of tubular neighbourhoods of facial-subgraphs are boundedly small.

LEMMA IV.6.6. *Let X be a connected, cocompact, locally finite graph. Let Γ be a connected, locally finite, planar graph. Suppose further that Γ is almost 2-connected. Let $\varphi : X \to \Gamma$ be a quasi-isometry with quasi-inverse $\psi : \Gamma \to X$. Then for every $r > 0$ there exists $m = m(r) > 0$ such that the following holds. Let $f_1, f_2 \in \mathcal{F}(\Gamma)$ be distinct. Then*

$$\operatorname{diam}_\Gamma \big(B_\Gamma(f_1; r) \cap B_\Gamma(f_2; r)\big) \leq m.$$

PROOF. We will prove this by contradicting Lemma IV.3.17. We first prove this under the stronger assumption that Γ is 2-connected.

Thus, assume Γ is 2-connected. Fix $r > 0$, and suppose that for every $m > 0$ there exists a pair of distinct $f_1, f_2 \in \mathcal{F}(\Gamma)$, and a pair of points

$$x, y \in B_\Gamma(f_1; r) \cap B_\Gamma(f_2; r)$$

such that $\mathrm{d}_\Pi(x, y) > m$. Note that since Γ is 2-connected, the closure of every facial subgraph of Γ in $|\Gamma|_{\mathrm{Fr}}$ is a simple closed curve Proposition IV.1.27. Let $\overline{f_i}$ denote the closure of f_i in $|\Gamma|_{\mathrm{Fr}}$.

We assume without loss of generality that m is much larger than r. It is clear that there exist (possibly degenerate) simple paths p, q in Γ such that

(1) The path p is contained within a bounded neighbourhood of x, and similarly q lies in a bounded neighbourhood of y.
(2) Both p and q begin on f_1 and end on f_2, and are otherwise disjoint from $f_1 \cup f_2$.
(3) The paths p and q are disjoint from each other.

It is easy to see pictorially that $|\Gamma|_{\mathrm{Fr}} \setminus (p \cup q)$ is disconnected, see Figure IV.12.

Consider f_1. Since $\overline{f_1} \cong \mathbf{S}^1$ and $f_1 \cap (p \cup q)$ consists of exactly two vertices, say u and v, we have that $\overline{f_1} \setminus (p \cup q) = \overline{f_1} \setminus \{u, v\}$ is a disjoint union of two open intervals, each lying in a distinct component of $|\Gamma|_{\mathrm{Fr}} \setminus (p \cup q)$. Let I_1, I_2 denote these components. Let e_1, e_2 be the unique two edges of Γ such that

(1) Both e_1, e_2 lie inside f_1, and in particular in (the closure of) I_1.
(2) We have that e_1 abuts p, and e_2 abuts q.

Let U_1 denote the connected component of $|\Gamma|_{\mathrm{Fr}} \setminus (p \cup q)$ containing I_1. This is a connected open subset of $|\Gamma|_{\mathrm{Fr}}$. By Proposition I.5.3, we have that $C_1 = U_1 \setminus \Omega(\Gamma)$ is connected. In particular, there is a combinatorial path through C_1 connects endpoints of e_1 and e_2. Similarly, we also see that there is a combinatorial path through $\Gamma \setminus C_1$ connecting endpoints of e_1 to e_2. Finally, let

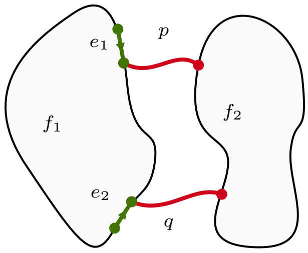

Figure IV.12

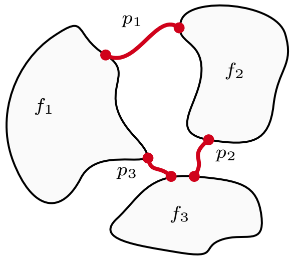

Figure IV.13

b denote the vertex set of C_1. We have that every edge in δb abuts p or q, so δb is finite and in particular uniformly bounded in size. Thus, $b \in \mathscr{B}(\Gamma)$. It is also clear that $e_1, e_2 \in \delta b$. By Proposition IV.3.14, there exists a tight $b' \in \mathscr{B}(\Gamma)$ such that $\delta b' \subset \delta b$, and $e_1, e_2 \in \delta b'$. In particular, $\mathrm{diam}_\Gamma(\delta b')$ is approximately at least m but $|\delta b'|$ is uniformly bounded. By taking $m \to \infty$, we eventually contradict Lemma IV.3.17.

Now suppose that Γ is just almost 2-connected, but not necessarily 2-connected. Let $\Gamma_0 \subset \Gamma$ be the 2-connected core. Once again, fix $r > 0$, and suppose that for every $m > 0$ there exists a pair of distinct $f_1, f_2 \in \mathcal{F}(\Gamma)$ such that

$$\mathrm{diam}_\Gamma \left(B_\Gamma(f_1; r) \cap B_\Gamma(f_2; r) \right) > m.$$

Recall (Remark IV.5.4) that every facial subgraph of Γ decomposes into its simple body (which is a facial subgraph of Γ_0) plus boundedly small adjoined cacti. It follows quickly from this that for some $r' > 0$ and for every $m' > 0$ there exists $f_1', f_2' \in \mathcal{F}(\Gamma_0)$ such that

$$\mathrm{diam}_{\Gamma_0} \left(B_{\Gamma_0}(f_1'; r) \cap B_{\Gamma_0}(f_2'; r) \right) > m'.$$

But we know that this cannot happen, since Γ_0 is also quasi-isometric to X and 2-connected. Thus, the lemma follows. ☐

We can also prove a similar result for three faces.

Lemma IV.6.7. *Let X be a connected, cocompact, locally finite graph. Let Γ be a connected, locally finite, planar graph. Let $\varphi : X \to \Gamma$ be a quasi-isometry with quasi-inverse $\psi : \Gamma \to X$. Suppose further that Γ is almost 2-connected. Then for every $r > 0$ there exists $n = n(r) > 0$ such that the following holds. Let $f_1, f_2, f_3 \in \mathcal{F}(\Gamma)$ be pairwise distinct. Suppose there exists*

$$x_1 \in B_\Gamma(f_1; r) \cap B_\Gamma(f_2; r), \quad x_2 \in B_\Gamma(f_2; r) \cap B_\Gamma(f_3; r), \quad x_3 \in B_\Gamma(f_3; r) \cap B_\Gamma(f_1; r).$$

Then $\mathrm{diam}_\Gamma(\{x_1, x_2, x_3\}) \leq n$.

Proof. This argument is very similar to that of Lemma IV.6.6, so we shall only give a sketch and leave the details to the reader.

As in the proof of Lemma IV.6.6, we can assume that Γ is 2-connected without any real loss of generality. Fix $r > 0$ and suppose that for every $n > 0$ there exists distinct facial subgraphs $f_1, f_2, f_3 \in \mathcal{F}(\Gamma)$ and points

$$x_1 \in B_\Gamma(f_1; r) \cap B_\Gamma(f_2; r), \quad x_2 \in B_\Gamma(f_2; r) \cap B_\Gamma(f_3; r), \quad x_3 \in B_\Gamma(f_3; r) \cap B_\Gamma(f_1; r)$$

such that $\operatorname{diam}_\Gamma(\{x_1, x_2, x_3\}) > n$. Since we assume Γ is 2-connected, the closure of each f_i in $|\Gamma|_{\mathrm{Fr}}$ is a simple closed curve by Proposition IV.1.27.

For each $i = 1, 2, 3$, let p_i be a simple path connecting f_i to f_{i+1} which is contained a bounded neighbourhood of x_i, and intersects each of f_i and f_{i+1} in exactly one vertex. Here, indices are to be taken modulo 3. It is clear pictorially that $|\Gamma|_{\mathrm{Fr}} \setminus (p_1 \cup p_2 \cup p_3)$ is disconnected (see Figure IV.13). Without loss of generality, p_1 and p_2 are at distance approximately n from each other. In fact, we may assume that all three p_i are pairwise far apart, as if p_2 is near p_3 then we deduce that the intersection of bounded neighbourhoods of f_1 and f_2 has large diameter, and we can apply Lemma IV.6.6. Thus, without loss of generality the p_i are pairwise disjoint.

As in the proof of Lemma IV.6.6 we now construct a tight element $b \in \mathscr{B}(\Gamma)$ such that $|\delta b|$ is uniformly bounded but $\operatorname{diam}_\Gamma(\delta b)$ is approximately at least n. This leads to a contradiction of Lemma IV.3.17, and so we are done. □

IV.6.3. Combinatorial Jordan curves. We now prove the following two results, which will be helpful for showing that two points lie near a common face later on.

LEMMA IV.6.8. *Let Γ be a connected, locally finite, planar graph with fixed good drawing ϑ. Given a finite collection $U_1, \dots, U_n \in \mathcal{D}(\Gamma)$, let V be a connected component of $\mathbf{S}^2 \setminus \bigcup_i \overline{U}_i$. Then $\overline{V}$ is locally (path) connected, and every connected component of ∂V is a simple closed curve.*

PROOF. In what follows, we will suppress ϑ from our notation for the sake of readability. That is, we will identify $|\Gamma|_{\mathrm{Fr}}$ with its ϑ-image in $\mathbf{S}^2$.

We first show that $\overline{V}$ is locally connected, i.e. that for every $x \in \overline{V}$ and every open subset $W \subset \mathbf{S}^2$ containing x, there exists an open subset $O \subset W$ with $x \in O$ and $O \cap \overline{V}$ connected. The only non-trivial case to check is when $x \in \Omega(\Gamma)$, i.e. x is an end of Γ. By Proposition I.5.3, we have that x is not a local cut-point of $|\Gamma|_{\mathrm{Fr}}$, so in particular we have that x lies inside the closure of at most one of the U_i. By replacing W with a smaller open subset, we may therefore reduce to the case where $n = 1$. Let us then simplify notation and write $U := U_1$.

Since $|\Gamma|_{\mathrm{Fr}}$ and $\mathbf{S}^2$ are locally path connected, let $O \subset W$ be an open sub-neighbourhood of X such that both O and $|\Gamma|_{\mathrm{Fr}} \cap O$ are (path) connected. It is a standard fact from plane topology that any open connected subset of the plane cannot be disconnected by the removal of any totally disconnected set. In particular, since $\partial V \subset |\Gamma|_{\mathrm{Fr}}$, we have that $\partial V \cap O$ is either empty or contains some $v_0 \in \Gamma$. That is, v_0 is not an end of Γ. After possibly subdividing an edge of Γ, we may assume without loss of generality that v_0 is a vertex of Γ.

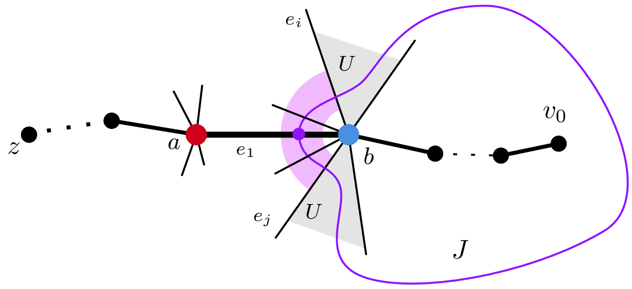

FIGURE IV.14. Jordan curve J separating v_0 from z. The purple region is disjoint from $\overline{V}$.

Now, let $z \in |\Gamma|_{\mathrm{Fr}} \cap \overline{V} \cap O$ be arbitrary. Since $O \cap |\Gamma|_{\mathrm{Fr}}$ is path connected, there exists a path p connecting z to v_0 contained in O. By an application of Proposition I.5.3, we may assume that p is a combinatorial (possibly one-way infinite) path. We claim that p is contained entirely in $\overline{V}$. Suppose not, then there exists a 'first' vertex a along p which is not contained $\overline{V}$. Let b be the vertex immediately preceding a along p, so $b \in \overline{V}$. Let $e \in E(\Gamma)$ be the edge in p connecting a to b, so $e \cap \overline{V} = \{b\}$. Consider now the clockwise cyclic sequence of edges $e_1, \ldots, e_m$ incoming to b, say $e = e_1$. Inspections reveals that there must exist $1 \leq i \leq j \leq m$ such that

- we have $e_i, e_j \subset \overline{U}$, and
- for all ℓ such that either $1 \leq \ell \leq i$ or $j \leq \ell \leq m$, the interior of e_ℓ is disjoint from $\overline{V}$.

Indeed, if this were not the case then it would necessarily follow that $a \in \overline{V}$. Note, it could be that $i = j$.

Using this observation, along with the fact that U is path connected (it is open and connected) we may draw a Jordan curve J which is disjoint from $\overline{V}$ and also separates v_0 from z, as depicted in Figure IV.14. This contradicts the fact that $\overline{V}$ is connected. It follows that $|\Gamma|_{\mathrm{Fr}} \cap \overline{V} \cap O$ is connected. Since O is connected and $\partial V \subset |\Gamma|_{\mathrm{Fr}}$ we have that $\overline{V} \cap O$ is connected and so $\overline{V}$ is locally connected.

Finally, as noted above we have that any connected open set of $\mathbf{S}^2$ cannot be disconnected by the removal of any totally disconnected subset. Thus V is 2-connected and so is $\overline{V}$. By Proposition IV.1.27, it follows that every connected component of ∂V is a simple closed curve. $\square$

PROPOSITION IV.6.9. *Let Γ be a connected, locally finite, planar graph with fixed good drawing ϑ. Let $x, y \in V(\Gamma)$, $r \geq 0$. Suppose there does **not** exist $f \in \mathcal{F}(\Gamma)$ such that $x, y \in B_\Gamma(f; r)$. Then there exists a simple combinatorial loop ℓ in Γ such that*

(1) *$\vartheta(x)$, $\vartheta(y)$ lie in distinct components of $\mathbf{S}^2 \setminus \vartheta(\ell)$, and*
(2) *$\{x, y\} \cap B_\Gamma(\ell; r) = \emptyset$.*

PROOF. To ease notation we will identify $|\Gamma|_{\mathrm{Fr}}$ with its image in $\mathbf{S}^2$ and suppress mention of ϑ.

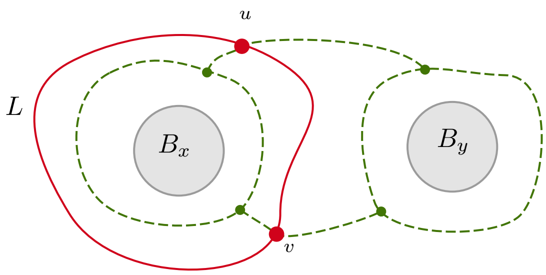

FIGURE IV.15. The figure Z (dashed) drawn in the plane.

Let $B_x = B_\Gamma(x;r)$, $B_y = B_\Gamma(y;r)$ be the closed r-balls about x and y. Let $U_1, \ldots, U_n \in \mathcal{D}(\Gamma)$ be those faces which intersect B_x. Let V be the connected component of $\mathbf{S}^2 \setminus \bigcup_i \overline{U}_i$ which contains y, and thus also contains all of B_y. Then by Lemma IV.6.8 we have that ∂V contains a simple closed curve L[6] which separates B_x from B_y. Note that $L \subset |\Gamma|_{\mathrm{Fr}}$. We will modify L and form a combinatorial loop ℓ with the same property.

Let W_x, W_y be the connected components of $\mathbf{S}^2 \setminus L$ containing x and y respectively. If L is not already a combinatorial loop then we must at least have that L contains two distinct vertices of Γ, say $u, v \in V(\Gamma)$, since $\Omega(\Gamma)$ is totally disconnected. Let γ_1, γ_2 be the connected components of $L \setminus \{v, u\}$. Now, in $\mathbf{S}^2$ (or equivalently, in the plane) draw a Jordan curve J_x contained in W_x separating L from B_x. Similarly, draw a Jordan curve J_y contained in W_y separating L from B_y. Let P_1 be an arc in W_x connecting J_x to u, otherwise disjoint from J_x and L, and P_2 an in W_x connecting J_x to v, also otherwise disjoint from J_x. Define Q_1, $Q_2 \subset \overline{W_y}$ similarly. Let

$$Z = J_x \cup J_y \cup P_1 \cup P_2 \cup Q_1 \cup Q_2.$$

Note that we do not require Z to be contained in $|\Gamma|_{\mathrm{Fr}}$. See Figure IV.15 for a cartoon. Note that $\mathbf{S}^2 \setminus Z$ is a disjoint union of four open disks: each containing exactly one of B_x, B_y, γ_1 and γ_2. Let U_i be the disk containing γ_i. Let $O_i \subset U_i$ denote the connected component of $|\Gamma|_{\mathrm{Fr}} \cap U_i$ containing γ_i. Since $|\Gamma|_{\mathrm{Fr}}$ is locally connected, we have that O_i is an open subset of $|\Gamma|_{\mathrm{Fr}}$. Now, using Proposition I.5.3 we may replace γ_i with a combinatorial path ℓ_i contained in O_i with the same endpoints, which is therefore drawn entirely inside of U_i. The loop $\ell = \ell_1 \cup \ell_2$ satisfies our requirements. □

IV.6.4. Splitting into well-behaved pieces. In this section we prove Theorem IV.6.5. This is one of the main technical achievements of this chapter, and indeed this thesis. We begin with some discussion, as what follows will be fairly technical.

It is possible to prove[7] that if $\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi)$ is a quasi-planar tuple with Γ 2-connected and $\mathrm{vs}(\Gamma)$ sufficiently large compared to the constants associated with the induced quasi-action, then

[6]In fact, since $\bigcup_i \overline{U}_i$ is connected it is not hard to deduce that $\partial V = L$.

[7]To prove this, one can adapt and simplify the arguments of this section. It suffices to take $\mathrm{vs}(\Gamma) > 3\lambda^2 + 3$.

$\mathfrak{Q}$ is well-behaved. The hope now is that we might pass to a vertex group G_v of some finite splitting of G which acts freely and cocompactly on a 'highly connected' subgraph X_v of X, and apply the aforementioned result. We must have that X_v will be quasi-isometric to some highly connected subgraph of Γ, surely?

Unfortunately, an important subtlety arises here. When we pass to a vertex group G_v of some finite splitting, simply restricting our quasi-isometries to the subgraph X_v will not quite produce a well-defined quasi-planar tuple, as we must also modify the restricted quasi-inverse slightly in order to ensure that the domains and codomains align. This slightly modifies the dynamics of the induced quasi-action and necessarily causes our error constants to increase. In particular, we move the goalposts on what we consider to be 'highly connected'. In order to break out of this cycle, some additional bookkeeping is required, which will allow us to refer back to the original quasi-action on the original super-graph. Roughly speaking, we will consider a tubular neighbourhood Y_v of our subgraph X_v, and approximate the facial subgraphs of $\varphi(Y_v)$ with facial subgraphs of $\varphi(X_v)$. The induced quasi-action on $\varphi(Y_v)$ will agree with the original quasi-action on Γ when we look specifically at its action on points in $\varphi(X_v)$. This helps us circumvent the problem of our quasi-isometry constants inflating, and is the motivation for defining friendly-faced subgraphs in § IV.5.

We now set up some notation which will follow us for the rest of § IV.6. In what follows, fix $\lambda \geq 1$ such that the induced quasi-action of G upon Γ is a λ-quasi-action (as in Definition I.5.2). We also assume that ψ is a λ-quasi-inverse to φ. Recall that by Theorem IV.3.19, we have that G splits as a graph of groups $G(\Theta) = (G_x, \alpha_e)$ such that the following hold:

(1) Each edge group G_e is finite.
(2) For each vertex group G_v acts freely and cocompactly on a subgraph $X_v \subset X$ with uniform coboundary.
(3) The quasi-isometry φ restricts to a quasi-isometry $\varphi_v : X_v \twoheadrightarrow \varphi(X_v)$.
(4) For each $v \in V(\Theta)$, we have that either:
    (a) X_v has at most one end, or
    (b) X_v is multi-ended, and $\mathrm{vs}(\varphi(X_v)) > 3\lambda^2 + 3$.

It is easy to see that we may replace each X_v above with a bounded uniform neighbourhood of itself and not affect any of the above properties. In particular, by Theorem IV.5.7 we may assume the following without loss of generality:

- For each $u \in V(\Theta)$, we have that $\varphi(X_u)$ is a friendly-faced subgraph of Γ.

Of course, if each X_u had at most one end then G would be accessible and we would be finished. Thus, let us now fix $u \in V(\Theta)$ such that X_u is multi-ended. The notation we fix here will follow us through the rest of this section. Let

$$Z := X_u, \quad \Lambda := \varphi(Z), \quad H := G_u.$$

We now define a certain tubular neighbourhood of Z. Let

$$Y := B_X(Z; \lambda), \quad \Pi := \varphi(Y).$$

Also, let $\mu : Y \twoheadrightarrow \Pi$ denote the restriction of φ to Y. By Proposition IV.1.10, μ is a quasi-isometry (where Y and Π are considered with their own intrinsic path-metric). We need to be a bit careful about how we choose a quasi-inverse to μ.

LEMMA IV.6.10. *There exists a quasi-isometry $\nu : \Pi \to Y$ such that the following hold:*

(1) *ν is a quasi-inverse to μ,*

(2) *$\nu|_\Lambda = \psi|_\Lambda$.*

In particular, for all $g \in H$ and $x \in \Lambda$ we have that $\varphi_g(x) = \mu_g(x)$.

PROOF. Since ψ is a λ-quasi-inverse to φ, we have that $\psi(\Lambda) \subset B_X(Z; \lambda) = Y$. It is immediate from this observation that such a ν exists. $\square$

Note that since Λ and Π have uniform coboundary, we have by Remark IV.3.5 that the restriction of ϑ to closure of one of these subgraphs in $|\Gamma|_{\mathrm{Fr}}$ is also a good drawing.

We record the following helpful consequence of Λ being friendly-faced.

LEMMA IV.6.11. *There exists $M > 0$ such that for every $f_1 \in \mathcal{F}(\Lambda)$ there exists $f_2 \in \mathcal{F}(\Pi)$ with*

$$\mathrm{Haus}_\Gamma(f_1, f_2) < M.$$

PROOF. By Lemma IV.5.5, we have that Λ is a friendly-faced subgraph of Π. Thus, for every $U \in \mathcal{D}(\Lambda)$ there exists $U' \in \mathcal{F}(\Pi)$ such that $U' \subset U$ and $\mathcal{F}[U]$ is contained in a bounded neighbourhood of $\mathcal{F}[U']$. But since Π is contained in a bounded neighbourhood of Λ, it is clear that actually we must have $\mathrm{Haus}_\Gamma(\mathcal{F}[U], \mathcal{F}[U'])$ is uniformly bounded. $\square$

The importance of this is that now we may study the quasi-action of H on the faces of Π by approximating them as faces of Λ. Then, we may apply Lemma IV.6.10 to inspect their quasi-translates by referring back to the quasi-action of G on Γ, for which we have 'small constants'.

For the rest of § IV.6, we will fix all of the above notation.

We now apply the results of the previous section and study the quasi-action of H on the faces of Π. We will need the following easy lemma, relating to finite planar graphs. We leave the proof as an exercise.

LEMMA IV.6.12. *Let Γ be a finite, connected planar graph. Suppose that Γ contains:*

(1) *Two disjoint connected subgraphs Λ_1, Λ_2.*

(2) *Three pairwise disjoint paths α_1, α_2, α_3 connecting Λ_1 to Λ_2.*

(3) *Three paths β_1, β_2, β_3 such that:*

  (a) *Each β_i connects α_i to α_{i+1}, but is disjoint from α_{i+2}, where indices are taken modulo 3, and*

  (b) *Each β_i is disjoint from both Λ_1 and Λ_2.*

Then the image of any embedding ϑ of Γ in $\mathbf{S}^2$ contains a Jordan curve separating $\vartheta(\Lambda_1)$ from $\vartheta(\Lambda_2)$. In particular, no $x \in \Lambda_1$ and $y \in \Lambda_2$ can lie on a common facial subgraph of Γ in any embedding.

See Figure IV.16 for an illustration of Lemma IV.6.12.

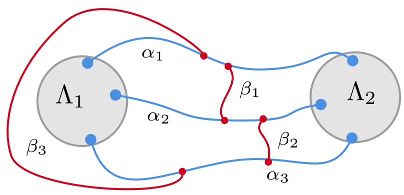

FIGURE IV.16

We now have the following application of MENGER's theorem, which motivates why we wanted $\mathrm{vs}(\Lambda)$ to be 'large'.

LEMMA IV.6.13. *The subgraph $\Lambda \subset \Gamma$ enjoys the following properties:*

(1) *For every distinct pair of ends $\omega_1, \omega_2 \in \Omega(\Lambda)$, there exists three distinct bi-infinite paths $\alpha_1, \alpha_2, \alpha_3 : \mathbf{R} \to \Lambda$ such that each α_i connects ω_1 to ω_2.*

(2) *There exists $S > 0$ such that for all $s \geq S$ and all $x, y \in \Lambda$ such that $\mathrm{d}_\Lambda(x, y) > 2s$, there exists three distinct paths $\alpha_1, \alpha_2, \alpha_3 : [0,1] \to \Lambda$ such that each α_i begins in $B_\Lambda(x; s)$ and ends in $B_\Lambda(y; s)$.*

Moreover, in each case we have that the α_i satisfy the following further property. We have that

$$\mathrm{d}_\Gamma(z, \alpha_i) > \lambda^2$$

for every $j \neq i$, $z \in \alpha_j$.

REMARK IV.6.14. It is very important to note that the lower bound given in the above lemma is about how far apart these rays are **in the super-graph** Γ, not just in Λ or Π.

PROOF OF LEMMA IV.6.13. (1): Let $\omega_1, \omega_2 \in \Omega(\Lambda)$. Since $\mathrm{vs}(\Lambda) \geq 3\lambda^2 + 3$, we have by MENGER's theorem (IV.1.5) that there exists $N := 3\lambda^2 + 3$ pairwise disjoint bi-infinite paths paths $\beta_1, \ldots, \beta_N \subset \Lambda$, between ω_1 and ω_2. Since $\Lambda \subset \Gamma$ is planar with a fixed drawing ϑ, we have that there is a fixed cyclic order in which these paths emerge from ω_1. Assume without loss of generality that β_i lies adjacent to $\beta_{i\pm 1}$ in this order, where indices are taken modulo N. Let

$$\alpha_1 = \beta_1, \quad \alpha_2 = \beta_{\lambda^2+2}, \quad \alpha_3 = \beta_{2\lambda^2+3}.$$

For every $i \neq j$, there are at least λ^2 distinct Jordan curves contained in $|\Gamma|_{\mathrm{Fr}}$ which separate α_i from α_j. Moreover, these curves only intersect in $\{\omega_1, \omega_2\}$. Thus, we see through an application of the Jordan curve theorem that any path in Γ from α_i to α_j must internally intersect at least λ^2 disjoint rays in Γ. Thus, we deduce that $\mathrm{d}_\Gamma(z, \alpha_i) > \lambda^2$ for every $z \in \alpha_j$. See Figure IV.17 for a cartoon.

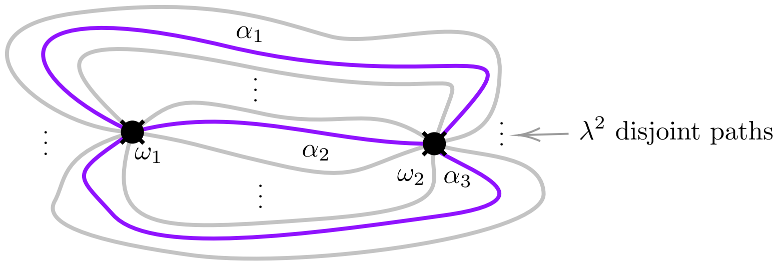

FIGURE IV.17. Any path through Γ between α_i and α_j must pass through λ^2 distinct β_k paths different from the α_i and α_j.

(2): Let $S \geq 0$ be sufficiently large so that $B_\Lambda(z; s)$ separates ends in Λ for every $s \geq S$, $z \in \Lambda$. Since Λ admits a cobounded quasi-action, such an S certainly exists. Fix $s \geq S$. Let $x, y \in \Lambda$ such that $\mathrm{d}_\Lambda(x, y) > 2s$. We have that both $B_1 := B_\Lambda(x; s)$ and $B_2 := B_\Lambda(y; s)$ separate ends in Λ. We have that B_1 and B_2 are disjoint. Let $\omega_1, \omega_2 \in \Omega(\Lambda)$ be distinct ends which lie in distinct components of both $\Lambda \setminus B_1$ and $\Lambda \setminus B_2$. By (1) above we have that there exists α_1, α_2, α_3 in Λ connecting ω_1 to ω_2 which sit disjoint from each others λ^2-neighbourhood in Γ. All three paths must intersect both B_1 and B_2, and we are done. □

LEMMA IV.6.15. *There exists a constant $r > 0$ such that the following holds. Given $f \in \mathcal{F}(\Pi)$, $x, y \in f$, and $g \in H$, there exist $f' \in \mathcal{F}(\Pi)$ such that $\mu_g(x), \mu_g(y) \in B_\Pi(f'; r)$.*

PROOF. Fix $g \in H$. To ease notation, write $x' = \mu_g(x)$, $y' = \mu_g(y)$. We will need some constants. Recall that $\lambda \geq 1$ has been fixed so that the quasi-action of G upon Γ is a λ-quasi-action. Fix $\eta \geq 1$ such that the quasi-isometric embeddings $\Lambda \hookrightarrow \Pi$, $\Pi \hookrightarrow \Gamma$, and $\Lambda \hookrightarrow \Gamma$ are all (η, η)-quasi-isometric embeddings, and also assume without loss of generality that $\mathrm{Haus}_\Pi(\Lambda, \Pi) \leq \eta$.

Let $a, b \in \Lambda$ be such that $\mathrm{d}_\Pi(x', a) \leq \eta$ and $\mathrm{d}_\Pi(y', b) \leq \eta$. Apply Lemma IV.6.13, and obtain three paths α_1, α_2, α_3 in Λ such that

(1) Each α_i has one endpoint in $B_\Lambda(a; S)$ and one in $B_\Lambda(b; S)$.
(2) We have that $\mathrm{d}_\Gamma(z, \alpha_i) > \lambda^2$ for every $j \neq i$, $z \in \alpha_j$,

where $S \geq 0$ is some fixed constant depending only on the graphs and quasi-isometries at play. Note that property (2) relates to the path metric of Γ, not Λ or Π. To ease notation, let

$$A_a = B_\Lambda(a; S), \ A_b = B_\Lambda(b; S),$$

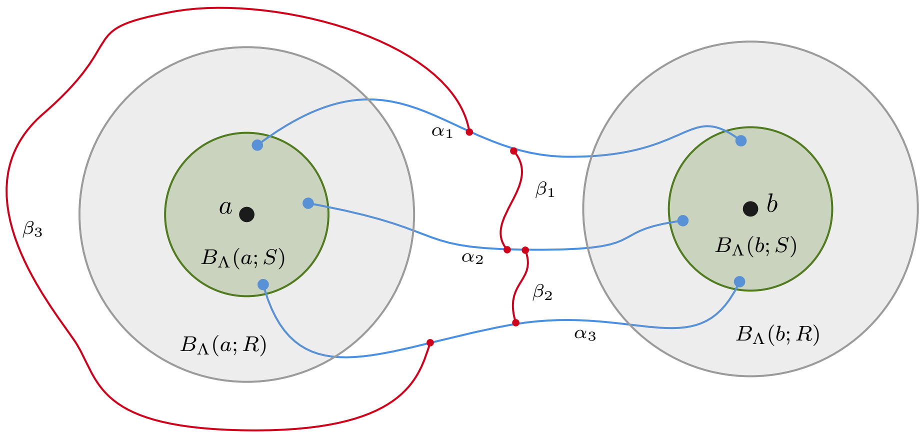

FIGURE IV.18. Construction of Ξ.

so each α_j begins in A_a and ends in A_b. Choose some large $R > 0$ such that

$$B_\Gamma(A_a;\lambda^2+\eta+1)\cap\Lambda\subset B_\Lambda(a;R),\qquad B_\Gamma(A_b;\lambda^2+\eta+1)\cap\Lambda\subset B_\Lambda(b;R).$$

Clearly such a uniform R exists, since Λ is quasi-isometrically embedded into Γ. We assume without loss of generality that $\mathrm{d}_\Lambda(a,b) > 2R$, lest x', y' lie a bounded distance apart in Π, and the statement is vacuously true.

Suppose there exists some facial subgraph $f_0 \in \mathcal{F}(\Lambda)$ such that

$$\{a,b\}\subset B_\Lambda(f_0;R).$$

Then by Lemma IV.6.11 there exists a $f_1 \in \mathcal{F}(\Pi)$ such that

$$\{x',y'\}\subset B_\Pi(f_1;R+M+\eta),$$

where $M > 0$ is some fixed constant. With this in mind, let $r = R + M + \eta$, and assume that there is no such $f_0 \in \mathcal{F}(\Lambda)$. We will show that x and y cannot lie on a common facial subgraph, and thus deduce a contradiction.

We have by Proposition IV.6.9 that there is a simple loop ℓ in Λ such that

(1) $\vartheta(a)$, $\vartheta(b)$ lie in distinct components of $\mathbf{S}^2 \setminus \vartheta(\ell)$, and
(2) $\{a,b\}\cap B_\Lambda(\ell;R)=\emptyset$.

Let β_i be a segment of ℓ connecting α_i to α_{i+1}, which is otherwise disjoint from all the α_j. Note that β_i sits outside of the λ^2-neighbourhood of α_{i+2} in Γ, by an easy application of the Jordan curve theorem, since any path between these two curves must intersect either α_i or α_{i+1}, or a bounded neighbourhood of a, b. Since ℓ is very far away from a and b, and the α_j are pairwise far apart, the claim follows. Let

$$\Xi := \beta_1\cup\beta_2\cup\beta_3\cup\alpha_1\cup\alpha_2\cup\alpha_3.$$

See Figure IV.18 for a cartoon of this construction.

We now apply $\mu_{g^{-1}}$ to Ξ. In particular, we claim that $\vartheta(\mu_{g^{-1}}(\Xi))$ contains a Jordan curve which separates $\vartheta(x)$ from $\vartheta(y)$ in $\mathbf{S}^2$. Since $\Xi \subset \Lambda$, we have by Lemma IV.6.10 that

$$\mu_{g^{-1}}(\Xi) = \varphi_{g^{-1}}(\Xi).$$

For $i = 1, 2, 3$, write $\alpha_i' = \mu_{g^{-1}}(\alpha_i)$, $\beta_i' = \mu_{g^{-1}}(\beta_i)$. Note that the α_i' are pairwise disjoint, and β_i' is disjoint from α_{i+2}'.

It is clear that if R above is chosen to be sufficiently large then in the Π-metric we have that x and y are strictly closer each of the α_i' than they are to any of the β_j'. In particular, if we draw a geodesic in Π from x or y to some α_i' then this geodesic will not intersect any of the β_j'. We now apply Lemma IV.6.12 and deduce that x and y cannot lie on a common facial subgraph of Π. $\square$

Lemma IV.6.16. *There exists $C, D \geq 0$ such that the following holds. For every $f \in \mathcal{F}(\Pi)$ such that* $\operatorname{diam}_\Pi(f) \geq D$ *and every $g \in H$, there exists $f' \in \mathcal{F}(\Pi)$ such that*

$$\mathrm{Haus}_\Pi(f', \mu_g(f)) \leq C.$$

Proof. Let $f \in \mathcal{F}(\Pi)$, $g \in H$. We first prove that there exists some $f' \in \mathcal{F}(\Pi)$ such that

$$\mu_g(f) \subseteq B_\Pi(f'; C'),$$

for some uniform constant $C' > 0$. We may assume without loss of generality that $\operatorname{diam}_\Pi(f)$ is bounded below by some constant, lest this claim is vacuously true. We will specify this lower bound later.

Using Lemma IV.6.6, let $m = m(r) > 0$ be such that if the r-neighbourhoods of any two distinct faces intersects at two points a, b, then $\mathrm{d}_\Pi(a, b) < m$. Similarly, let $n = n(r) > 0$ be as in Lemma IV.6.7, so if the r-neighbourhoods of three distinct faces all pairwise intersect, then the union of these intersections has diameter at most n. Let $M = \max\{n, m\}$, and fix $\eta \geq 1$ so that the induced quasi-action of H upon Π is a η-quasi-action.

We assume without loss of generality that $\operatorname{diam}_\Pi(f) \geq \eta(3M + \eta)$ so that

$$\operatorname{diam}_\Pi(\mu_g(f)) \geq 3M.$$

Fix $x \in \mu_g(f)$, and let $y \in \mu_g(f)$ be such that $\mathrm{d}_\Pi(x, y) \geq M$. By Lemma IV.6.15, there exists $f_1 \in \mathcal{F}(\Pi)$ such that both x and y are contained in the r-neighbourhood of f_1. We claim that $\mu_g(f)$ is contained in the $(r + M)$-neighbourhood of f_1. Suppose to the contrary that there exists some $z \in \mu_g(f)$ such that z is not contained in the $(r+M)$-neighbourhood of f_1. Applying Lemma IV.6.15 again, we see that there exists $f_2, f_3 \in \mathcal{F}(\Pi)$ such that

$$x, z \in B_\Pi(f_2; r), \quad y, z \in B_\Pi(f_3; r).$$

Clearly, both f_2 and f_3 must be distinct from f_1, since f_1 is far away from z. Now, either f_2 and f_3 are equal or they are distinct. If they are equal then we contradict Lemma IV.6.6 since $\mathrm{d}_\Pi(x, y)$ is large. If they are distinct then we contradict Lemma IV.6.7 since $\operatorname{diam}_\Pi(\{x, y, z\})$ is large. In any case, we deduce that $\mu_g(f)$ is contained in the $(r+M)$-neighbourhood of f_1, and so our earlier claim is proven by choosing $C' = r + M$.

We now prove that the Hausdorff distance between f_1 and $\mu_g(f)$ must be uniformly bounded, provided $\mathrm{diam}_\Pi(f)$ is sufficiently large. By the above claim, there exists $f_0 \in \mathcal{F}(\Pi)$ such that

$$\mu_h(f_1) \subset B_\Pi(f_0; C').$$

In particular, f is contained in the r'-neighbourhood of f_0, where $r' = C' + \eta C' + 2\eta$. Let $m' = m(r') > 0$ be as in Lemma IV.6.6. If $\mathrm{diam}_\Pi(f) > m'$ then this forces $f_0 = f$. It follows quickly from this observation that f_1 is contained in a bounded neighbourhood of $\mu_g(f)$, and this proves the Lemma. □

By Remark IV.3.5, the closure of Π in the Freudenthal compactification of Γ is homeomorphic to that very compactification of Π. Thus, ϑ restricts to a good drawing of Π. Let $\vartheta' : |\Pi|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$ denote this restriction. Let $\mathfrak{P}$ denote the quasi-planar tuple

$$\mathfrak{P} = (H, Y, \Pi, \vartheta', \mu, \nu).$$

Note that Lemma IV.6.16 immediately implies the following.

Lemma IV.6.17. *The quasi-planar tuple $\mathfrak{P}$ defined above satisfies (GB1).*

We now address the fact that Π may not be 2-connected. This turns out to be an easy fix since Π is already almost 2-connected.

Lemma IV.6.18. *Let $\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi)$ be a quasi-planar tuple which satisfies (GB1). Suppose that Γ is almost 2-connected. Then there exists a **well-behaved** quasi-planar tuple*

$$\mathfrak{Q}' = (G, X, \Gamma', \vartheta', \varphi', \psi').$$

We may take Γ' to be the 2-connected core of Γ and ϑ' to be the restriction of ϑ.

Proof. Let Γ_0 be the 2-connected core of Γ. Recall (see Remark IV.5.4) that every facial subgraph of Γ decomposes as a simple body (which is a facial subgraph of Γ_0) plus some adjoined cacti of bounded diameter.

Recall that the inclusion $\iota : \Gamma_0 \hookrightarrow \Gamma$ is a quasi-isometry. Let $\rho : \Gamma \to \Gamma_0$ be some choice of quasi-inverse to the inclusion ι. Let $\varphi' = \rho \circ \varphi$, and let ψ' be any choice of quasi-inverse to μ. Note that the map ϑ restricts to a good drawing ϑ' of Γ_0. Let $\mathfrak{Q}' = (G, X, \Gamma_0, \vartheta', \varphi', \psi')$. Trivially, $\mathfrak{Q}'$ satisfies (GB2) since Γ_0 is 2-connected. It is also easy to verify that $\mathfrak{Q}'$ satisfies (GB1), since the induced quasi-action of G on Γ_0 differs from that of G upon on Γ only by a bounded amount, and the (large enough) faces of Γ_0 lie a bounded Hausdorff distance from faces of Γ. □

Combining Lemmas IV.6.17 and IV.6.18, we immediately deduce Theorem IV.6.5 which we restate below for the convenience of the reader.

THEOREM IV.6.5 (Splittings with well-behaved vertex groups). *Let G be a finitely generated group, and suppose that G is quasi-isometric to some connected planar graph Γ with good drawing ϑ. Then G splits as a graph of groups $G(\Theta) = (G_x, \alpha_e)$ such that*

(1) *Each edge group G_e is finite,*
(2) *For each vertex group G_u, one of the following holds:*
   (a) *G_u has at most one end, or*
   (b) *There is a well-behaved quasi-planar tuple*

$$\mathfrak{Q}_u = (G_u, X_u, \Gamma_u, \vartheta_u, \varphi_u, \psi_u),$$

where $\Gamma_u \subset \Gamma$ and ϑ_u is the restriction of ϑ to the closure of Γ_u in $|\Gamma|_{\mathrm{Fr}}$.

This can be interpreted as a coarse version of the statement that every finitely generated group with a planar Cayley graph admits as a finite splitting such that every vertex group admits a 3-connected planar Cayley graph, which thus admits a unique drawing in $\mathbf{S}^2$.

IV.7. Main results

IV.7.1. Accessibility and applications. In this section we will prove the main result of this chapter, that a finitely generated group which is quasi-isometric to a planar graph is accessible. In fact, we essentially prove on the way that our group is finitely presented, and thus reduce our problem to the theorem of DUNWOODY [**46**]. We briefly state this theorem now, for reference below.

THEOREM IV.7.1 (DUNWOODY). *Every finitely presented group is accessible.*

In order to effectively reduce our problem to the above, we will need the following definition.

DEFINITION IV.7.2 (Bad loops). Let Γ be a connected, locally finite, planar graph with good drawing $\vartheta : |\Gamma|_{\mathrm{Fr}} \hookrightarrow \mathbf{S}^2$. We say that a simple combinatorial loop ℓ in Γ is a *bad loop* if there exists two ends $\omega_1, \omega_2 \in \Omega(\Gamma)$ such that $\vartheta(\omega_1)$ and $\vartheta(\omega_2)$ lie in distinct components of $\mathbf{S}^2 \setminus \vartheta(\ell)$.

LEMMA IV.7.3. *Let Γ be a connected, bounded valence, planar graph such that every $f \in \mathcal{F}^{\mathrm{c}}(\Gamma)$ has uniformly bounded diameter. Then there exists $\varepsilon > 0$ such that every closed loop in Γ is either a bad loop or null-homotopic in $K_\varepsilon(\Gamma)$. In particular, if Γ contains no bad loops then Γ is coarsely simply connected.*

PROOF. Let ℓ be a closed loop in Γ. We assume without loss of generality that ℓ is simple, and not a bad loop. Thinking of ℓ as a Jordan curve in $\mathbf{S}^2$ we have that one side of ℓ contains only finitely many (necessarily finite) faces. Each face uniformly small, and hence ℓ is null-homotopic in $K_\varepsilon(\Gamma)$ for some uniform $\varepsilon > 0$. □

LEMMA IV.7.4. *Let G be a finitely generated group, and suppose that G is quasi-isometric to a connected, locally finite, planar graph with no bad loops. Then G is accessible.*

Proof. Assume without loss of generality that G is infinite-ended. Let $X = \mathrm{Cay}(G, S)$ be some choice of Cayley graph of G. Using Propositions IV.1.7, IV.1.9, we may assume without loss of generality that Γ has bounded valence, and fix a quasi-isometry $\varphi : X \to \Gamma$ which is continuous and surjective, with continuous quasi-inverse ψ. This may involve passing to a subgraph of Γ, but this subgraph (considered together with the appropriate restriction of the original drawing) will also contain no bad loops. In particular, we may assume G, X, and Γ form part of a quasi-planar tuple $\mathfrak{Q}$.

We now apply Theorem IV.6.5 and Proposition IV.6.4, and obtain a splitting of G as a graph of groups $G(\Theta) = (G_x, \alpha_e)$ such that each edge group is finite, and each vertex group G_v is quasi-isometric to a planar graph $\Gamma_v \subset \Gamma$ with no bad loops and finite facial subgraphs of uniformly bounded size. In particular, by Lemma IV.7.3 we have that Γ_v is coarsely simply connected. By Proposition IV.4.28 we have G_v is finitely presented. Now, a graph of finitely presented groups is easily seen to be finitely presented. This follows immediately from the standard presentation of a the fundamental group of a graph of groups; see [**7**], for example. In particular, we deduce that G is finitely presented, and thus accessible by Theorem IV.7.1. $\square$

Lemma IV.7.5. *Let Γ be a connected, locally finite planar graph with a good drawing ϑ, and let $\omega_1, \omega_2 \in \Omega(\Gamma)$. Suppose there exists a bad loop $\ell \subset \Gamma$ such that $\vartheta(\ell)$ separates $\vartheta(\omega_1)$, and $\vartheta(\omega_2)$. Then there is some $b \in \mathscr{B}_{\mathcal{F}}(\Gamma)$ which separates ω_1 from ω_2, where $\mathcal{F} = \mathcal{F}^\infty(\Gamma)$.*

Proof. Assume without loss of generality that ℓ is simple. Let U_1, U_2 be the two connected components of $\mathbf{S}^2 \setminus \vartheta(\ell)$. For each $f \in \mathcal{F}$, we have that $\vartheta(f)$ is contained in the closure of exactly one of the U_i. Let

$$b = \vartheta^{-1}(U_1) \cap V(\Gamma).$$

The coboundary δb contains only edges which intersect ℓ, which is finite, and so $b \in \mathscr{B}(\Gamma)$. Then, for every $f \in \mathcal{F}$ we have that either $f \cap V(\Gamma) \subset b^*$, or $f \cap b^* \subset \ell$. Since ℓ contains finitely many vertices, it follows that $b \in \mathscr{B}_{\mathcal{F}}(\Gamma)$. $\square$

Lemma IV.7.6. *Let $\mathfrak{Q} = (G, X, \Gamma, \vartheta, \varphi, \psi)$ be a well-behaved quasi-planar tuple. Then G is accessible.*

Proof. Write $\mathcal{F} = \mathcal{F}^\infty(\Gamma)$ to ease notation. Our immediate goal is to construct a system of subgraphs $\mathcal{H}$ of X such that X is coarsely CHomP relative to $\mathcal{H}$. This system will roughly correspond to the image of $\mathcal{F}$ under ψ.

Given $f \in \mathcal{F}$, $g \in G$, define $g(f)$ as the unique $f' \in \mathcal{F}$ such that $\mathrm{Haus}_\Gamma(f, f')$ is finite. Indeed, the existence of such an f' is given by (GB1), and uniqueness was noted in Remark IV.6.3. This induces a well-defined action of G by permutations upon the set $\mathcal{F}$. Given $f \in \mathcal{F}$, denote by G_f the stabiliser of f with respect to this action. Note that since $\mathcal{F}$ is locally finite (in the sense of

Definition IV.4.9), it is immediate that G acts on $\mathcal{F}$ with finitely many orbits. Let $f_1, \ldots, f_n \in \mathcal{F}$ be orbit representatives. Let $R > 0$ be some large constant which we will choose shortly. For every f_i, let

$$Z_{f_i} = G_{f_i} \psi(f_i).$$

For each $i = 1, \ldots, n$ let T_i be a transversal of G_{f_i} containing the identity. If $f = t(f_i)$ for some nontrivial $t \in t_i$, then define $Z_f = tZ_{f_i}$. It is immediate from (GB1) that the Hausdorff distance between Z_f and $\psi(f)$ is uniformly bounded above for all $f \in \mathcal{F}$, say $\mathrm{Haus}_X(Z_f, \psi(f)) < R$ for some fixed $R > 0$. For each $f \in \mathcal{F}$, let

$$Y_f = B_X(Z_f; R).$$

We have for every $f \in \mathcal{F}$ that $\psi(f) \subset Y_f$ and every component of Y_f intersects $\psi(f)$. Let

$$\mathcal{H} = \{Y_f : f \in \mathcal{F}\}.$$

We have that $\mathcal{H}$ is a G-invariant system of infinite subgraphs. Since $\mathcal{F}$ is locally finite, $\mathcal{H}$ is too. We also note that $\mathcal{H}$ is tame by, in the sense of Definition IV.4.13, by Proposition IV.4.15, since every connected component of every $Y \in \mathcal{H}$ is clearly infinite.

By Theorem IV.4.35 and Proposition IV.6.4 we have that Γ is coarsely CHomP relative to $\mathcal{F}$. Combining this with the above observations and Theorem IV.4.30, we deduce that X is coarsely CHomP relative to $\mathcal{H}$. By Corollary IV.4.22 and Proposition IV.4.23, there exists a G-invariant, nested, G-finite subset $\mathcal{E}$ of $\mathscr{B}(X)$ consisting of tight elements, such that $\mathscr{B}_{\mathcal{H}}X$ is contained in the subring generated by $\mathcal{E}$. Let $T = T(\mathcal{E})$ be the structure tree of $\mathcal{E}$. Note that G acts on T. We may assume this action is without inversions by simply subdividing each edge; this does not create any trouble. Clearly the edge stabilisers of this action are finite, since the action of G upon X was free.

We now apply Theorem IV.3.10, and for any given $v \in V(T)$ we get a subgraph $X_v \subset X$ such that:

(1) The vertex stabiliser G_v acts freely and cocompactly on X_v.
(2) Each X_v has uniform coboundary in X, and so we may canonically identify $\Omega(X_v)$ with a subset of $\Omega(X)$.
(3) If ω_1, ω_2 lie in $\Omega(X_v)$ then there is no $b \in \mathcal{E}$ which separates ω_1 and ω_2.

Now, let $\Gamma_v = \varphi(X_v)$. Since φ is continuous, we certainly have Γ_v is quasi-isometric to X_v.

Claim IV.7.7. *The subgraph $\Gamma_v \subset \Gamma$ constructed above cannot contain any bad loops, with respect to the drawing induced by the given drawing ϑ of Γ.*

Proof. Suppose Γ_v contains a bad loop separating $\vartheta(\omega_1)$ and $\vartheta(\omega_2)$, where $\omega_1, \omega_2 \in \Omega(\Gamma_v)$. Since quasi-isometries induces bijections on the sets of ends, we have that $\Omega(\Gamma_v)$ can also be canonically indentified with a subset of $\Omega(\Gamma)$. Then, since the drawing of Γ_v is just the restriction of the drawing of Γ, we find that this bad loop in Γ_v is also a bad loop in Γ, which separates this same pair of ends.

By Lemma IV.7.5 we find some $b \in \mathscr{B}_{\mathcal{F}}(\Gamma)$ separating ω_1 and ω_2. Then, by Proposition IV.4.26 we find some $b' \in \mathscr{B}_{\mathcal{H}}(X)$ which separates $\psi(\omega_1)$ and $\psi(\omega_2)$. By Proposition IV.3.15 we have that $\psi(\omega_1)$ and $\psi(\omega_2)$ are separated by some $b'' \in \mathcal{E}$. But this contradicts property (3) of X_v given above. Thus, Γ_v has no bad loops. □

Combining Lemmas IV.7.7 and IV.7.4, we immediately see that each vertex stabiliser G_v is accessible. This implies that G splits as a graph of groups with finite edge groups and accessible vertex groups. In particular, G is accessible. □

We now deduce our main theorem.

Theorem IV.7.8. *Let G be a finitely generated group. Suppose G is quasi-isometric to a planar graph. Then G is accessible.*

Proof. Using Propositions IV.1.7, IV.1.9, we may assume without loss of generality that Γ has bounded valence, and that the quasi-isometry $\varphi : X \to \Gamma$ is continuous and onto with continuous quasi-inverse ψ. In other words, X and Γ form part of a quasi-planar tuple.

By Theorem IV.6.5, we have that G splits as a graph of groups $G(\Theta)$ with finite edge groups, where every vertex group either has at most one end or forms part of a well-behaved quasi-planar tuple. By Lemma IV.7.6 we see that every vertex group is accessible. Thus, G is accessible. □

We now prove our main application.

Corollary IV.7.9. *Let G be a finitely generated group which is quasi-isometric to a planar graph. Then a finite-index subgroup of G admits a planar Cayley graph.*

Proof. Let G be a finitely generated group which is quasi-isometric to a planar graph. By Theorem IV.7.8, we have that G is accessible, and so G splits as a finite graph of groups with finite edge groups and where each vertex group is either finite or one-ended and quasi-isometric to a planar graph. By Theorem IV.2.8, the one-ended vertex groups are virtual surface groups.

This it is an easy exercise to check that G is residually finite and thus virtually torsion-free. Note that a torsion-free virtually free group is necessarily free, and a torsion-free virtual surface group is a surface group. The former follows easily from the fact that virtually free groups are finitely presented and thus accessible [**44**], and the latter follows from [**51**, Cor. 2]. In particular, some finite-index subgroup of G is isomorphic to a finite free product of free and surface groups, and thus admits a planar Cayley graph. □

REMARK IV.7.10. The above, together with a theorem of PAPASOGLU and WHYTE [**111**, Thm. 0.4], implies that there are precisely eight quasi-isometry classes of finitely generated groups quasi-isometric to planar graphs, since every surface group is quasi-isometric to either $\mathbf{R}^2$ or $\mathbf{H}^2$. In particular, every such group is quasi-isometric to one of:

$$1, \quad \mathbf{Z}, \ F_2, \quad \mathbf{Z}^2, \quad \Sigma, \quad \mathbf{Z}^2 * \mathbf{Z}^2, \quad \Sigma * \Sigma, \quad \mathbf{Z}^2 * \Sigma,$$

where F_2 is the free group of rank two, and Σ denotes the fundamental group of the closed orientable surface of genus 2. In fact, one can upgrade this from quasi-isometric to commensurable, by applying the argument presented in [**9**].

IV.7.2. Beyond Cayley graphs. In this chapter we proved Theorem IV.7.8 for Cayley graphs of finitely generated groups. It should be noted however, that the proof extends to a more general setting. Call a graph X *quasi-transitive* (or *vertex-transitive*) if $\mathrm{Aut}(X)$ acts upon X with finitely many orbits (exactly one orbit) of vertices. The most general version of this theorem then states:

THEOREM IV.7.11 ([**89**]). *Let X be a connected, locally finite, quasi-transitive graph. If X is quasi-isometric to some planar graph then X is an accessible graph.*

By an *accessible graph*, we mean in the sense of THOMASSEN–WOESS [**129**]; see Definition IV.1.2.

One can say more about the coarse geometry of a quasi-transitive graph X which is quasi-isometric to a planar graph.

THEOREM IV.7.12 ([**91**]). *Let X be a connected, locally finite, quasi-transitive graph. If X is quasi-isometric to some planar graph, then X is quasi-isometric to a planar Cayley graph.*

This is not obvious. It was asked by WOESS [**138**] whether every vertex-transitive graph is quasi-isometric to a Cayley graph. A family of potential counterexamples was conjectured by DIESTEL and LEADER [**41**]. This conjecture was later proven true by ESKIN, FISHER, and WHYTE across two celebrated papers [**52, 53**]. With this in mind, what Theorem IV.7.12 tells us is that the coarse geometry of planar graphs somehow forbids similar examples from arising within our setting.

IV.7.3. Beyond planarity. We conclude by noting that, for finitely presented groups, one can strengthen Corollary IV.7.9 even further. In particular, in [**?**] we prove the following.

THEOREM IV.7.13 ([**?**]). *Let G be a finitely presented group. Then G is virtually planar if and only if G is asymptotically minor-excluded.*

The definition of *asymptotically minor-excluded* is central to the new field of *coarse graph theory*, and is beyond the scope of this thesis. See [**63**] for survey on this topic. We will simply note the following corollary, which follows immediately from Theorem IV.7.13 and is a direct strengthening of Corollary IV.7.9 for finitely presented groups.

Corollary IV.7.14. *Let G be a finitely presented group. If G is quasi-isometric to a minor-excluded graph, then G is virtually planar.*

Recall that a graph is said to be *minor-excluded* if there exists a finite graph which is not a minor of it. By the classical Kuratowski theorem, this class includes all planar graphs. It also includes, for example, those graphs with topological embeddings into finite genus surfaces.

It is conjectured by Georgakopoulos and Papasoglu [**63**] that Theorem IV.7.13 should extend to all finitely generated groups, though this is currently wide-open. In particular, it is unclear how to deal with the issue of accessibility, which was the main technical challenge of this chapter.

Bibliography

[1] S. I. Adian, *Unsolvability of some algorithmic problems in the theory of groups*, Trudy Moskovskogo Matematicheskogo Obshchestva, 6 (1957), pp. 231–298.

[2] A. Akhmedov, *Big tiles in hyperbolic groups*, arXiv preprint arXiv:2309.02607, (2023).

[3] A. Akhmedov and D. Fulghesu, *Arithmetic sets in groups*, Math. Z., 292 (2019), pp. 1195–1206.

[4] G. N. Arzhantseva, *On quasiconvex subgroups of word hyperbolic groups*, Geometriae dedicata, 87 (2001), pp. 191–208.

[5] G. N. Arzhantseva and P.-A. Cherix, *On the Cayley graph of a generic finitely presented group*, Bulletin of the Belgian Mathematical Society-Simon Stevin, 11 (2004), pp. 589–601.

[6] B. Barrett, *Computing JSJ decompositions of hyperbolic groups*, Journal of Topology, 11 (2018), pp. 527–558.

[7] H. Bass, *Covering theory for graphs of groups*, Journal of pure and applied algebra, 89 (1993), pp. 3–47.

[8] G. Baumslag, C. F. Miller III, and H. Short, *Unsolvable problems about small cancellation and word hyperbolic groups*, Bulletin of the London Mathematical Society, 26 (1994), pp. 97–101.

[9] J. Behrstock, T. Januszkiewicz, and W. Neumann, *Commensurability and qi classification of free products of finitely generated abelian groups*, Proceedings of the American Mathematical Society, 137 (2009), pp. 811–813.

[10] M. Bestvina and M. Feighn, *A hyperbolic* $\mathrm{Out}(F_n)$*-complex*, Groups Geom. Dyn, 4 (2010), pp. 31–58.

[11] M. Bestvina and G. Mess, *The boundary of negatively curved groups*, Journal of the American Mathematical Society, 4 (1991), pp. 469–481.

[12] W. W. Boone, *The word problem*, Proceedings of the National Academy of Sciences, 44 (1958), pp. 1061–1065.

[13] A. Borisov and M. Sapir, *Polynomial maps over finite fields and residual finiteness of mapping tori of group endomorphisms*, Invent. Math., 160 (2005), pp. 341–356.

[14] B. H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta mathematica, 180 (1998), pp. 145–186.

[15] ———, *Splittings of finitely generated groups over two-ended subgroups*, Transactions of the American Mathematical Society, 354 (2002), pp. 1049–1078.

[16] ———, *Planar groups and the Seifert conjecture*, Journal für die reine und angewandte Mathematik, 576 (2004), pp. 11–62.

[17] ———, *Tight geodesics in the curve complex*, Invent. Math., 171 (2008), pp. 281–300.

[18] ———, *Notes on Maskits planarity theorem*, LEnseignement Mathématique, 68 (2022), pp. 1–24.

[19] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, vol. 319, Springer Science & Business Media, 2013.

[20] M. R. Bridson and H. Wilton, *The triviality problem for profinite completions*, Inventiones mathematicae, 202 (2015), pp. 839–874.

[21] T. Brown, *Uncountably many quasi-isometry classes of groups of type FP via graphical small cancellation theory*, PhD thesis, University of Southampton, 1 2022.

[22] J. O. Button and R. P. Kropholler, *Nonhyperbolic free-by-cyclic and one-relator groups*, New York J. Math., 22 (2016), pp. 755–774.

[23] M. Calvez, *Euclidean Artin-Tits groups are acylindrically hyperbolic*, Groups Geom. Dyn., 16 (2022), pp. 963–983.

[24] A. Casson and D. Jungreis, *Convergence groups and Seifert fibered 3-manifolds*, Inventiones mathematicae, 118 (1994), pp. 441–456.

[25] A. Cayley, *On the theory of groups, as depending on the symbolic equation* $\theta^n = 1$, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 7 (1854), pp. 40–47.

[26] ———, *Desiderata and suggestions: No. 2. The Theory of groups: graphical representation*, American journal of mathematics, 1 (1878), pp. 174–176.

[27] B. Chandler and W. Magnus, *The history of combinatorial group theory: a case study in the history of ideas*, vol. 9, Springer Science & Business Media, 2012.

[28] R. Charney and R. Morris-Wright, *Artin groups of infinite type: trivial centers and acylindrical hyperbolicity*, Proc. Amer. Math. Soc., 147 (2019), pp. 3675–3689.

[29] C. Chou, *Elementary amenable groups*, Illinois J. Math., 24 (1980), pp. 396–407.

[30] L. Ciobanu, D. Holt, and S. Rees, *Equations in groups that are virtually direct products*, J. Algebra, 545 (2020), pp. 88–99.

[31] F. Dahmani and D. Groves, *Detecting free splittings in relatively hyperbolic groups*, Transactions of the American Mathematical Society, (2008), pp. 6303–6318.

[32] F. Dahmani and V. Guirardel, *The isomorphism problem for all hyperbolic groups*, Geometric and Functional Analysis, 21 (2011), pp. 223–300.

[33] F. Dahmani, V. Guirardel, and D. Osin, *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*, Mem. Amer. Math. Soc., 245 (2017), pp. v+152.

[34] M. Dehn, *Über die Topologie des dreidimensionalen Raumes*, Mathematische Annalen, 69 (1910), pp. 137–168.

[35] ———, *Über unendliche diskontinuierliche Gruppen*, Mathematische Annalen, 71 (1911), pp. 116–144.

[36] ———, *Transformation der Kurven auf zweiseitigen Flächen*, Mathematische Annalen, 72 (1912), pp. 413–421.

[37] ———, *Die beiden Kleeblattschlingen*, Mathematische Annalen, 75 (1914), pp. 402–413.

[38] G.-A. Diao and M. Feighn, *The Grushko decomposition of a finite graph of finite rank free groups: an algorithm*, Geometry & Topology, 9 (2005), pp. 1835–1880.

[39] W. Dicks and M. J. Dunwoody, *Groups acting on graphs*, vol. 17, Cambridge University Press, 1989.

[40] R. Diestel, *Graph Theory*, Springer-Verlag, 5th ed., 2017.

[41] R. Diestel and I. Leader, *A conjecture concerning a limit of non-Cayley graphs*, Journal of algebraic combinatorics, 14 (2001), pp. 17–25.

[42] C. Droms, *Infinite-ended groups with planar Cayley graphs*, Journal of Group Theory, 9 (2006), pp. 487–496.

[43] C. Druţu and M. Kapovich, *Geometric group theory*, vol. 63, American Mathematical Soc., 2018.

[44] M. J. Dunwoody, *Accessibility and groups of cohomological dimension one*, Proceedings of the London Mathematical Society, 3 (1979), pp. 193–215.

[45] ———, *Cutting up graphs*, Combinatorica, 2 (1982), pp. 15–23.

[46] ———, *The accessibility of finitely presented groups.*, Inventiones mathematicae, 81 (1985), pp. 449–458.

[47] ———, *An inaccessible group*, Geometric group theory, 1 (1993), pp. 75–78.

[48] ———, *Planar graphs and covers*, arXiv preprint arXiv:0708.0920, (2007).

[49] W. Dyck, *Gruppentheoretische Studien*, Mathematische Annalen, 20 (1882), pp. 1–44.

[50] B. Eckmann and P. Linnell, *Poincaré duality groups of dimension two, II*, Commentarii Mathematici Helvetici, 58 (1983), pp. 111–114.

[51] B. Eckmann and H. Müller, *Poincaré duality groups of dimension two*, Commentarii Mathematici Helvetici, 55 (1980), pp. 510–520.

[52] A. Eskin and D. Fisher, *Quasi-isometric rigidity of solvable groups*, in Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures, World Scientific, 2010, pp. 1185–1208.

[53] A. Eskin, D. Fisher, and K. Whyte, *Coarse differentiation of quasi-isometries I: Spaces not quasi-isometric to Cayley graphs*, Annals of Mathematics, (2012), pp. 221–260.

[54] L. Esperet, U. Giocanti, and C. Legrand-Duchesne, *The structure of quasi-transitive graphs avoiding a minor with applications to the domino problem*, arXiv preprint arXiv:2304.01823, (2023).

[55] H. Fleischner and W. Imrich, *Transitive planar graphs*, Mathematica Slovaca, 29 (1979), pp. 97–106.

[56] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Mathematische Zeitschrift, 33 (1931), pp. 692–713.

[57] D. Gabai, *Convergence groups are Fuchsian groups*, Annals of Mathematics, 136 (1992), pp. 447–510.

[58] S. Gao, S. Jackson, and B. Seward, *Group colorings and Bernoulli subflows*, Mem. Amer. Math. Soc., 241 (2016), pp. vi+241.

[59] R. Geoghegan, *Topological methods in group theory*, vol. 243, Springer Science & Business Media, 2007.

[60] A. Georgakopoulos, *On planar Cayley graphs and Kleinian groups*, Transactions of the American Mathematical Society, 373 (2020), pp. 4649–4684.

[61] A. Georgakopoulos and M. Hamann, *The planar Cayley graphs are effectively enumerable I: consistently planar graphs*, Combinatorica, 39 (2019), pp. 993–1019.

[62] ———, *The planar Cayley graphs are effectively enumerable II*, arXiv preprint arXiv:1901.00347, (2019).

[63] A. Georgakopoulos and P. Papasoglu, *Graph minors and metric spaces*, arXiv preprint arXiv:2305.07456, (2023).

[64] R. Gitik, *On intersections of conjugate subgroups*, International Journal of Algebra and Computation, 27 (2017), pp. 403–419.

[65] M. Greendlinger, *Dehn's algorithm for the word problem*, Communications on Pure and Applied Mathematics, 13 (1960), pp. 67–83.

[66] M. Gromov, *Groups of polynomial growth and expanding maps (with an appendix by Jacques Tits)*, Publications Mathématiques de l'IHÉS, 53 (1981), pp. 53–78.

[67] ———, *Hyperbolic groups*, in Essays in group theory, Springer, 1987, pp. 75–263.

[68] ———, *Geometric Group Theory: Asymptotic invariants of infinite groups*, vol. 2, Cambridge University Press, 1993.

[69] J. R. J. Groves and G. A. Swarup, *Remarks on a technique of Dunwoody*, Journal of Pure and Applied Algebra, 75 (1991), pp. 259–269.

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

[71] V. Guirardel and G. Levitt, *JSJ decompositions of groups*, arXiv preprint arXiv:1602.05139, (2016).

[72] R. Halin, *A note on Menger's theorem for infinite locally finite graphs*, in Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, vol. 40, 1974, pp. 111–114.

[73] M. Hamann, *Group actions on metric spaces: fixed points and free subgroups*, Abh. Math. Semin. Univ. Hambg., 87 (2017), pp. 245–263.

[74] G. Higman, B. H. Neumann, and H. Neuman, *Embedding theorems for groups*, Journal of the London Mathematical Society, 1 (1949), pp. 247–254.

[75] H. Hopf, *Enden offener Räume und unendliche diskontinuierliche Gruppen*, Commentarii Mathematici Helvetici, 16 (1943), pp. 81–100.

[76] C. H. Houghton, *Ends of locally compact groups and their coset spaces*, Journal of the Australian Mathematical Society, 17 (1974), pp. 274–284.

[77] W. Jaco, D. Letscher, and J. H. Rubinstein, *Algorithms for essential surfaces in 3-manifolds*, Contemporary Mathematics, 314 (2002), pp. 107–124.

[78] W. Jaco and U. Oertel, *An algorithm to decide if a 3-manifold is a Haken manifold*, Topology, 23 (1984), pp. 195–209.

[79] I. Kapovich, *Detecting quasiconvexity: algorithmic aspects*, Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994), 25 (1996), pp. 91–99.

[80] I. Kapovich and N. Benakli, *Boundaries of hyperbolic groups*, in Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), vol. 296 of Contemp. Math., Amer. Math. Soc., Providence, RI, 2002, pp. 39–93.

[81] I. Kapovich and A. Myasnikov, *Stallings foldings and the subgroup structure of free groups*, arXiv preprint math/0202285, (2002).

[82] I. Kapovich and H. Short, *Greenberg's theorem for quasiconvex subgroups of word hyperbolic groups*, Canadian Journal of Mathematics, 48 (1996), pp. 1224–1244.

[83] F. Klein, *Vergleichende Betrachtungen über neuere geometrische Forschungen*, Mathematische Annalen, 43 (1893), pp. 63–100.

[84] P. H. Kropholler, *Baumslag-Solitar groups and some other groups of cohomological dimension two*, Comment. Math. Helv., 65 (1990), pp. 547–558.

[85] P. H. Kropholler and M. A. Roller, *Relative ends and duality groups*, Journal of Pure and Applied Algebra, 61 (1989), pp. 197–210.

[86] H. Levinson and B. Maskit, *Special embeddings of Cayley diagrams*, Journal of Combinatorial Theory, Series B, 18 (1975), pp. 12–17.

[87] M. Linton and C.-F. Nyberg-Brodda, *The theory of one-relator groups: history and recent progress*, arXiv preprint arXiv:2501.18306, (2025).

[88] R. C. Lyndon, P. E. Schupp, R. Lyndon, and P. Schupp, *Combinatorial group theory*, vol. 89, Springer, 1977.

[89] J. P. MacManus, *Accessibility, planar graphs, and quasi-isometries*, arXiv preprint arXiv:2310.15242, (2023).

[90] ———, *Deciding if a hyperbolic group splits over a given quasiconvex subgroup*, Groups, Geometry, and Dynamics, (2024).

[91] ———, *A note on quasi-transitive graphs quasi-isometric to planar (Cayley) graphs*, arXiv preprint arXiv:2407.13375, (2024).

[92] J. P. MacManus and L. Mineh, *Tiling in some nonpositively curved groups*, arXiv preprint arXiv:2401.09545, (2024).

[93] W. Magnus, *Über diskontinuierliche Gruppen mit einer definierenden Relation. (Der Freiheitssatz).*, Journal für die reine und angewandte Mathematik, 163 (1930), pp. 141–165.

[94] S. Maillot, *Quasi-isometries of groups, graphs and surfaces*, Commentarii Mathematici Helvetici, 76 (2001), pp. 29–60.

[95] H. Maschke, *The representation of finite groups, especially of the rotation groups of the regular bodies of three-and four-dimensional space, by Cayley's color diagrams*, American Journal of Mathematics, 18 (1896), pp. 156–194.

[96] B. Maskit, *A theorem on planar covering surfaces with applications to 3-manifolds*, Annals of Mathematics, (1965), pp. 341–355.

[97] H. A. Masur and Y. N. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Inventiones mathematicae, 138 (1999), pp. 103–149.

[98] K. Menger, *Zur allgemeinen kurventheorie*, Fundamenta Mathematicae, 10 (1927), pp. 96–115.

[99] G. Mess, *The Seifert conjecture and groups which are coarse quasiisometric to planes*, preprint, (1988).

[100] A. Minasyan and D. Osin, *Acylindrical hyperbolicity of groups acting on trees*, Math. Ann., 362 (2015), pp. 1055–1105.

[101] R. G. Möller, *Ends of graphs*, in Math. Proc. Camb. Phil. Soc, vol. 111, 1992, pp. 255–266.

[102] H. M. Morse, *A fundamental class of geodesics on any closed surface of genus greater than one*, Transactions of the American Mathematical Society, 26 (1924), pp. 25–60.

[103] B. H. Neumann, *On the number of generators of a free product*, Journal of the London Mathematical Society, 1 (1943), pp. 12–20.

[104] G. Niblo, *The singularity obstruction for group splittings*, Topology and its Applications, 119 (2002), pp. 17–31.

[105] G. Niblo, M. Sageev, P. Scott, and G. A. Swarup, *Minimal cubings*, International Journal of Algebra and Computation, 15 (2005), pp. 343–366.

[106] P. S. Novikov, *On the algorithmic unsolvability of the word problem in group theory*, Trudy Matematicheskogo Instituta imeni VA Steklova, 44 (1955), pp. 3–143.

[107] D. Osin, *Acylindrically hyperbolic groups*, Trans. Amer. Math. Soc., 368 (2016), pp. 851–888.

[108] ———, *Groups acting acylindrically on hyperbolic spaces*, in Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. II. Invited lectures, World Sci. Publ., Hackensack, NJ, 2018, pp. 919–939.

[109] P. Papasoglu, *An algorithm detecting hyperbolicity*, Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994), 25 (1996), pp. 193–200.

[110] ———, *Quasi-isometry invariance of group splittings*, Annals of mathematics, (2005), pp. 759–830.

[111] P. Papasoglu and K. Whyte, *Quasi-isometries between groups with infinitely many ends*, Commentarii Mathematici Helvetici, 77 (2002), pp. 133–144.

[112] H. Poincaré, *Analysis situs*, Gauthier-Villars Paris, France, 1895.

[113] M. O. Rabin, *Recursive unsolvability of group theoretic problems*, Annals of Mathematics, (1958), pp. 172–194.

[114] F. Raymond, *The end point compactification of manifolds.*, Pacific J. Math., 10 (1960), pp. 947–963.

[115] R. Richter and C. Thomassen, *3-connected planar spaces uniquely embed in the sphere*, Transactions of the American Mathematical Society, 354 (2002), pp. 4585–4595.

[116] O. Schreier, *Die Untergruppen der freien Gruppen*, in Abhandlungen aus dem Mathematischen Seminar der universität Hamburg, vol. 5, Springer, 1927, pp. 161–183.

[117] P. Scott, *Ends of pairs of groups*, Journal of Pure and Applied Algebra, 11 (1977), pp. 179–198.

[118] ———, *The symmetry of intersection numbers in group theory*, Geometry & Topology, 2 (1998), pp. 11–29.

[119] P. Scott and G. A. Swarup, *Splittings of groups and intersection numbers*, Geometry & Topology, 4 (2000), pp. 179–218.

[120] ———, *Regular neighbourhoods and canonical decompositions for groups*, Société mathématique de France, 2003.

[121] Z. Sela, *The isomorphism problem for hyperbolic groups I*, Annals of Mathematics, (1995), pp. 217–283.

[122] ———, *Acylindrical accessibility for groups*, Inventiones mathematicae, 129 (1997), pp. 527–565.

[123] J.-P. Serre, *Trees*, Springer Science & Business Media, 2002.

[124] B. Seward, *Burnside's Problem, spanning trees and tilings*, Geom. Topol., 18 (2014), pp. 179–210.

[125] J. R. Stallings, *On torsion-free groups with infinitely many ends*, Annals of Mathematics, (1968), pp. 312–334.

[126] ———, *Group theory and three-dimensional manifolds*, Yale Mathematical Monographs, 4, (1971).

[127] G. A. Swarup, *A note on accessibility*, Geometric Group Theory I (Sussex, 1991), 181 (1993), pp. 204–207.

[128] ———, *On the cut point conjecture*, Electronic Research Announcements of the American Mathematical Society, 2 (1996), pp. 98–100.

[129] C. Thomassen and W. Woess, *Vertex-transitive graphs and accessibility*, Journal of Combinatorial Theory, Series B, 58 (1993), pp. 248–268.

[130] H. Tietze, *On the topological invariants of multidimensional manifolds*, Monatshefte fur Mathematik und Physik, 19 (1908), pp. 1–118.

[131] N. Touikan, *Detecting geometric splittings in finitely presented groups*, Transactions of the American Mathematical Society, 370 (2018), pp. 5635–5704.

[132] P. Tukia, *Homeomorphic conjugates of Fuchsian groups.*, Journal für die reine und angewandte Mathematik, 390 (1988), pp. 1–54.

[133] N. Vaskou, *Acylindrical hyperbolicity for Artin groups of dimension 2*, Geom. Dedicata, 216 (2022), pp. Paper No. 7, 28.

[134] A. Vonseel, *Ends of Schreier graphs of hyperbolic groups*, Algebraic & Geometric Topology, 18 (2018), pp. 3089–3118.

[135] C. Wall, *Pairs of relative cohomological dimension one*, Journal of Pure and Applied Algebra, 1 (1971), pp. 141–154.

[136] B. Weiss, *Monotileable amenable groups*, in Topology, ergodic theory, real algebraic geometry, vol. 202 of Amer. Math. Soc. Transl. Ser. 2, Amer. Math. Soc., Providence, RI, 2001, pp. 257–262.

[137] H. C. Wilkie, *On non-Euclidean crystallographic groups*, Mathematische Zeitschrift, 91 (1966), pp. 87–102.

[138] W. Woess, *Topological groups and infinite graphs*, Discrete mathematics, 95 (1991), pp. 373–384.

[139] H. Zieschang, E. Vogt, and H.-D. Coldewey, *Flächen und ebene diskontinuierliche Gruppen*, Springer, 1970.

Index