ALGORITHMIC AND STATISTICAL PROPERTIES OF FILLING ELEMENTS OF A FREE GROUP, AND QUANTITATIVE RESIDUAL PROPERTIES OF Γ-LIMIT GROUPS

BY

BRENT BRADFORD SOLIE

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2011

Urbana, Illinois

Doctoral Committee:

Associate Professor Christopher J. Leininger, Chair
Associate Professor Ilya Kapovich, Director of Research
Associate Professor Igor Mineyev
Professor Emeritus Derek J. S. Robinson

Abstract

A *filling subgroup* of a finitely generated free group $F(X)$ is a subgroup which does not fix a point in any very small action free action on an $\mathbb{R}$-tree. For the free group of rank two, we construct a combinatorial algorithm to determine whether or not a given finitely generated subgroup is filling. In higher ranks, we discuss two types of non-filling subgroups: those contained in loop vertex subgroups and those contained in segment vertex subgroups. We construct a combinatorial algorithm to determine whether or not a given finitely generated subgroup is contained in a segment vertex subgroup. We further give a combinatorial algorithm which identifies a certain kind of subgroup contained in a loop vertex subgroup. Finally, we show that the set of filling elements of $F(X)$ is exponentially generic in the sense of Arzhantseva-Ol'shanskiĭ, refining a result of Kapovich and Lustig.

Let Γ be a fixed hyperbolic group. The Γ*-limit groups* of Sela are exactly the finitely generated, fully residually Γ groups. We give a new invariant of Γ-limit groups called Γ-discriminating complexity and show that the Γ-discriminating complexity of any Γ-limit group is asymptotically dominated by a polynomial. Our proof relies on an embedding theorem of Kharlampovich-Myasnikov which states that a Γ-limit group embeds in an iterated extension of centralizers over Γ. The result then follows from our proof that if G is an iterated extension of centralizers over Γ, the G-discriminating complexity of a rank n extension of a cyclic centralizer of G is asymptotically dominated by a polynomial of degree n.

Table of Contents

List of Figures iv

Chapter 1 Introduction 1
1.1 Filling Elements and Filling Subgroups of Free Groups 1
1.2 Residual Properties of Γ-Limit Groups 5

Chapter 2 Filling Elements and Filling Subgroups of Free Groups 8
2.1 Background 8
2.1.1 Stallings Graphs 8
2.1.2 Whitehead's Algorithm 10
2.1.3 Outer space 14
2.2 Main Results 16
2.2.1 Algorithmic Properties of Filling Subgroups 16
2.2.1.1 The Rank Two Case 18
2.2.1.2 Segment vertex subgroups in higher rank 20
2.2.1.3 Loop vertex subgroups in higher rank 26
2.2.2 Genericity of Filling Elements 28
2.2.2.1 The set TS' 28
2.2.2.2 Filling elements 29

Chapter 3 Residual Properties of Limit Groups 31
3.1 Background 31
3.1.1 Γ-limit groups 31
3.1.2 Relative hyperbolicity 33
3.1.3 Iterated extensions of centralizers over Γ are relatively hyperbolic 34
3.1.4 Relative hyperbolic geometry 35
3.2 Main Results 37
3.2.1 Relative hyperbolic geometry 37
3.2.2 Discriminating complexity 46
3.2.2.1 Free Abelian groups 48
3.2.2.2 Extensions of centralizers 50
3.2.2.3 Iterated extensions of centralizers 52
3.2.2.4 Arbitrary Γ-limit groups 52

References 54

Index 57

List of Figures

2.1 The local effect of $\phi = (C, m)$ on an X-digraph 12
2.2 Stallings graphs of standard vertex subgroups. 16
2.3 Graphs associated to the subgroup $\langle a, a^b \rangle$. 19
2.4 Stallings graph of a standard segment vertex subgroup. 21
2.5 If $(\Delta \cup m \cup \Pi, m)$ reduces S with $m \in A^{\pm}$, then so must $(\Delta \cup m, m)$. 23
2.6 If (C, m) reduces S with $m \in B^{\pm}$, then either $(C \cap B^{\pm}, m)$ or $(\Delta \cup b_r, b_r)$ also reduces S. 24
2.7 A Stallings graph satisfying Property (L). 27

3.1 Producing the relation $w^{-1} u^a w = u^b$ in the proof of Lemma 3.28 38
3.2 Finding a shorter coset representative in Lemma 3.30 39
3.3 A relative geodesic triangle. 40
3.4 A decomposition of Q and one of its sides. 41
3.5 The decomposition of α. 43
3.6 β_i and β_{i+1} cannot fellow travel too far without causing γ_i and γ_{i+1} to fellow travel. 44

Chapter 1

Introduction

1.1 Filling Elements and Filling Subgroups of Free Groups

The action of a group on a metric space is one of the fundamental tools of geometric group theory. When such an action is sufficiently well-behaved, it can reveal many different properties of the group which may be otherwise difficult to discover. These geometric methods are integral to the modern understanding of free groups, surface groups, and their corresponding automorphism groups.

An $\mathbb{R}$-*tree* is a geodesic metric space in which every pair of points is connected by a unique injective path. Free groups and surface groups admit many actions on $\mathbb{R}$-trees, and the study of these actions is a central concern in modern group theory. Such important spaces as the compactification of Culler-Vogtmann outer space (in the case of free groups) and the compactification of Teichmüller space (in the case of surface groups) have important characterizations in terms of actions on trees.

Our results in Chapter 2 are concerned with the algorithmic and statistical properties of actions of a free group on $\mathbb{R}$-trees. The central notion here is that of a filling element, an element which has a non-degenerate action in every sufficiently nice action on an $\mathbb{R}$-tree. Filling elements are the free group analogue of filling curves in a closed, orientable, hyperbolic surface. Filling curves have played an important role in the theory of surface groups, which we will briefly review here.

Let Σ be a closed, orientable surface of genus at least two. By a *surface group* we mean the fundamental group of such a surface Σ. Let α and β be closed curves on Σ. The *geometric intersection number*, denoted $i(\alpha, \beta)$, is the least number of intersections between members of the free homotopy classes of α and β. If β is such that $i(\alpha, \beta) > 0$ for every essential simple closed curve α, then we say that β is a *filling curve* .

Recall that the *dual tree* T_α associated to an essential simple closed curve α on Σ is a simplicial tree equipped with a small minimal isometric action by $\pi_1(\Sigma)$. (From now on, we will assume all our surface group actions on trees are minimal and isometric.) It is well-known that if β is a (not necessarily simple) closed curve on Σ, then the *translation length* of β on T_α, denoted $||\beta||_{T_\alpha}$, is equal to $i(\alpha, \beta)$. Therefore, a closed curve β is filling if and only if it has positive translation length on T_α for every essential simple closed curve α.

As a consequence of Skora's duality theorem, any simplicial tree equipped with a small action by $\pi_1(\Sigma)$ can be collapsed down into a tree T_α for some essential simple closed curve α. Therefore, a closed curve β is filling if and only if it has positive translation length in every small action of $\pi_1(\Sigma)$ on a simplicial tree. An application of Bass-Serre theory shows that a closed curve β is filling if and only if it is not conjugate into a vertex subgroup in any elementary cyclic splitting of $\pi_1(\Sigma)$.

We now move from surface groups to consider a finitely generated non-Abelian free group $F(X)$. The Culler-Vogtmann outer space, denoted $CV(F(X))$, is the projectivized space of free discrete actions of $F(X)$ on simplicial trees. Outer space is the free group counterpart to Teichmüller space in the sense that it admits a properly discontinuous isometric action by the outer automorphism group of $F(X)$ [16]. Moreover, outer space also admits a Thurston-type compactification $\overline{CV}(F(X))$, the projectivized space of very small actions of $F(X)$ on $\mathbb{R}$-trees [5].

In [25], Kapovich and Lustig introduce the notion of a filling element as a free group analogue for a filling curve. A *filling element* is an element $w \in F(X)$ that has positive translation length in every very small action of $F(X)$ on an $\mathbb{R}$-tree. The cyclic subgroup generated by a filling element is an instance of a finitely generated *filling subgroup* of $F(X)$: a subgroup which does not fix a point in any very small action of $F(X)$ on an $\mathbb{R}$-tree.

In the same paper in which they introduce filling elements, Kapovich and Lustig prove the following theorem, which serves as the inspiration for our investigation of the filling property:

Proposition 1.1 ([25, Theorem 13.6])**.** *With respect to the uniform measure on $\partial F(X)$, for almost every infinite geodesic ray in the Cayley graph of $F(X)$, every sufficiently long initial segment of that ray represents a filling element of $F(X)$.*

The proof of Kapovich and Lustig's theorem is non-constructive, so while it is a strong indication that filling elements should be common in a free group, it cannot be used to show that the filling elements are common in any formal sense, nor does it provide a method for identifying such an element. Furthermore, their result does not address the more general concept of a filling subgroup.

The first part of Chapter 2 is dedicated to finding a partial solution to the following decision problem:

> Let $F(X)$ be a finitely generated, non-Abelian free group. Given a finitely generated subgroup H of $F(X)$, is H a filling subgroup?

Decision problems such as these have long played an important role in group theory. Dehn's three major group-theoretic decision problems, the word, conjugacy, and isomorphism problems, have been studied extensively and have lead to such important concepts as the small cancellation conditions and word-hyperbolicity.

Our first main result is:

Theorem A (c.f. Theorem 2.26). *Let $F(a,b)$ denote the free group of rank two. There is an algorithm to determine, given a finitely generated subgroup H of $F(a,b)$, whether or not H is a filling subgroup.*

In higher rank cases, we have two different types of non-filling subgroup: those which are elliptic in a cyclic segment splitting of $F(X)$ and those which are elliptic in a cyclic loop splitting of $F(X)$ (see Definitions 2.14 and 2.17.)

A vertex subgroup in a cyclic segment splitting of $F(X)$ has a highly structured Stallings graph. Specifically, up to automorphism, the Stallings graph of such a subgroup consists of a bouquet of circles labeled by the elements of some proper subset of X plus a loop labeled by the remaining elements of X. This structure is encoded by the combinatorial Property (S) (Definition 2.27). Immersions onto a graph with Property (S) are preserved by the Whitehead minimization process (Proposition 2.28), a consequence of which is our second main result:

Theorem B (c.f. Theorem 2.34). *Let $F(X)$ be a free group of finite rank at least three. There is an algorithm to determine, given a finitely generated subgroup H in $F(X)$, whether or not H is elliptic in a cyclic segment splitting of $F(X)$.*

As in the segment case, a vertex subgroup in a cyclic loop splitting of $F(X)$ also has a very rigid combinatorial structure which can be characterized in terms of Stallings graphs. Briefly, a Stallings graph has Property (L) if it has an edge with a unique label which separates the graph into two components, at least one of which is rank one (see Definition 2.35.) Up to automorphism, a subgroup of a loop vertex subgroup has a Stallings graph which admits an immersion onto a graph with Property (L). However, unlike the case with Property (S), the Whitehead minimization process does not preserve immersions onto graphs with Property (L), so the previous technique cannot algorithmically detect whether or not a subgroup is contained in a loop vertex subgroup. However, Property (L) itself can be detected up to automorphism.

Theorem C (c.f. Theorem 2.37). *Let $F(X)$ be a free group of finite rank at least three. There is an algorithm to determine, given a finitely generated subgroup H of $F(X)$, whether or not there exists $\phi \in \operatorname{Aut} F(X)$ such that the Stallings graph of $\phi(H)$ satisfies Property (L).*

Another aspect of filling elements we would like to address is their statistical properties. Kapovich and Lustig's theorem indicates that filling elements should be fairly common in a free group. However, the non-constructive nature of the proof does not allow us to formalize the sense in which filling elements are common.

The second part of Chapter 2 is dedicated to investigating whether or not the set of filling elements of $F(X)$ is generic in the following sense of Arzhantseva and Ol'shanskiĭ. Let S be a subset of elements of $F(X)$. We say that S is $F(X)$-*generic* if

$$\lim_{R\to\infty} \frac{\#(S \cap B_R)}{\#B_R} = 1,$$

where B_R is the set of elements of $F(X)$ with X-length at most R [29]. If the limit converges exponentially fast, we say that S is *exponentially* $F(X)$-*generic.* (We will give a slightly more general definition of genericity in Definition 2.39.)

Historically, the earliest appearance of the notion of genericity seems to be due to Guba [21]. Shortly afterwards, Gromov gave a formal definition in the context of finitely presented groups [20]. In the same paper, Gromov asserts that almost every finitely presented group is hyperbolic, a fact first proved by Ol'shanskiĭ [39] and later by Champetier [14, 15]. Subsequent results in statistical group theory include the work of Arzhantseva [1, 2], Arzhantseva and Ol'shanskiĭ [3], and Ollivier [35, 34, 36, 37]. The surveys by Ghys [19] and Ollivier [38] provide an excellent overview of genericity with a focus on random groups. More recent results in statistical group theory apply the notion of genericity to computational group theory. Some group-theoretic decision problems with high worst-case complexity have been shown to have low complexity on a generic set of inputs [28, 29]. These results have furthered the understanding of the average-case complexity of these problems [26, 27].

Our final main result on the filling property is:

Theorem D (c.f. Theorem 2.45)**.** *Let $F(X)$ be a free group of finite rank at least two.*

1. *Let $w \in F(X)$. If the stabilizer of w in $\operatorname{Aut} F(X)$ is infinite cyclic, then w is filling.*
2. *The set of filling elements of $F(X)$ is exponentially $F(X)$-generic.*
3. *There exists an exponentially $F(X)$-generic subset S of $F(X)$ such that every element of S is filling and the membership problem for S is solvable in linear time.*

This result recalls the result of Bonahon that filling is the typical behavior of closed curves on a closed orientable hyperbolic surface [7, 8, 9]. Genericity of the filling property is therefore another example of the symmetry between free and surface groups.

We briefly note the following application of Theorem D to the work of Reynolds [43]. An injective endomorphism $\phi : F(X) \to F(X)$ is *admissible* if $\phi(F(X))$ is a filling subgroup. Reynolds shows that an admissible injective endomorphism of $F(X)$ acts on $\overline{CV}(F(X))$ with a single attracting fixed point.

As a corollary to Theorem D, we see that admissibility is the typical behavior of injective endormorphisms. Specifically, let N be the cardinality of X and consider the set of N-tuples $(w_1, \ldots, w_N)$ of elements of $F(X)$. Let B_R^N denote the set of such tuples with $|w_i|_X \leq R$ for each $i = 1, \ldots, N$. We may extend the notion of genericity to subsets of $F(X)^N$: a subset $S \subseteq F(X)^N$ is generic if

$$\lim_{R \to \infty} \frac{\#(S \cap B_R^N)}{\#B_R^N} = 1,$$

and exponentially generic if the above limit converges exponentially fast. Since the set of filling elements is exponentially generic in $F(X)$, the set of tuples of filling elements is exponentially generic in $F(X)^N$. After restricting to the set of tuples representing injective endomorphisms, we obtain an exponentially generic subset of admissible injective endomorphisms.

1.2 Residual Properties of Γ-Limit Groups

Quantitative analysis of group properties is an increasingly active field in modern group theory. In particular, the various residual properties of groups have proven themselves quite suitable for investigation through quantitative means.

Let P be a property of groups, and recall that a group G is *residually* P if for every nontrivial element $g \in G$, there is a homomorphism $\phi : G \to H$ such that H is a group with property P and $\phi(g) \neq 1$. We say that a group is *fully residually* P if for every finite subset of nontrivial elements $S \subseteq G - 1$, there is a homomorphism $\phi : G \to H$ such that H is a group with property P and $1 \notin \phi(S)$.

(An alternate definition of fully residually P insists that the homomorphism ϕ not just avoid 1 but actually be injective on S. Note that ϕ is injective on S if and only if the image under ϕ of the set $\{uv^{-1} : u, v \in S, u \neq v\}$ does not include 1, so these definitions are equivalent. Also note that we also do not require our homomorphisms to be surjective, as may sometimes be the case when discussing residual properties.)

For instance, let G be a residually finite group with finite generating set X. Let $f : \mathbb{N} \to \mathbb{N}$ be such that whenever $g \in G - 1$ has X-length at most R, then there exists $\phi : G \to H$ such that $\phi(g) \neq 1$ and $|H| \leq f(R)$. When f is the smallest such function, then we think of f as measuring the complexity of the residual finiteness of G; we may also think of f as measuring the growth of the number of subgroups of G with respect to index. This version of complexity has been studied extensively by Bou-Rabee in [11], with additional results by Kassabov and Matucci [30].

Bou-Rabee has obtained further results by restricting his attention to finite nilpotent or finite solvable quotients. This yields group invariants known as the *nilpotent Farb growth* and the *solvable Farb growth*, and Bou-Rabee has obtained new characterizations of algebraic group properties in terms of the asymptotic properties of these growth functions. For instance, Bou-Rabee has shown that a finitely generated group G is nilpotent if and only if it has nilpotent Farb growth which is polynomial in $\log(n)$ [11]. Similarly, a finitely generated group is solvable and virtually nilpotent if and only if it has solvable Farb growth that is polynomial in $\log(n)$ [12].

Rather than considering residually finite groups, we will study another well-known class of groups with strong residual properties: the Γ-limit groups of Sela. Let Γ be a fixed torsion-free hyperbolic group. A Γ-*limit group* G is a finitely generated, fully residually Γ group: for any finite subset $S \subseteq G - 1$, there exists a homomorphism $\phi : G \to \Gamma$ such that $1 \notin \phi(S)$. We say that the set S is Γ-*discriminated* by ϕ.

Fix finite generating sets X and Y for G and Γ, respectively. Let the homomorphism $\phi_R : G \to \Gamma$ discriminate $B_R(G,X)-1$, where $B_R(G,X)$ is the closed ball of radius R in G with respect to X. Here, we measure the complexity of ϕ_R by the maximum Y-length over all images of elements of X. The minimum complexity required to discriminate each set $B_R(G,X)-1$, as a function of R, is called the Γ-discriminating complexity of G, and it is an invariant of G up to asymptotic equivalence. (See Definition 3.45.)

Our main result on the Γ-discriminating complexity of Γ-limit groups is the following:

Theorem E (c.f. Theorem 3.57)**.** *The* Γ-*discriminating complexity of a* Γ-*limit group is asymptotically dominated by a polynomial.*

In order to prove Theorem E, we must first start with the simplest examples of Γ-limit groups: the finitely generated, free Abelian groups. The free Abelian group $\mathbb{Z}^n$ is fully residually $\mathbb{Z}$, and our next main result establishes its $\mathbb{Z}$-discriminating complexity.

Theorem F (c.f. Theorem 3.53)**.** *The* $\mathbb{Z}$-*discriminating complexity of* $\mathbb{Z}^n$ *is asymptotically equivalent to a polynomial of rank* $n-1$.

The fundamental construction in our study of Γ-limit groups is the extension of a centralizer. Informally, if G is a Γ-limit group, we may construct another Γ-limit group G' by extending a centralizer of G by a free Abelian group of finite rank. (See Definition 3.4.)

Our next result is motivated by the well-known "big powers" property of hyperbolic groups. If Γ is a hyperbolic group and $u \in \Gamma$ generates its own centralizer, then for any tuple of elements $(g_1, g_2, \ldots, g_k)$ of elements of $G - \langle u \rangle$, there is an integer N such that

$$u^{n_0} g_1 u^{n_1} g_2 u^{n_2} \ldots u^{n_{k-1}} g_k u^{n_k}$$

is nontrivial in Γ whenever $|n_i| > N$ for $i = 1, \ldots, k-1$ and either $|n_i| > N$ or $n_i = 0$ for $i = 0, k$.

The big powers property seems to appear first due to B. Baumslag in his study of fully residually free groups [4]; a later version appears due to Ol'shanskiĭ in the context of hyperbolic groups [40]. Most recently, the big powers property is proven by Kharlampovich and Myasnikov for relatively hyperbolic groups in [31] using the techniques of Osin from [41, 42]. Our main technical lemma, Lemma 3.42, is an analysis of the big powers property for relatively hyperbolic groups with the goal of analyzing the dependence of N on the group G, generating set X, and the elements g_i and u.

By iterating the extension of centralizer construction, we obtain a group known as an *iterated extension of centralizers* (see Definition 3.7). Iterated extensions of centralizers are relatively hyperbolic and therefore have the big powers property. By combining Theorem F with our analysis of the big powers property, we obtain our third main result.

Theorem G (c.f. Theorem 3.55)**.** *Let G be an iterated extension of centralizers over Γ. Let G' be a rank n extension of a cyclic centralizer of G. Then the G-discriminating complexity of G' is asymptotically dominated by a polynomial of degree n.*

Repeated application of Theorem G gives us our final main result, a bound on the discriminating complexity of an arbitrary iterated extension of centralizers over Γ.

Theorem H (c.f. Theorem 3.56)**.** *The Γ-discriminating complexity of an iterated extension of centralizers over Γ is asymptotically dominated by a polynomial with degree equal to the product of the ranks of the extensions.*

Theorem H then directly implies Theorem E via a theorem of Kharlampovich and Myasnikov, which states that every Γ-limit group embeds in some iterated extension of centralizers over Γ [31].

Chapter 2

Filling Elements and Filling Subgroups of Free Groups

2.1 Background

Let X be a finite set with at least two elements. Define $X^{-1} := \{x^{-1} : x \in X\}$ to be the set of *formal inverses* of elements of X, and set $X^{\pm} := X \sqcup X^{-1}$. We denote the set of words on the letters $X^{\pm}$ by $(X^{\pm})^*$. A word in $(X^{\pm})^*$ is *freely reduced* if it has no subword of the form xx^{-1} or $x^{-1}x$ for any $x \in X$. A word in $(X^{\pm})^*$ is *cyclically reduced* if every cyclic permutation of that word is freely reduced.

Let $F(X)$ be the free group on the letters X. The X*-length* of $w \in F(X)$, denoted $|w|_X$, is the length of the freely reduced word in $(X^{\pm})^*$ which represents w. We will indicate that H is a finitely generated subgroup of $F(X)$ by $H \leq_{fg} F(X)$.

2.1.1 Stallings Graphs

Definition 2.1 (X-digraph)**.** Given a finite set X, an X*-digraph* is given by the data $(V, E, \cdot_+, \cdot_-, \lambda)$, where:

- V and E are sets;
- $\cdot_+, \cdot_- : E \to V$; and
- $\lambda : E \to X$.

We call V the *vertex set* and E the *edge set*. For $e \in E$, we say that e_- is the *initial vertex* of e and e_+ is the *terminal vertex* of e. We call *lambda* the labeling function.

Let S be an X-digraph. By VS and ES we denote the vertex and edge sets of S, respectively. For $v \in VS$, we define the *in-link* of v to be $\mathrm{lk}_+(v) := \{e \in ES : e_+ = v\}$, and we say the *in-hyperlink* of v is $\mathrm{hl}_+(v) := \{\lambda(e) : e \in \mathrm{lk}_+(v)\}$. Likewise, we define the *out-link* of v to be $\mathrm{lk}_-(v) := \{e \in ES : e_- = v\}$ and the *out-hyperlink* of v to be $\mathrm{hl}_-(v) := \{\lambda(e)^{-1} : e \in \mathrm{lk}_-(v)\}$. The *link* of v is $\mathrm{lk}(v) := \mathrm{lk}_-(v) \cup \mathrm{lk}_+(v)$ and the *hyperlink* of v is $\mathrm{hl}(v) := \mathrm{hl}_-(v) \cup \mathrm{hl}_+(v)$.

We say that the *degree* of $v \in VS$ is $\deg_S(v) := \#\,\mathrm{lk}(v)$. If $\deg_S(v) = 0$ we say that v is *isolated*, and if $\deg_S(v) = 1$ we say that v is a *leaf*. If no vertex of S is a leaf, we say that S is *cyclically reduced*.

If $u, v \in VS$ are such that there is $e \in ES$ with $e_- = v$ and $e_+ = u$, then we say that u and v are *adjacent* and that e is *incident* to both u and v. If $e, f \in ES$ are such that $e_- = f_-$, then we say that e and f are *coinitial*; if $e_+ = f_+$, then e and f are *coterminal*. The edges e and f are *coincident* if e shares an endpoint with f.

Let $Y \subseteq X$. A Y*-edge* is any edge with label in Y. The set of Y-edges of S is denoted $E_Y S$.

Convention. To aid readability, we will always denote singleton sets by their unique element. For instance, if $x \in X$, we will write x-edges rather than $\{x\}$-edges, and the set of x-edges of S will be denoted $E_x S$ rather than $E_{\{x\}} S$.

We say that S is *folded at* v if λ induces a bijection $\mathrm{lk}(v) \to \mathrm{hl}(v)$. The graph S is *folded* if S is folded at every vertex. If S is not folded, then some pair of coterminal or coinitial edges $e, f \in ES$ share the same label. We *fold* these edges by identifying the pair of edge e and f and either the vertices e_- and f_- if e and f are coinitial or the vertices e_+ and f_+ if e and f are coterminal. The process of performing folds in S until none remain is called *folding*.

We may *delete* a leaf v of S by deleting v and the unique edge incident to it. By repeatedly deleting leaves, we eventually arrive at a cyclically reduced X-digraph. We call this process *cyclic reduction*.

Lastly, if S and T are X-digraphs, an X*-map* is a map $S \to T$ which sends vertices to vertices, edges to edges, and preserves both orientation and label. We will assume that all of our maps of X-digraphs are X-maps. An X-map is an *immersion* if the induced map on the link of a vertex is injective for every vertex of S.

Definition 2.2 (Dual digraph)**.** Given an X-digraph S, we may construct the *dual* of S, denoted S^*, by adding a set of *formal inverse edges* $\overline{E}S := \{\bar{e} : e \in ES\}$ and extending $\cdot_-, \cdot_+$, and λ as follows:

$$
\begin{aligned}
(\bar{e})_- &:= e_+, \\
(\bar{e})_+ &:= e_-, \text{ and} \\
\lambda(\bar{e}) &:= \lambda(e)^{-1}.
\end{aligned}
$$

By defining $\overline{(\bar{e})} = e$, the above equations are satisfied for any $e \in ES^*$.

A *path* p in an X-digraph S is a sequence of edges $p = e_1 e_2 \dots e_l$ in the dual S^* such that $(e_i)_+ = (e_{i+1})_-$ for all $i = 1, \dots, l-1$. The path p is a *loop* if we further have that $(e_l)_+ = (e_1)_-$. The path is *immersed* if $e_{i+1} \neq \bar{e}_i$ for all $i = 1, \dots, l-1$. The *label* of p is $\lambda(p) := \lambda(e_1) \dots \lambda(e_l)$. The *length* of p is l. We say that an X-digraph S is *connected* if there exists a path between any two vertices.

Remark. For the purposes of this chapter, we will not distinguish much between an X-digraph S and its dual S^*. Specifically, for $x \in X$, we will regard an x^{-1}-edge as simply an x-edge with the opposite orientation. If e is an x-edge, we will therefore consider it as an x-edge in that it contributes the label x to the hyperlink of e_+, and also as an x^{-1} edge as it contributes the label x^{-1} to the hyperlink of e_-.

Definition 2.3 (Stallings graph)**.** Let $H \leq_{fg} F(X)$. The *Stallings graph* representing H with respect to X, denoted $S_X(H)$, is the unique X-digraph with basepoint such that a freely reduced word in $(X^{\pm})^*$ represents an element of H if and only if it occurs as the label of an immersed loop of $S_X(H)$ beginning and ending at the basepoint.

Recall that we may construct $S_X(H)$ as follows. Let $h_1, \dots, h_k$ be elements of $(X^{\pm})^*$ representing a finite set of generators of H. Beginning with a basepoint, denoted 1, we construct a loop beginning and ending at 1 with label h_i for each i; let S_0 be the resulting graph. We then perform all possible folds in S_0 (in any order) to obtain a folded graph S_1. Finally, we repeatedly delete leaves of S_1 different from 1 until no leaves remain except possibly the basepoint. The resulting graph is $S_X(H)$, and it is well-known that $S_X(H)$ is invariant with respect to the choice of generating set for H as well as the order of the folds and leaf deletions.

Suppose that $S_X(H)$ is cyclically reduced. For any $g \in F(X)$, we may construct $S_X(H^g)$ from $S_X(H)$ by adding a new basepoint $1'$, a path from 1 to $1'$ labeled by g, and then folding and deleting non-basepoint leaves. Therefore whenever $S_X(H)$ is cyclically reduced, we will forget the basepoint and think of $S_X(H)$ as representing $H^{\operatorname{Aut} F(X)}$, the conjugacy class of H in $F(X)$, rather than the single subgroup H.

2.1.2 Whitehead's Algorithm

Definition 2.4 (Whitehead automorphism)**.** A *type I Whitehead automorphism* is an automorphism $\phi \in \operatorname{Aut} F(X)$ which is induced by permutations and inversions of the set $X^{\pm}$.

A *type II Whitehead automorphism* is an automorphism $\phi \in \operatorname{Aut} F(X)$ for which there exists $m \in X^{\pm}$ such that $\phi(m) = m$ and

$$\phi(x) \in \{x, m^{-1}x, xm, x^m\}$$

for all $x \in X$. We call m the *multiplier* for ϕ.

Given a type II Whitehead automorphism ϕ with multiplier m, define

$$C := \left\{x \in X^{\pm} : \phi(x) \in \{m, xm, x^m\}\right\}.$$

Then ϕ is determined completely by the pair (C, m), and we refer to C as the *cut* for ϕ.

Let $C \subseteq X^{\pm}$ be such that $m \in C$ and $m^{-1} \notin C$. We call such a C an *m-cut*. For any $m \in X^{\pm}$ and m-cut C, the pair (C, m) defines a type II Whitehead automorphism of $F(X)$.

More generally, if $C, D \subseteq X^{\pm}$, we say that C *cuts* D if D contains an element of both C and $C' := X^{\pm} - C$.

Definition 2.5 (Hypergraph)**.** A *hypergraph* is a tuple (V, E, ι), where V and E are sets and $\iota : E \to \mathcal{P}(V)$, where $\mathcal{P}(V)$ denotes the power set of V. The elements of V are called *vertices* and the elements of E are called *hyperedges*. We call ι the *incidence function*.

Let Γ be a hypergraph. We refer to the vertex and hyperedge sets of Γ by $V\Gamma$ and $E\Gamma$, respectively. We will refer to the incidence function by simply ι when Γ is clear from context. We say that a hyperedge $e \in E\Gamma$ is *incident* to a vertex $v \in V\Gamma$ if $v \in \iota(e)$. A pair of hyperedges $e, e' \in E\Gamma$ are *coincident* if $\iota(e) \cap \iota(e') \neq \emptyset$. Two vertices $v, v' \in V\Gamma$ are *adjacent* if there is a hyperedge $e \in E\Gamma$ with $v, v' \in \iota(e)$.

More generally, if $Y \subset V\Gamma$, we say that a hyperedge $e \in E\Gamma$ is *incident* to Y if $\iota(e) \cap Y \neq \emptyset$. Let $Y_1, \ldots, Y_n, Z$ be subsets of $V\Gamma$. We say that a hyperedge $e \in E\Gamma$ has *type* $(Y_1, Y_2, \ldots, Y_n; Z)$ if e is incident to each Y_i for $i = 1, \ldots, n$ but e is not incident to Z. When Z is empty, we will write $(Y_1, Y_2, \ldots, Y_n)$ instead of $(Y_1, Y_2, \ldots, Y_n; \emptyset)$. We denote by $[Y_1, Y_2, \ldots, Y_n; Z]_\Gamma$ the number of hyperedges of Γ of type $(Y_1, Y_2, \ldots, Y_n; Z)$.

Let $Y \subseteq V\Gamma$, and let Y' denote the complement $V\Gamma - Y$. We define the *capacity* of Y in Γ to be the number of hyperedges of Γ incident to both Y and its complement; in the above notation,

$$\mathrm{cap}_\Gamma(Y) = [Y, Y']_\Gamma.$$

Let $v \in V\Gamma$. The *degree* of v in Γ is the number of edges incident to v; in the above notation,

$$\deg_\Gamma(v) = [v]_\Gamma.$$

Definition 2.6 (Whitehead hypergraph)**.** Let S be a cyclically reduced X-digraph. We define the *Whitehead hypergraph of* S to be the hypergraph $\Gamma(S) := (X^{\pm}, VS, \mathrm{hl} : VS \to \mathcal{P}(X^{\pm}))$.

Given a cyclically reduced X-digraph S and an automorphism $\phi \in \operatorname{Aut} F(X)$, one may construct $\phi(S)$ from S as follows. First, for all $x \in X$, we subdivide every x-edge in S into a path and relabel this path with $\phi(x)$. We fold the resulting graph and then delete leaves until none remain; the final graph is $\phi(S)$. When S represents the conjugacy class $H^{\operatorname{Aut} F(X)}$, we have that $\phi(S)$ represents $\phi(H)^{\operatorname{Aut} F(X)}$.

When $\phi = (C, m)$ is a type II Whitehead automorphism, this construction has the special feature of being "local". Let $v \in VS$ be such that $m \in \mathrm{hl}(v)$, and let e be the m-edge with endpoints v and u for some $u \in VS$. We "unhook" each edge in $\mathrm{lk}(v)$ with label in $C - m$ and reconnect that edge to u instead. If $v \in VS$ is such that $m \notin \mathrm{hl}(v)$, we then construct an *auxiliary vertex* v_{aux} and an auxiliary m-edge with initial vertex v_{aux} and terminal vertex v. We again

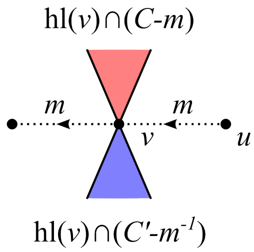

(a) The neighborhood of v before applying $\phi = (C, m)$.

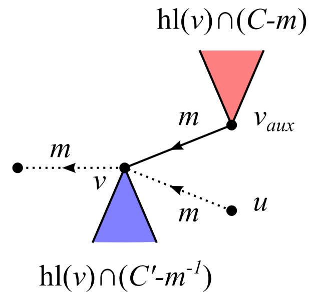

(b) Constructing the auxiliary vertex v_{aux}.

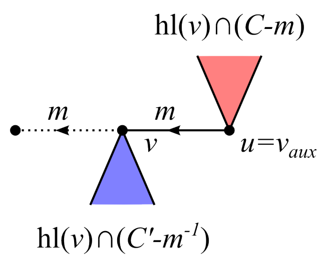

(c) Folding identifies v_{aux} and the vertex u at the other end of the edge corresponding to m in $\mathrm{hl}(v)$. If no such vertex u exists, no folding occurs.

Figure 2.1: Locally, the Whitehead automorphism $\phi = (C, m)$ moves the edges in $\mathrm{hl}(v) \cap (C - m)$ across the edge corresponding to $m \in \mathrm{hl}(v)$ (if present).

"unhook" the edges of $\mathrm{lk}(v)$ with label in $C - \{m\}$ and reconnect them to v_{aux}. The result of performing these moves at every vertex is the graph $\phi_{\text{aux}}(S)$, and we obtain $\phi(S)$ from $\phi_{\text{aux}}(S)$ by cyclic reduction. (See Figure 2.1; the dotted edges represent edges which may or may not be present.)

We make the following observations about $\phi_{\text{aux}}(S)$:

1. There is an injection from the vertex set of S to the set of non-auxiliary vertices of $\phi_{\text{aux}}(S)$; we will refer to this injection simply as ϕ_{aux}.

2. (a) Let e be an m-edge of S such that $e_- = u$ and $e_+ = v$. Then

$$\mathrm{hl}(\phi_{\text{aux}}(v)) = (\mathrm{hl}(v) \cap (C' \cup m) \cup (\mathrm{hl}(u) \cap (C - m). \qquad (2.1)$$

   (b) Let $v \in VS$ with $m \notin \mathrm{hl}(v)$. Then

$$\mathrm{hl}(v_{\text{aux}}) = m^{-1} \cup (\mathrm{hl}(v) \cap (C - m)) \, .$$

3. A vertex of $\phi_{\text{aux}}(S)$ is a leaf if and only if it is one of the following:

(a) v_{aux} for $v \in VS$ with $\text{hl}(v) \subseteq C'$; or

(b) $\phi_{\text{aux}}(v)$ for $v \in VS$ with $\text{hl}(v) \subseteq (C - m)$.

As a result, note that if $\phi = (\{m\}, m)$, then $\phi(S) = S$. If $\phi = (C, m)$ and $\phi' = (C', m^{-1})$, then note that $\phi(S) = \phi'(S)$. This latter observation allows us to assume that, without loss of generality, $m \in X$.

By keeping careful track of the construction for $\phi(S)$, it is possible to describe the change in the number of vertices between S and $\phi(S)$.

Proposition 2.7 ([44])**.** *Let S be a connected, cyclically reduced X-digraph with Whitehead hypergraph $\Gamma = \Gamma(S)$, and let $\phi = (C, m)$ be a type II Whitehead automorphism with $m \in X$. Then we have:*

$$\#V\phi(S) - \#VS = \text{cap}_\Gamma(C) - \text{deg}_\Gamma(m).$$

We will find it useful to recast Proposition 2.7 in terms of change in number of edges.

Proposition 2.8. *Let S be a connected, cyclically reduced X-digraph with Whitehead hypergraph $\Gamma = \Gamma(S)$, and let $\phi = (C, m)$ be a type II Whitehead automorphism with $m \in X$. Then we have:*

1. $\#E\phi(S) - \#ES = \text{cap}_\Gamma(C) - \text{deg}_\Gamma(m)$
2. *For a Whitehead automorphism $\phi = (C, m)$ with $m \in X$, we have*

$$\#E_x S = \#E_x \phi(H)$$

for all $x \neq m$.

Proof. Let S represent the conjugacy class $H^{\text{Aut}\,F(X)}$. Since S is connected, we have the well-known relation $\#ES = \#VS - 1 + R$, where R is the rank of H as a free group. Since $\phi(S)$ represents the class $\phi(H)^{\text{Aut}\,F(X)}$ and $\phi(H)$ must also have rank R, we then have $\#E\phi(S) = \#V\phi(S) - 1 + R$, and part 1 follows immediately.

Part 2 follows from the "local" version of the construction of $\phi(S)$. The only positive edges introduced in the subdivision stage have label m, and the only leaves which arise after subdivision and folding are leaves with hyperlink $\{m\}$. Therefore, the only positive edges added or removed in the application of ϕ to S are those labeled m. $\square$

We will recast Gersten's version of Whitehead's algorithm in graph-theoretic terms first seen in [24] and used later in [44] to analyze the complexity of the Whitehead reduction process.

Let S be a connected, cyclically reduced X-digraph and let $\phi \in \text{Aut}\,F(X)$. We call $\phi(S)$ an *automorphic image* of S. We say that ϕ *reduces* S if $\#VS < \#V\phi(S)$ (or equivalently, $\#ES < \#E\phi(S)$), and that ϕ *expands* S if $\#VS >$

$\#V\phi(S)$ (or equivalently, $\#ES > \#E\phi(S)$). Where S is clear from context, we will say that ϕ is *reducing* or *expanding*. If no automorphism reduces S, then we say that S is *minimal*.

Theorem 2.9 (Whitehead's Theorem [18]). *Let S be a connected, cyclically reduced X-digraph.*

1. *If S is not minimal, then some Whitehead automorphism reduces S.*

2. *Let S be minimal, and suppose there is $\phi \in \operatorname{Aut} F(X)$ such that $\phi(S)$ is also minimal. Then there exists a sequence of type II Whitehead automorphisms $\phi_1, \ldots, \phi_k$ such that ϕ_i does not expand $\phi_{i-1} \circ \cdots \circ \phi_1(S)$ and $\phi_k \circ \cdots \circ \phi_1(S) = \phi(S)$.*

Let S be a cyclically reduced X-digraph. Let $\min(S)$ denote the set of minimal automorphic images of S. Whitehead's Theorem gives an effective algorithm for constructing $\min(S)$ given S. Let S and T be cyclically reduced X-digraphs representing conjugacy classes $H^{\operatorname{Aut} F(X)}$ and $K^{\operatorname{Aut} F(X)}$; then there exists $\phi \in \operatorname{Aut} F(X)$ such that $K = \phi(H)$ if and only if $\min(S) = \min(T)$. This gives us Gersten's extension of Whitehead's famous algorithm to finitely generated subgroups.

Theorem 2.10 (Whitehead's algorithm). *There is an algorithm to decide, given $H, K \leq_{fg} F(X)$, whether or not there exists $\phi \in \operatorname{Aut} F(X)$ such that $\phi(H) = K$.*

2.1.3 Outer space

Definition 2.11 ($\mathbb{R}$-tree). An $\mathbb{R}$-*tree* is a geodesic metric space in which every two points are connected by a unique injective path and this path is a geodesic.

We say that the action of $F(X)$ on an $\mathbb{R}$-tree T is:

- *isometric* if each element $w \in F(X)$ acts as an isometry on T;
- *minimal* if there exists no $F(X)$-invariant subtree of T;
- *very small* if the stabilizer of any tripod is trivial and the stabilizer of any arc is either trivial or maximal cyclic in the stabilizers of the endpoints of the arc;
- *simplicial* if T has the topological structure of a simplicial complex.

We will now assume that all actions of $F(X)$ on $\mathbb{R}$-trees are isometric and minimal.

Definition 2.12 (Filling subgroup). Let $H \leq_{fg} F(X)$. We say that H is a *filling subgroup* if H fixes no point in any very small action of $F(X)$ on an $\mathbb{R}$-tree, and that H is a *non-filling subgroup* if H fixes a point in some very small action of $F(X)$ on an $\mathbb{R}$-tree. We say that $w \in F(X)$ is a *filling element* if w generates a filling subgroup of $F(X)$.

The work of Guirardel allows one to approximate the very small action of $F(X)$ on a given $\mathbb{R}$-tree by a very small action on a simplicial tree. In particular, if H fixes a point in the $\mathbb{R}$-tree, then we may have H fix a point in the simplicial approximation [22, Theorem 1].

Proposition 2.13. *A subgroup $H \leq_{fg} F(X)$ is non-filling if and only if H fixes a point in some very small action of $F(X)$ on a simplicial tree T.*

A very small action of $F(X)$ on a simplicial tree gives a particular type of decomposition of $F(X)$ called a graph of groups decomposition, the details of which can be found in [47]. We briefly review the associated terminology.

Definition 2.14 (Cyclic splitting)**.** A *cyclic splitting* of $F(X)$ is the decomposition of $F(X)$ as the fundamental group of a graph of groups with cyclic edge groups. A *free splitting* of $F(X)$ is the decomposition of $F(X)$ as the fundamental group of a graph of groups with trivial edge groups. An *edge map* refers to a homomorphism from an edge group to a vertex subgroup in a particular graph of groups. A splitting is *elementary* if the corresponding graph of groups is connected and has exactly one edge. An elementary splitting is a *segment splitting* if the underlying graph of groups has two distinct vertices and is a *loop splitting* if it has only one vertex. An elementary splitting is *nontrivial* if it is either a loop splitting or a segment splitting in which neither edge map is an isomorphism. An elementary cyclic splitting is *very small* if the image of the edge group is maximal cyclic in the vertex subgroup(s).

We say that $H \leq_{fg} F(X)$ is *elliptic* in a splitting of $F(X)$ if H is conjugate to a subgroup of a vertex subgroup. Subgroups which are not elliptic in a given splitting are said to be *hyperbolic*.

Proposition 2.15. *A subgroup $H \leq_{fg} F(X)$ is non-filling if and only if H is elliptic in either a nontrivial elementary free splitting of $F(X)$ or a nontrivial, very small, elementary cyclic splitting of $F(X)$.*

Proposition 2.16. *The vertex subgroups in a nontrivial, very small, elementary cyclic segment splitting of $F(X)$ have the form*

$$\langle A, b\rangle \text{ and } \langle B\rangle,$$

where $A \sqcup B$ is a basis for $F(X)$, $\#A \geq 1, \#B \geq 2$, and $b \in \langle B\rangle$ is not a proper power.

The vertex subgroup in a cyclic loop splitting of $F(X)$ has the form

$$\langle U, u^v\rangle,$$

where $U \sqcup \{v\}$ is a basis for $F(X)$ and $u \in \langle U\rangle$ is not a proper power.

Proof. This is a straightforward application of a lemma of Bestvina-Feighn [5, Lemma 4.1]. Similar results also appear in [48, 51, 52]. □

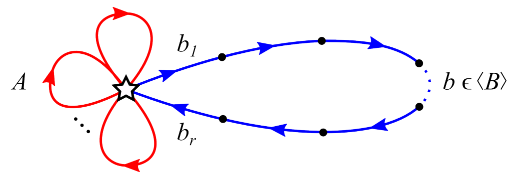

(a) A standard segment vertex subgroup.

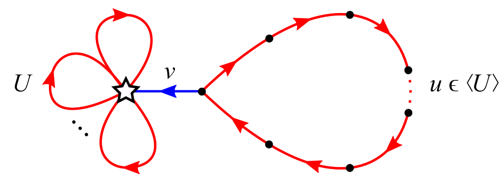

(b) A standard loop vertex subgroups.

Figure 2.2: Stallings graphs of standard vertex subgroups.

Definition 2.17 (Segment, loop vertex subgroups)**.** We call a subgroup $\langle A, b\rangle$ as in Proposition 2.16 a *segment vertex subgroup*. A subgroup $\langle U, u^v\rangle$ is called a *loop vertex subgroup*. When $A \sqcup B = X$ or $U \sqcup \{v\} = X$, we say that these vertex subgroups are *standard*. By $\mathcal{SV}$ and $\mathcal{LV}$ we denote the sets of standard segment and standard loop vertex subgroups, respectively.

If H is a segment vertex subgroup, then for an automorphism ϕ induced by a bijection $A \sqcup B \to X$, $\phi(H) \in \mathcal{SV}$. Thus, every segment vertex subgroup has an automorphic image in the set $\mathcal{SV}$. Likewise, every loop vertex subgroup has an automorphic image in $\mathcal{LV}$ and every proper free factor has an automorphic image in $\mathcal{SF}$.

Let $H \leq_{fg} F(X)$ be a proper free factor of $F(X)$. If $H = \langle Y\rangle$ where $Y \subset X$, then we say that H is a *standard free factor*. By $\mathcal{SF}$ we denote the set of standard free factors of $F(X)$. Note that every standard free factor is a subgroup of a standard loop vertex subgroup.

Proposition 2.18. *A subgroup $H \leq_{fg} F(X)$ is non-filling if and only if there exist $\phi \in \operatorname{Aut} F(X)$ and $K \in \mathcal{SV} \cup \mathcal{LV}$ such that $\phi(H) \leq K$.*

2.2 Main Results

2.2.1 Algorithmic Properties of Filling Subgroups

Definition 2.19 (Automorphic subgroup problem)**.** Let $\mathcal{K}$ be a (possibly infinite) collection of subgroups of $F(X)$. The *automorphic subgroup problem*, denoted ASP($\mathcal{K}$), is the problem:

> Given $H \leq_{fg} F(X)$, do there exist $K \in \mathcal{K}$ and $\phi \in \operatorname{Aut} F(X)$ such that $\phi(H) \leq K$?

Little seems to be known about ASP($\mathcal{K}$), even in the case where $\mathcal{H}$ consists of a single subgroup. However, in the case where $\mathcal{H}$ consists of a single cyclic subgroup, ASP($\mathcal{H}$) can be solved by Whitehead's algorithm.

Recall that $\mathcal{SF}$ is the set of standard free factors of $F(X)$, and that $H \leq_{fg} F(X)$ is contained in a proper free factor of $F(X)$ if and only if H has some automorphic image which is a subgroup of of a standard free factor.

Proposition 2.20. ASP($\mathcal{SF}$) *can be decided for any free group $F(X)$ of finite rank.*

Proof. Let $H \leq_{fg} F(X)$. Let Y be a basis for $F(X)$ such that $H \leq \langle Y' \rangle$ for some $Y' \subset Y$. Let $\phi \in \operatorname{Aut} F(X)$ be induced by a bijection $Y \to X$, so that $\phi(Y') := X' \subset X$. The graph $S = S_X(\phi(H))$ therefore omits some element $m \in X$ as an edge label.

Let $\psi = (C, m)$ be a Whitehead automorphism, where S omits m as an edge label. Proposition 2.8 states that applying ψ to S changes only the number of positive edges labeled m, and so ψ must expand S. We conclude that any reducing Whitehead automorphism for S must have a multiplier which occurs as an edge label in S. Therefore if the X-digraph S omits m as an edge label, S can be minimized without ever introducing m as an edge label. Every element of $\min(S) = \min(S_X(H))$ must therefore omit some letter of X from its set of edge labels. Conversely, it is straightforward to see that if some (every) element of $\min(S_X(H))$ omits a letter from X, then H is contained in a proper free factor. □

Corollary 2.21. *There is an algorithm to determine, given $H \leq_{fg} F(X)$, whether or not H is contained in a proper free factor of $F(X)$.*

Proof. The subgroup H is contained in a proper free factor if and only if there exist $\phi \in \operatorname{Aut} F(X)$ and $K \in \mathcal{SF}$ such that $\phi(H) \leq K$. However, $\phi(H) \leq K \in \mathcal{SF}$ if and only if some (every) element of $\min(S_X(\phi(H)))$ omits an element of X as an edge label. Since $\min(S_X(\phi(H))) = \min(S_X(H))$, our algorithm is as follows.

Algorithm 2.22. *Given $H \leq_{fg} F(X)$, we may determine whether or not H is contained in a proper free factor of $F(X)$ as follows:*

1. *Construct the finite graph $S_X(H)$.*

2. *Construct an element T of* $\min(S_X(H))$.

3. *Determine whether T omits some element of X as an edge label. If T omits some element of X as an edge label, conclude that H is contained in a proper free factor of $F(X)$. Otherwise, conclude that H is contained in no proper free factor of $F(X)$.*

□

2.2.1.1 The Rank Two Case

Let $F(a,b)$ denote the free group of rank two. The following characterization of the standard vertex subgroups of $F(a,b)$ follows directly from Proposition 2.16.

Proposition 2.23. *For $F(a,b)$, we have $\mathcal{SV} = \emptyset$ and $\mathcal{LV} = \{\langle a, a^b\rangle\}$.*

Clearly, an element $w \in F(a,b)$ is non-filling if and only if $\phi(w) \in \langle a, a^b\rangle$ for some $\phi \in \mathrm{Aut}\, F(a,b)$. The problem of identifying the non-filling elements of $F(a,b)$ is therefore equivalent to $\mathrm{ASP}(\langle a, a^b\rangle)$.

Theorem 2.24. *The problem $\mathrm{ASP}(\langle a, a^b\rangle)$ is decidable.*

Proof. Suppose that $H \leq_{fg} \langle a, a^b\rangle$ is cyclically reduced and that H is contained in no proper free factor of $F(a,b)$. We have an immersion $S_X(H) \to S_X(\langle a, a^b\rangle)$. Note that if this immersion is not a surjection, then H is contained in a proper free factor of $F(a,b)$ (either $\langle a\rangle$ or $\langle a^b\rangle$).

Every vertex of $S = S_X(H)$ therefore has a hyperlink which is a subset of either $\{a, a^{-1}, b\}$ or $\{a, a^{-1}, b^{-1}\}$. In particular, note that the set of initial vertices of b-edges in $S_X(H)$ is disjoint from the set of terminal vertices of b-edges, and so $S_X(H)$ has at least $2\#E_bS$ vertices.

Suppose that $S - E_bS$ has k connected components. Since S is cyclically reduced, each of these components has at least one a-edge. Since $S - E_bS$ is an a-digraph, each connected component of $S - E_bS$ is either a path or a cycle. We therefore have $\#E_aS \geq 2\#E_bS - k$, hence $\#E_aS + k \geq 2\#E_bS$. Since each of the connected components of $S - E_bS$ must have at least one a-edge, $\#E_aS \geq k$ and therefore $\#E_aS \geq \#E_bS$. In terms of the Whitehead hypergraph $\Gamma(S)$, we have $\deg_\Gamma(a) \geq \deg_\Gamma(b)$.

Now suppose that $\phi = (C, m)$ is a reducing Whitehead automorphism for H. Since C is an m-cut, if C has one or three elements, $\phi(S) = S$. C therefore has two elements; without loss of generality, we may assume that $C = \{a, b\}$.

Suppose that $m = a$, so that $\phi = (\{a,b\}, a)$ is reducing. Clearly ϕ leaves $S_X(\langle a, a^b\rangle)$ invariant, so $\phi(S)$ admits an immersion onto $S_X(\langle a, a^b\rangle)$.

Suppose that $m = b$, so that $\phi = (\{a,b\}, b)$ is reducing. Since $\deg_\Gamma(a) \geq \deg_\Gamma(b)$, the Whitehead automorphism $\phi' = (\{a,b\}, a)$ is also reducing for S. By the above observation, $\phi(S)$ also admits an immersion onto $S_X(\langle a, a^b\rangle)$.

Therefore, if S admits an immersion onto $S_X(\langle a, a^b\rangle)$, there is at least one element of $\min(S)$ which also admits such an immersion. It follows directly that an arbitrary subgroup $H \leq_{fg} F(a,b)$ has some automorphic image which

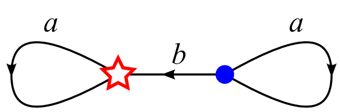

(a) $S_X(\langle a, a^b \rangle)$

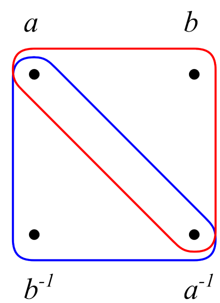

(b) $\Gamma(S_X(\langle a, a^b \rangle))$

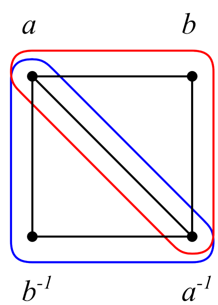

(c) $\Gamma(S_X(H))$ for $H \leq_{fg} F(X)$, showing every possible edge.

Figure 2.3: Graphs associated to the subgroup $\langle a, a^b \rangle$.

is a subgroup of $\langle a, a^b \rangle$ if and only if some element of $\min(S_X(H))$ admits an immersion onto $S_X(\langle a, a^b \rangle)$. Our algorithm is therefore the following:

Algorithm 2.25. *Given $H \leq_{fg} F(a,b)$, we may determine whether or not there exists $\phi \in \operatorname{Aut} F(a,b)$ such that $\phi(H) \leq \langle a, a^b \rangle$ as follows:*

1. *Construct the finite graph $S_X(H)$;*
2. *Construct the finite set $\min(S_X(H))$;*
3. *If some member of $\min(S_X(H))$ admits an immersion onto $S_X(\langle a, a^b \rangle)$, conclude that there is $\phi \in \operatorname{Aut} F(a,b)$ such that $\phi(H) \leq \langle a, a^b \rangle$. Otherwise, conclude that no such $\phi \in \operatorname{Aut} F(X)$ exists.*

$\square$

For $F(a,b)$, a subgroup $H \leq_{fg} F(X)$ is non-filling if and only if H has an automorphic image in $\langle a, a^b \rangle$. We may therefore solve the membership problem for the set of filling subgroups of $F(a,b)$.

Theorem 2.26. *There is an algorithm to determine, given $H \leq_{fg} F(a,b)$, whether or not H is a filling subgroup.*

2.2.1.2 Segment vertex subgroups in higher rank

Let $F(X)$ be a free group of rank at least three.

Recall that $\mathcal{SV}$ is the set of subgroups of $F(X)$ of the form $\langle A, b \rangle$ where $A \sqcup B = X$, $\#A \geq 1$, $\#B \geq 2$, and $b \in \langle B \rangle$ is not a proper power.

Let S be an X-digraph and let $Y \subseteq X$. The subgraph of S *spanned* by the Y-edges is the subgraph consisting of all Y-edges and all vertices having an incident Y-edge.

Definition 2.27 (Property (S))**.** We say that an X-digraph *satisfies property* (S) if:

1. S is connected and cyclically reduced; and
2. There is a partition $X = A \sqcup B$ such that
   (a) The subgraph spanned by the A-edges is a bouquet of single edge loops; and
   (b) The subgraph spanned by the B-edges is rank one.

Any element of $\mathcal{SV}$ has a Stallings graph which satisfies Property (S). We immediately obtain the following.

Proposition 2.28. *Let $H \leq_{fg} F(X)$. Then H is a subgroup of some element of $\mathcal{SV}$ if and only if $S_X(H)$ admits an immersion into a graph satisfying Property (S).*

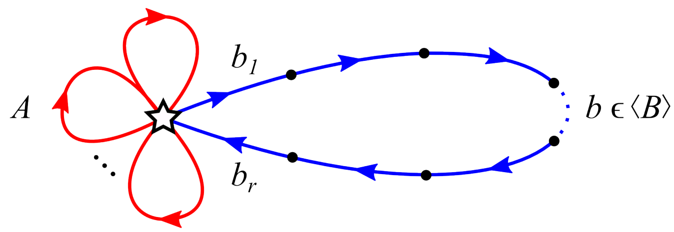

Figure 2.4: Stallings graph of a standard segment vertex subgroup.

Lemma 2.29. *Let S be a connected, cyclically reduced X-digraph admitting an immersion onto a graph satisfying property (S). Suppose that S represents (a conjugacy class of) a subgroup contained in no proper free factor of $F(X)$. If S is not minimal, then some element of* $\min(S)$ *also admits an immersion onto a graph satisfying property (S).*

Proof. Suppose T is an X-digraph satisfying Property (S) such that $\pi : S \to T$ is an immersion. Let $A \sqcup B = X$ be the partition given in the definition of Property (S). Let the *basepoint* of T be the unique vertex whose hyperlink meets $A^{\pm}$, and let $b = b_1 \dots b_r$ be the label of the loop in T labeled by B-edges, beginning and ending at the basepoint. Note that the hyperlink of the basepoint is $A^{\pm} \cup \{b_1^{-1}, b_r\}$.

Since $\pi : S \to T$ is an immersion, there is a k such that, for any non-basepoint $v \in VT$, the preimage $\pi^{-1}(v)$ is a set of exactly k vertices. More, the subgraph of S spanned by the B-edges is the union of exactly k paths labeled by b, any two of which are either disjoint or intersect only at one or both endpoints.

The following technical proposition will provide useful sufficient conditions for Property (S) to be preserved.

Proposition 2.30.

1. *Let $\phi = (C, m)$ be a Whitehead automorphism with $m \in A^{\pm}$ and let T be as above. If C does not cut the hyperlink of any non-basepoint vertex of T, then $\phi(T)$ also satisfies property (S).*

2. *Let $\phi = (C, m)$ be a Whitehead automorphism with $m \in B^{\pm}$ and let T be as above. If C does not cut the set $A^{\pm}$, then $\phi(T)$ also satisfies property (S).*

Proof. Suppose $\phi = (C, m)$ is as in part 1 of the proposition. In the construction of $\phi(S)$, new m-edges are only introduced at vertices of S whose hyperlinks are cut by C. Therefore, the only new edges introduced in the application of ϕ are incident to the basepoint; since every A-edge of T has the basepoint as its initial and terminal vertex, these new edges are folded away, leaving a B-labeled loop beginning and ending at the basepoint.

Suppose $\phi = (C, m)$ is as in part 2 of the proposition. If $A^{\pm}$ is not cut by C, then the effect of ϕ on T is to replace the loop labeled b with a loop labeled $\phi(b)$ or possibly $\phi(b)^{m^{-1}}$. The resulting graph satisfies Property (S). $\square$

Convention. To simplify notation, we define the following sets.

- $\Delta := (A^{\pm} \cap C) - \{m, m^{-1}, b_1^{-1}, b_r\}$
- $\Sigma := (A^{\pm} \cap C') - \{m, m^{-1}, b_1^{-1}, b_r\}$
- $\Pi := (B^{\pm} \cap C) - \{m, m^{-1}, b_1^{-1}, b_r\}$
- $\Omega := (B^{\pm} \cap C') - \{m, m^{-1}, b_1^{-1}, b_r\}$

Note that we do not necessarily have that m, m^{-1}, b_1^{-1}, and b_r are pairwise distinct.

Suppose that $b_1^{-1} = b_r$, and consider $\Gamma(S)$. Since in $\Gamma(S)$, the only element of $B^{\pm}$ adjacent to some element of $A^{\pm}$ is b_r, a direct calculation shows that the Whitehead automorphism $\phi = (B^{\pm} - b_r^{-1}, b_r)$ reduces S. By Proposition 2.30, $\phi(T)$ satisfies Property (S). We will therefore assume from now on that $b_1^{-1} \neq b_r$.

Suppose that $\phi = (C, m)$ reduces S, where $m \in A^{\pm}$. First note that if C does not cut $\{b_1^{-1}, b_r\}$, then either $(C \cup B^{\pm}, m)$ or $(C - B^{\pm}, m)$ is reducing for S, since only b_1^{-1} and b_r are adjacent to elements of $A^{\pm}$ in $\Gamma(S)$. By Proposition 2.30, the image of T under either of these Whitehead automorphisms again satisfies Property (S).

Now assume that, without loss of generality, $b_r \in C$ and $b_1^{-1} \in C' := X^{\pm} - C$ and that C cuts the hyperlink of some non-basepoint vertex of T. Since the preimage under π of a non-basepoint vertex is a set of k internal vertices in $S_X(H)$, we have $[\Pi, \Omega]_{\Gamma(S)} \geq k$. Therefore, passing from $(\Delta \cup \{m, b_r\} \cup \Pi, m)$ to $(\Delta \cup m, m)$ reduces the capacity by at least k (since at least k hyperedges contributing to capacity came from the hyperlink of an internal vertex) at the cost of adding $[b_r, \Delta \cup m]_{\Gamma(S)}$ to the capacity. However, a b_r-edge is coincident to an A-edge in at most k vertices of S, so $[b_r, \Delta \cup m]_{\Gamma(S)} \leq k$. Therefore, $\mathrm{cap}_{\Gamma(S)}(\Delta \cup m) \leq \mathrm{cap}_{\Gamma(S)}(\Delta \cup \{m, b_r\} \cup \Pi)$, and so $\phi' = (\Delta \cup m, m)$ must reduce S. Again, by Proposition 2.30, $\phi'(T)$ satisfies Property (S). (See Figure 2.5.)

Now suppose that $\phi = (\Delta \cup \{m, b_r\} \cup \Pi, m)$ reduces S, where $m \in B^{\pm}$. Once again, if C does not cut $\{b_1^{-1}, b_r\}$, then either $(\Delta \cup \Sigma \cup m \cup \Pi, m)$ or $(m \cup \Pi, m)$ will also reduce S. By Proposition 2.30, each of these Whitehead automorphisms preserve Property (S). We may therefore assume that $b_r \in C$ and $b_1^{-1} \in C'$.

Consider the quantities $[\Delta, b_r; \Sigma \cup \Omega \cup \{b_1^{-1}, m^{-1}\}]_{\Gamma(S)}$ and $[\Delta, \Sigma \cup \Omega \cup \{b_1^{-1}, m^{-1}\}; b_r]_{\Gamma(S)}$. Suppose that $[\Delta, b_r; \Sigma \cup \Omega \cup \{b_1^{-1}, m^{-1}\}]_{\Gamma(S)} \leq [\Delta, \Sigma \cup \Omega \cup$

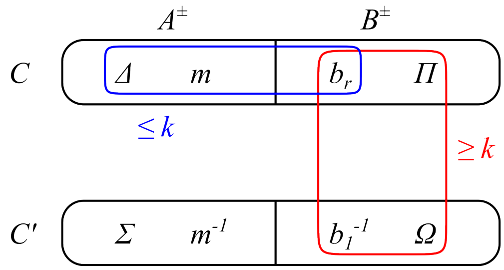

(a) $\Gamma(S)$ with $(\Delta \cup m \cup \Pi, m)$ reducing, $m \in A^{\pm}$.

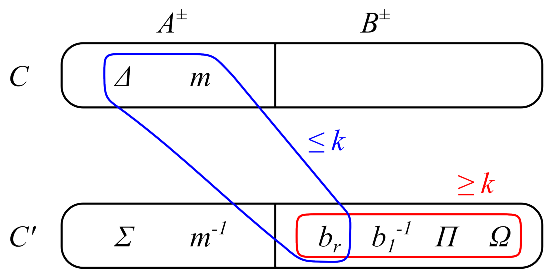

(b) $\Gamma(S)$ with $(\Delta \cup m, m)$ reducing, $m \in A^{\pm}$.

Figure 2.5: If $(\Delta \cup m \cup \Pi, m)$ reduces S with $m \in A^{\pm}$, then so must $(\Delta \cup m, m)$.

$\{b_1^{-1}, m^{-1}\}; b_r]_{\Gamma(S)}$. Moving Δ into C' must therefore not increase the capacity of the cut, and so $\phi' = (\{m, b_r\} \cup \Pi, m)$ must be reducing for S.

Suppose that $[\Delta, b_r; \Sigma \cup \Omega \cup \{b_r^{-1}, m^{-1}\}]_{\Gamma(S)} > [\Delta, \Sigma \cup \Omega \cup \{b_1^{-1}, m^{-1}\}; b_r]_{\Gamma(S)}$. A straightforward calculation then shows that the Whitehead automorphism $\phi' = (\Delta \cup b_r, b_r)$ reduces S.

Proposition 2.31. *For $i = 1, \dots, r$, define $\phi_i := (\Delta \cup b_i, b_i)$. If ϕ_r reduces S, then for all $i = 1, \dots, r-1$, the Whitehead automorphism ϕ_i reduces $\phi_{i+1} \cdots \phi_r(S)$. Furthermore, $\phi_1 \dots \phi_r(S)$ immerses onto a graph satisfying Property (S).*

Proof. First notice that since S immerses onto T with property (S), then $\phi_{i+1} \cdots \phi_r(S)$ immerses onto $\phi_{i+1} \cdots \phi_r(T)$. Set $\Pi_i := B^{\pm} - b_i$ and $\Pi_i^{-1} := B^{\pm} - b_i^{-1}$.

Suppose $v \in VS$ has type $(\Delta, b_r; \Sigma \cup \Pi_r)$. The vertex adjacent to v via the b_r edge will then have hyperlink type $(\Delta, b_{r-1}; \Sigma \cup \Pi_{r-1})$ in $\phi_r(S)$. Moreover, this is the only way in which a vertex of $\phi_r(S)$ may have type $(\Delta, b_{r-1}; \Sigma \cup \Pi_{r-1})$. Therefore

$$[\Delta, b_r; \Sigma \cup \Pi_r]_{\Gamma(S)} = [\Delta, b_{r-1}; \Sigma \cup \Pi_{r-1}]_{\Gamma(\phi_r(S))}.$$

Suppose $v \in VS$ has type $(\Delta, \Sigma \cup \Pi_r; b_r)$. Then v contributes an auxiliary vertex with hyperlink of type $(\Delta, b_r^{-1}; \Sigma \cup \Pi_r^{-1})$. Again, the only way a hyperlink

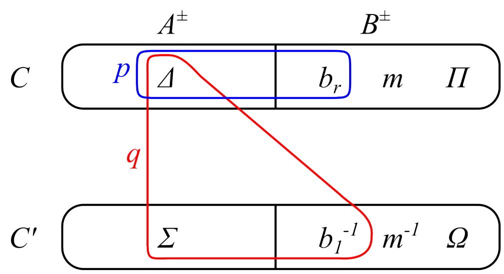

(a) $\Gamma(S)$ with (C, m) reducing, $m \in B^{\pm}$.

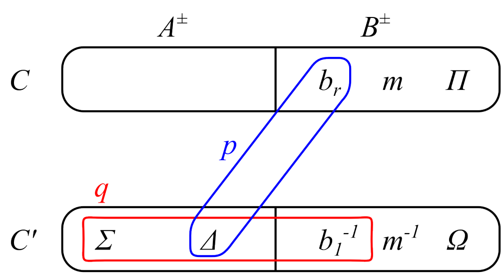

(b) If $p \leq q$, then $(\{b_r, m\} \cup \Pi, b_r)$ is reducing.

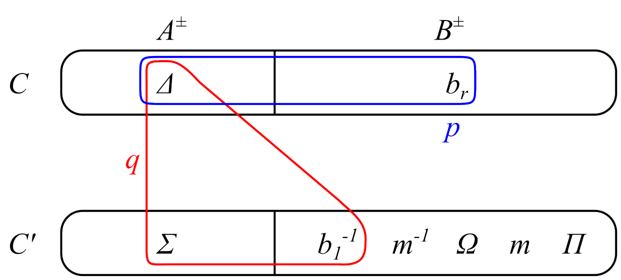

(c) If $p > q$, then $(\Delta \cup b_r, b_r)$ is reducing.

Figure 2.6: If (C, m) reduces S with $m \in B^{\pm}$, then either $(C \cap B^{\pm}, m)$ or $(\Delta \cup b_r, b_r)$ also reduces S.

of type $(\Delta, b_r^{-1}; \Sigma \cup \Pi_r^{-1})$ may arise is as such an auxiliary vertex, so

$$[\Delta, b_r^{-1}; \Sigma \cup \Pi_r^{-1}]_{\Gamma(\phi_r(S))} = [\Delta, \Sigma \cup \Pi_r; b_r]_{\Gamma(S)}.$$

However, since a vertex of $\phi_r(S)$ whose hyperlink meets Δ must have hyperlink contained in $\Delta \cup \{b_r^{-1}, b_{r-1}\}$, a vertex of $\phi_r(S)$ is of type $(\Delta, b_r^{-1}; \Sigma \cup \Pi_r^{-1})$ if and only if it is of type $(\Delta, \Sigma \cup \Pi_{r-1}; b_{r-1})$. We therefore have

$$[\Delta, \Sigma \cup \Pi_r; b_r]_{\Gamma(S)} = [\Delta, \Sigma \cup \Pi_{r-1}; b_{r-1}]_{\Gamma(\phi_r(S))}.$$

Given that $\phi_r = (\Delta \cup b_r, b_r)$ reduces S, it follows immediately that

$$[\Delta, b_r; \Sigma \cup \Pi_r]_{\Gamma(S)} > [\Delta, \Sigma \cup \Pi_r; b_r]_{\Gamma(S)}.$$

Using the above equivalences, we then have

$$[\Delta, b_{r-1}; \Sigma \cup \Pi_{r-1}]_{\Gamma(\phi_r(S))} > [\Delta, \Sigma \cup \Pi_{r-1}; b_{r-1}]_{\Gamma(\phi_r(S))},$$

which is equivalent to saying that $\phi_{r-1} = (\Delta \cup b_{r-1}, b_{r-1})$ reduces $\phi_r(S)$.

To see that ϕ_i reduces $\phi_{i+1} \cdots \phi_r(S)$, note that any hyperedge of $\Gamma(\phi_{i+1} \cdots \phi_r(S))$ incident to Δ is contained in $\Delta \cup \{b_{i+1}^{-1}, b_i\}$. A similar argument shows that any vertex whose hyperlink contributes to capacity and not degree in $\phi_{i+1} \cdots \phi_r(S)$ came from a vertex with hyperlink contributing to capacity and not degree in $\phi_{i+2} \cdots \phi_r(S)$, and similarly for vertices contributing to degree and not capacity. It then follows that ϕ_i reduces $\phi_{i+1} \cdots \phi_r(S)$.

Since the net effect of $\phi_1 \cdots \phi_r$ is to multiply the edges in Δ by the entire word b, it is immediate that $\phi_1 \cdots \phi_r(S)$ immerses onto T. □

By Proposition 2.31, if $\phi_r = (\Delta \cup b_r, b_r)$ reduces S, then we have an entire sequence of reducing Whitehead automorphisms which, when applied to S, yield an X-digraph that again immerses onto a graph with Property (S). Therefore, whenever S admits an immersion onto a graph satisfying Property (S), some element of $\min(S)$ is guaranteed to also admit an immersion onto a graph satisfying Property (S). □

Corollary 2.32. *Let $F(X)$ be a free group with $\#X \geq 3$. Then* $\mathrm{ASP}(\mathcal{SV})$ *is decidable.*

Proof. If $H \leq_{fg} F(X)$ is such that $\phi(H) \leq K \in \mathcal{SV}$, then $S_X(\phi(H))$ immerses onto an X-digraph satisfying Property (S). Therefore some element of $\min(S_X(\phi(H)))$ immerses onto an X-digraph satisfying Property (S); equivalently, some element of $\min(S_X(\phi(H)))$ has a principal quotient satisfying Property (S). Since $\min(S_X(\phi(H))) = \min(S_X(H))$, some element of $\min(S_X(H))$ has a principal quotient satisfying Property (S).

Our algorithm is therefore:

Algorithm 2.33. *Given $H \leq_{fg} F(X)$, we may determine whether or not there exist $\phi \in \operatorname{Aut} F(X)$ and $K \in \mathcal{SV}$ such that $\phi(H) \leq K$ as follows:*

1. *Construct the finite X-digraph $S_X(H)$.*

2. *Construct the finite set $\min(S_X(H))$.*

3. *Construct the finite set $\mathcal{PQ}(\min(S_X(H))) := \bigcup_{M \in \min(S_X(H))} \mathcal{PQ}(M)$.*

4. *For each $P \in \mathcal{PQ}(\min(S_X(H)))$, determine whether or not P satisfies Property (S). If a P satisfying Property (S) is found, conclude that there exist $\phi \in \operatorname{Aut} F(X)$ and $K \in \mathcal{SV}$ such that $\phi(H) \leq K$. Otherwise, conclude that no such $\phi \in \operatorname{Aut} F(X)$ and $K \in \mathcal{SV}$ exist.*

$\square$

Theorem 2.34. *Let $F(X)$ be a free group of finite rank at least three. There is an algorithm to determine, given $H \leq_{fg} F(X)$, whether or not H is elliptic in a nontrivial, very small, elementary cyclic splitting of $F(X)$.*

2.2.1.3 Loop vertex subgroups in higher rank

Let $F(X)$ be a free group with rank at least three.

Recall that the set $\mathcal{LV}$ is the set of standard loop vertex subgroups; in other words, groups of the form

$$\langle U, u^v \rangle$$

where $U \sqcup \{v\} = X$ and $u \in \langle U \rangle$ is not a proper power.

Observe that $S_X(\langle U, u^v \rangle)$ has a unique v-edge, and that the complement of this edge has at least one component of rank one.

Definition 2.35 (Property (L))**.** Let S be a Stallings graph. We say that S satisfies *Property* (L) if

1. There is some $x \in X$ for which $S_X(H)$ has a unique x-edge e; and

2. $S_X(H) - e$ has two connected components, at least of which is a rank one graph.

We say that $H \leq_{fg} F(X)$ satisfies Property (L) if $S_X(H)$ does.

We note that $S_X(\langle U, u^v \rangle)$ satisfies Property (L). Suppose T is a cyclically reduced subgraph of $S_X(\langle U, u^v \rangle)$. It is straightforward to verify that either T satisfies Property (L) or omits some $x \in X$ as an edge label.

Lemma 2.36. *Let S be a cyclically reduced X-digraph satisfying property (L), and let $\phi = (C, m)$ be a reducing Whitehead automorphism for S. Then either $\phi(S)$ satisfies Property (L) or $\phi(S)$ omits some element of X as an edge label.*

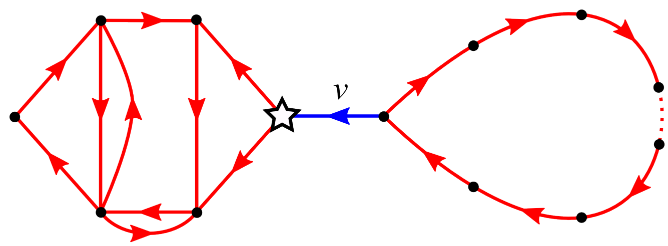

Figure 2.7: A Stallings graph satisfying Property (L). The unlabeled edges have labels different from v.

Proof. Let $\phi = (C, m)$ reduce S. Let $v \in X$ be such that S has a unique v-edge e and that $S - e$ has two components, at least one of which is rank one.

First, suppose that $m \neq v^{\pm}$. Then $\phi(S)$ also has a unique v-edge, since ϕ changes only the number of m-edges.

Recall that $\phi(S)$ is constructed from S in three stages: subdivision, folding, and leaf deletion. Let S_1 and S_2 be the connected components of $S - e$. Let u_i be the endpoint of e in S_i for $i = 1, 2$.

We may construct $\phi(S)$ by first subdividing and folding each S_i to obtain a graph S_i'. It is clear that, since ϕ is an automorphism, S_i and S_i' have the same rank. We then connect the vertices u_1 and u_2 via an appropriately oriented path labeled with $\phi(v)$, to obtain a graph T. By construction, T has at most two unfolded vertices, u_1 and u_2, and has a unique v-edge which is also a cut edge. Making the final two folds at u_1 and u_2 and deleting any leaves which T may have introduces no new paths between the endpoints of the unique v-edge, and so T satisfies Property (L).

Now suppose that $m = v^{\pm}$. Since ϕ reduces S, it must reduce the number of v-edges in S. Since S has exactly one v-edge, $\phi(S)$ must have no v-edges, so $\phi(S)$ omits some element of X as an edge label. $\square$

Theorem 2.37. *There is an algorithm to determine, given $H \leq_{fg} F(X)$, whether or not there exists $\phi \in \operatorname{Aut} F(X)$ such that $S_X(\phi(H))$ satisfies Property (L).*

Proof. If there exists such a $\phi \in \operatorname{Aut} F(X)$ such that $S_X(\phi(H))$ satisfies Property (L), then by Lemma 2.36, $\min(S_X(\phi(H)) = \min(S_X(H))$ has an element which satisfies Property (L). Since elements of $\min(S_X(H))$ represent subgroups which are automorphic images of H, the converse also holds. Therefore, the algorithm is as follows:

Algorithm 2.38. *Given $H \leq_{fg} F(X)$, we may determine whether or not there exists $\phi \in \operatorname{Aut} F(X)$ such that $S_X(\phi(H))$ satisfies Property (L) as follows:*

1. *Construct the finite graph* $S_X(H)$*.*

2. *Construct the finite set* $\min(S_X(H))$*.*

3. *If some element of the finite set* $\min(S_X(H))$ *satisfies Property* (L)*, conclude that there exists* $\phi \in \operatorname{Aut} F(X)$ *such that* $S_X(\phi(H))$ *satisfies Property* (L)*. Otherwise, conclude that no such* ϕ *exists.*

□

2.2.2 Genericity of Filling Elements

We will now turn our attention to the statistical properties of the set of filling elements of a non-Abelian free group $F(X)$. The following may be found in the author's preprint [50].

Definition 2.39 (Genericity [29])**.** Let $S \subseteq T \subseteq F(X)$. We say that S is *T-generic* in the sense of Arzhantseva-Ol'shanskiĭ if

$$\lim_{R\to\infty} \frac{\#(S \cap B_R)}{\#(T \cap B_R)} = 1,$$

where B_R denotes the set of elements of $F(X)$ whose X-length does not exceed R.

If the above limit converges exponentially quickly, then we say that S is *exponentially T-generic*. We say that S is *(exponentially) T-negligable* if its complement $T - S$ is (exponentially) T-generic.

2.2.2.1 The set TS'

In [29], Kapovich, Schupp, and Shpilrain construct an exponentially $F(X)$-generic set with several important properties related to Whitehead's algorithm.

Definition 2.40 (The set TS')**.** Let $C \subseteq F(X)$ be the set of cyclically and freely reduced elements of $F(X)$. The set TS is the set of $w \in C$ which are not proper powers, whose cyclic length is increased by every non-inner type II Whitehead automorphism, and whose conjugacy class is fixed by no type I Whitehead automorphism. The set TS' is the set of elements $w \in F(X)$ whose cyclic reductions are in TS.

Proposition 2.41 ([29, Theorem 8.5])**.** *Let* $\#X \geq 2$ *and let* $TS' \subseteq F(X)$ *be as above.*

1. *The set* TS' *is exponentially* $F(X)$*-generic.*

2. *For any nontrivial* $w \in TS'$*, the stabilizer of* w *in* $\operatorname{Aut} F(X)$ *is the infinite cyclic group generated by right-conjugation by* w*.*

3. *The membership problem for* TS' *is solvable in linear time.*

2.2.2.2 Filling elements

We first consider the case where $w \in F(X)$ is elliptic in an elementary splitting of $F(X)$ over the trivial group.

Lemma 2.42. *Let $w \in F(X)$ be elliptic in an elementary splitting of $F(X)$ over a trivial group. Then w has non-cyclic stabilizer in* $\operatorname{Aut} F(X)$.

Proof. To say that w is elliptic in an elementary splitting of $F(X)$ over a trivial group is equivalent to saying that w is contained in a proper free factor of $F(X)$.

Suppose that w is not a proper power. Let $A \sqcup B$ be a basis for $F(X)$ such that $\#A, \#B \geq 1$ and $w \in \langle A \rangle$. Let $\sigma : F(X) \to F(X)$ be right-conjugation by w. Define $\tau : F(X) \to F(X)$ via

$$\tau(x) = \begin{cases} x^w & \text{if } x \in A, \\ x & \text{if } x \in B, \end{cases}$$

where $x^w := w^{-1}xw$. Since $w \in \langle A \rangle$, τ is an automorphism of $F(X)$. Both σ and τ fix w. However, σ fixes exactly $\langle w \rangle$, while τ fixes $\langle w, B \rangle$. Thus σ must be distinct from every power of τ, so the $\operatorname{Aut} F(X)$ stabilizer of w cannot be cyclic.

If $w = z^r$ where $r > 1$ and z is not a proper power, then z is elliptic in an elementary cyclic splitting if and only if w is elliptic in that same splitting. We may therefore pass from w to its root z, which is also elliptic in the given splitting. The argument above shows that z has a non-cyclic stabilizer, and since the stabilizer of w contains that of z, the element w must have non-cyclic stabilizer in $\operatorname{Aut} F(X)$ as well. ☐

Since the set TS' is an exponentially $F(X)$-generic set whose elements all have cyclic stabilizers in $\operatorname{Aut} F(X)$, any set consisting of elements with non-cyclic stabilizers is exponentially $F(X)$-negligable.

Corollary 2.43. *The set of elements of $F(X)$ which lie in a proper free factor of $F(X)$ is exponentially $F(X)$-negligable.*

Remark. This is a slight generalization of results appearing in [10] and [13], which show that the set of primitive elements of $F(X)$ is $F(X)$-negligable.

Lemma 2.44. *Let $w \in F(X)$ be elliptic in some elementary cyclic splitting of $F(X)$. Then w has a non-cyclic stabilizer in* $\operatorname{Aut} F(X)$.

Proof. Suppose that w is not a proper power.

Let $w \in F(X)$ be elliptic in a segment of groups. Then there must exist a basis $A \sqcup B$ of $F(X)$ such that $\#A \geq 1$, $\#B \geq 2$, $b \in \langle B \rangle$, and either $w \in \langle A, b \rangle$ or $w \in \langle B \rangle$. Note that if b is a proper power of some $c \in F(X)$, then we would have $w \in \langle A, c \rangle$, so w would remain elliptic in a splitting of the same type.

Hence we may assume that b is not a proper power. Define an automorphism $\sigma : F(X) \to F(X)$ by

$$\sigma(y) = \begin{cases} y, & \text{if } y \in A \\ y^b, & \text{if } y \in B. \end{cases}$$

Any power of σ fixes the rank 2 subgroup $\langle A, b \rangle$ pointwise and so also fixes w, whereas right-conjugation by w fixes exactly the cyclic subgroup $\langle w \rangle$. Right-conjugation by w must therefore differ from every power of σ, so the stabilizer of w in $\operatorname{Aut} F(X)$ cannot be cyclic.

If $w \in \langle B \rangle$, since $\#A \geq 1$, w lies in a proper free factor of $F(X)$. Lemma 2.42 states that such an element has a non-cyclic stabilizer in $\operatorname{Aut} F(X)$.

Let $w \in F(X)$ be elliptic in a loop of groups. There then exists a basis $U \sqcup \{v\}$ of $F(X)$ such that $w \in \langle U, u^v \rangle$ for some $u \in \langle U \rangle$. We define the homomorphism $\tau : F(X) \to F(X)$ by

$$\begin{aligned} \tau(y) &= y \text{ for } y \in U \\ \tau(v) &= uv. \end{aligned}$$

Since $u \in \langle U \rangle$, τ is an automorphism. In particular, τ fixes the subgroup $\langle U, u^v \rangle$ pointwise, so no power of τ equals right-conjugation by x, which fixes only the cyclic subgroup $\langle w \rangle$. Again, the stabilizer of w in $\operatorname{Aut} F(X)$ therefore cannot be cyclic.

We handle the case where w is a proper power in the same way it was handled in the proof of Lemma 2.42. □

Theorem 2.45. *Let $F(X)$ be a finitely generated non-Abelian free group.*

1. *Let $w \in F(X)$. If the stabilizer of w in $\operatorname{Aut} F(X)$ is infinite cyclic, then w is filling.*

2. *The set of filling elements of $F(X)$ is exponentially $F(X)$-generic.*

3. *There exists an exponentially $F(X)$-generic subset S of $F(X)$ such that every element of S is filling and the membership problem for S is solvable in linear time.*

Proof. Part 1 follows from Lemmas 2.42 and 2.44. Since every element of TS' has a cyclic stabilizer in $\operatorname{Aut} F(X)$ (Proposition 2.41, part 1), every element of TS' must be filling. Part 2 then follows from the fact that TS' is exponentially $F(X)$-generic (Proposition 2.41, part 2). Finally, part 3 follows from Proposition 2.41, part 3, taking S to be exactly TS'. □

Chapter 3

Residual Properties of Limit Groups

3.1 Background

Let G be a group with a generating set X.

Definition 3.1 (Cayley graph)**.** The *Cayley graph* of G with respect to the generating set X, denoted $\mathrm{Cayley}(G, X)$, is an oriented graph with vertex set in bijection with G. The edge set is in bijection with $G \times X$, where the pair (g, x) corresponds to an edge having initial vertex g, terminal vertex gx, and label x.

For a fixed set X, an *X-word* is a finite sequence of elements of X. By X^* we denote the set of all X-words, including the empty word. When X is a generating set for a group G, then every element of X^* represents an element of G. Where it is necessary to distinguish between them, we will denote by $\overline{w}$ the element of G represented by $w \in X^*$.

Recall that for an element $g \in G$, the *word length with respect to X* or *X-length*, of g, denoted $|g|_X$, is number of letters in the shortest X-word representing g. Equivalently, $|g|_X$ is the number of edges in the shortest path from 1 to g in $\mathrm{Cayley}(G, X)$.

For an integer $R \geq 0$, the *ball of radius R with respect to generating set X* is the set $B_R(G, X) = \{g \in G : |g|_X \leq R\}$. Where G and X are clear from context, we will denote this set simply by B_R. Note that when X is a finite set, then B_R is also finite for any integer $R \geq 0$.

Finally, for elements $g, h \in G$, the *right-conjugate of h by g* is the element $h^g := g^{-1}hg$.

3.1.1 Γ-limit groups

Sela first introduced the notion of a limit group in [45] in his investigation of groups having the elementary theory of a non-Abelian free group. Sela later generalized this notion to that of a Γ-limit group, where Γ is some fixed torsion-free hyperbolic group [46].

Definition 3.2 (Residual properties)**.** Fix a group H. We say that a group G is *residually H* if for any $g \in G - 1$, there exists a homomorphism $\phi_g : G \to H$ such that $\phi_g(g) \neq 1$. A group G is *fully residually H* if for any finite set S of nontrivial elements of G, there exists a homomorphism $\phi_S : G \to H$ such that $1 \notin \phi_S(S)$. The homomorphisms ϕ_g and ϕ_S are called *H-discriminating homomorphisms* for g and S, respectively.

For the remainder of this chapter, Γ will denote a non-Abelian, torsion-free hyperbolic group.

Definition 3.3 (Γ-limit group [46])**.** We say that a group G is a *Γ-limit group* if G is finitely generated and fully residually Γ.

A trivial example of a Γ-limit group is Γ itself. For a more complicated example, it is well-known that fundamental groups of closed, orientable hyperbolic surfaces are F_2-limit groups, where F_2 denotes the free group of rank two.

We may produce new Γ-limit groups from existing limit groups through a construction called an *extension of a centralizer*. Extensions of centralizers will provide the basis for our analysis of the residual properties of limit groups.

Let G be a group, and given $g \in G$, let $C_G(u) = \{g \in G : u^g = u\}$ denote the centralizer of u in G.

Definition 3.4 (Extension of a centralizer [32])**.** Suppose that for some $u \in G$, the centralizer $C = C_G(u)$ is Abelian and that $\phi : C \to A$ is injective for some Abelian group A. We call the amalgamated product

$$G(u, A) := G *_{C=\phi(C)} A$$

the *extension of the centralizer C by A with respect to ϕ*. We will call the extension *direct* if $A = \phi(C) \times B$ for some subgroup $B \leq A$. A direct extension is *free of rank n* if $B \cong \mathbb{Z}^n$.

Having given the most general definition, we will now assume that all extensions of centralizers are free and of finite rank. We will omit reference to the homomorphism ϕ when it is clear from context.

The following proposition is well-known and will serve as the starting point for our investigation of the residual properties of Γ-limit groups.

Proposition 3.5. *The extension of centralizer $G(u, A)$ is a G-limit group.*

Proposition 3.6 ([32, Corollary 3])**.** *A maximal Abelian subgroup of $G(u, A)$ is either conjugate to a subgroup of G, conjugate to A, or cyclic.*

Definition 3.7 (Iterated extension of centralizers)**.** Let G be a group. An *iterated extension of centralizers over G* is a group H for which there exists a finite series

$$G = G_0 \leq G_1 \leq \cdots \leq G_k = H$$

such that for $i = 0, \ldots, k-1$, each G_{i+1} is an extension of a centralizer of G_i.

Since each G_{i+1} is fully residually G_i, we immediately obtain the following:

Proposition 3.8. *An iterated extension of centralizers over G is fully residually G.*

The following theorem of Kharlampovich and Myasnikov will allow us to approach the residual properties of arbitrary Γ-limit groups by considering iterated extensions of centralizers.

Proposition 3.9 ([31, Theorems D, E])**.** *Every Γ-limit group embeds into some iterated extension of centralizers over Γ.*

Recall that a subgroup $H \leq G$ is *malnormal* if $H \cap H^g = 1$ for all $g \in G - H$.

Definition 3.10 (CSA group [32])**.** A group G is called a *CSA-group* if every maximal Abelian subgroup of G is malnormal. G is called a *CSA*-group* if it is a CSA-group and has no elements of order 2.

We summarize some of the important properties of CSA- and CSA*-groups.

Proposition 3.11 ([32])**.**

1. *Any torsion-free hyperbolic group is a CSA*-group.*
2. *The class of CSA*-groups is closed under iterated extensions of centralizers.*
3. *Let G be a CSA-group and let $A \leq G$ be a maximal Abelian subgroup. Then there is $u \in G$ for which $A = C_G(u)$.*
4. *Let G be a CSA-group. For any maximal Abelian subgroup A, $N_G(A) = A$.*
5. *Let G be a CSA-group. Then commutativity is a transitive relation on the set $G - 1$.*

3.1.2 Relative hyperbolicity

The following discussion is taken from Osin [42] with some minor modifications to notation inspired by Hruska [23].

By a pair $(G, \mathbb{P})$ we denote a group G with a distinguished set of subgroups $\mathbb{P} = \{P_\lambda\}_{\lambda \in \Lambda}$. A subgroup $H \leq G$ is called *parabolic* if it is conjugate into some $P \in \mathbb{P}$, and *hyperbolic* otherwise. We call the conjugates of the elements of $\mathbb{P}$ *maximal parabolic subgroups.*

Definition 3.12 (Relative generating set)**.** Let $\mathcal{P} = \bigcup_{\lambda \in \Lambda} (P_\lambda - \{1\})$. We say that $X \subseteq G$ is a *relative generating set for* $(G, \mathbb{P})$ if G is generated by $X \cup \mathcal{P}$. If X is finite, we call it a *finite relative generating set.*

Definition 3.13 (Relative presentation)**.** We may consider G as a quotient of the group

$$F := (*_{\lambda \in \Lambda} P_\lambda) * F(X),$$

where $F(X)$ is the free group with basis X. Note that the group F is generated by $X \cup \mathcal{P}$.

For each $\lambda \in \Lambda$, let S_λ denote all the words in $(P_\lambda - 1)^*$ which represent the identity in P_λ. Further denote

$$\mathcal{S} := \bigcup_{\lambda \in \Lambda} S_\lambda.$$

Let $\mathcal{R} \subseteq (X \cup \mathcal{P})^*$ be such that the normal closure of $\mathcal{R}$ generates the kernel of the homomorphism $F \to G$. We say that $(G, \mathbb{P})$ has the *relative presentation*

$$\langle X, \mathcal{P} \mid \mathcal{R}, \mathcal{S} \rangle. \tag{3.1}$$

If X and $\mathcal{R}$ are finite, then we say that the relative presentation (3.1) is *finite*. If $(G, \mathbb{P})$ has a finite relative presentation, we say that $(G, \mathbb{P})$ is *finitely relatively presented*.

Suppose that $(G, \mathbb{P})$ has a relative presentation as in (3.1). If $W \in (X \cup \mathcal{P})^*$ represents the identity in G, then there is an expression

$$W =_F \prod_{i=1}^{k} R_i^{f_i} \tag{3.2}$$

with equality in the group F and such that $R_i \in \mathcal{R}$ and $f_i \in F$ for each i.

Definition 3.14 (Relative isoperimetric function)**.** Let $\theta : \mathbb{N} \to \mathbb{N}$. We say that θ is a *relative isoperimetric function* for $(G, \mathbb{P})$ if there exists a finite relative presentation with X and $\mathcal{R}$ as above such that for any $W \in (X \cup \mathcal{P})^*$ with $|W|_{X \cup \mathcal{P}} \leq n$, there exists an expression of the form (3.2) such that $k \leq \theta(n)$.

Definition 3.15 (Relative Dehn function)**.** We call the smallest relative isoperimetric function for a relative presentation the *relative Dehn function* of that relative presentation. If a relative presentation has no finite relative isoperimetric function, then we say that the relative Dehn function for that relative presentation is not well-defined.

Definition 3.16 (Relatively hyperbolic group)**.** We say that $(G, \mathbb{P})$ is a *relatively hyperbolic group* if $(G, \mathbb{P})$ has a finite relative presentation with a well-defined, linear relative Dehn function.

3.1.3 Iterated extensions of centralizers over Γ are relatively hyperbolic

We will now fix a non-Abelian, torsion-free hyperbolic group Γ. Our goal is next to show that an iterated extension of centralizers over Γ is hyperbolic relative to its maximal non-cyclic Abelian subgroups. We begin by noting the following results which may both be found in [17].

Proposition 3.17 ([17]). *Let $(G, \mathbb{P})$ be a torsion-free relatively hyperbolic group. Let U be a cyclic hyperbolic subgroup such that $N_G(U) = U$. Then $(G, \mathbb{P} \cup \{U\})$ is also a torsion-free relatively hyperbolic group.*

Proposition 3.18 ([17]). *Let $(G_1, \mathbb{P}_1)$ and $(G_2, \mathbb{P}_2)$ be relatively hyperbolic groups. Let $P \in \mathbb{P}_1$, and suppose that P is isomorphic to a parabolic subgroup of $(G_2, \mathbb{P}_2)$. Let $G = G_1 *_P G_2$. Then $(G, (\mathbb{P}_1 - \{P\}) \cup \mathbb{P}_2))$ is relatively hyperbolic.*

Corollary 3.19. *An iterated extension of centralizers over a torsion-free hyperbolic group Γ is hyperbolic relative a set of representatives of conjugacy classes of maximal non-cyclic Abelian subgroups.*

Proof. We induct on k, the number of steps in the iterated extension. If $k = 0$, $G_k = \Gamma$ is hyperbolic and we are done.

Suppose that $(G_k, \mathbb{P}_k)$ is relatively hyperbolic, where $\mathbb{P}_k$ is a set of representatives of conjugacy classes of maximal non-cyclic Abelian subgroups of G_k. Without loss of generality, we may assume that G_{k+1} is constructed by extending the centralizer $C(u) = C_{G_k}(u)$ of a hyperbolic element $u \in G_k$ by a rank n free Abelian group A, so that

$$G_{k+1} = G_k *_{C(u)} A.$$

Since u is hyperbolic in the CSA-group $(G_k, \mathbb{P}_k)$, the centralizer $C(u)$ is maximal Abelian and $N_{G_k}(C(u)) = C(u)$ by Proposition 3.11. Moreover, $C(u)$ is cyclic; otherwise, u would be contained in a maximal non-cyclic Abelian subgroup of $(G_k, \mathbb{P}_k)$, contradicting that u is hyperbolic. Therefore, by Proposition 3.17, $(G_k, \mathbb{P}_k \cup \{C(u)\})$ is relatively hyperbolic. The free Abelian group A may be viewed as the relatively hyperbolic group $(A, \{A\})$, so $C(u) \leq A$ is parabolic. By Proposition 3.18, $(G_{k+1}, \mathbb{P}_k \cup \{A\})$ is therefore a relatively hyperbolic group. Finally, Proposition 3.6 states that every maximal non-cyclic Abelian subgroup of G_{k+1} is conjugate to some member of $\mathbb{P}_k \cup \{A\}$, so G_{k+1} is indeed hyperbolic relative to its maximal non-cyclic Abelian subgroups. $\square$

3.1.4 Relative hyperbolic geometry

Fix a relatively hyperbolic group $(G, \mathbb{P})$ with finite relative generating set X. We call $\mathrm{Cayley}(G, X \cup \mathcal{P})$ the *relative Cayley graph.*

Recall that a metric space (X, d_X) is *δ-hyperbolic*, or simply *hyperbolic*, if it satisfies the *thin triangles condition*: for any geodesic triangle with sides α, β, γ, every point of α is δ-close in the metric d_X to some point of $\beta \cup \gamma$.

Proposition 3.20 ([42]). *Let $(G, \mathbb{P})$ be a relatively hyperbolic group. Then for any finite relative generating set X, the relative Cayley graph* $\mathrm{Cayley}(G, X \cup \mathcal{P})$ *is hyperbolic.*

We have two distinct metrics on $\text{Cayley}(G, X \cup \mathcal{P})$. The *relative metric* is denoted $d_{X\cup\mathcal{P}}$, and for $u, v \in \text{Cayley}(G, X \cup \mathcal{P})$, we define $d_{X\cup\mathcal{P}}(u, v)$ to be the least number of edges in any path in $\text{Cayley}(G, X \cup \mathcal{P})$ having u and v as endpoints. The *absolute metric* is denoted d_X, and for $u, v \in \text{Cayley}(G, X \cup \mathcal{P})$, we define $d_X(u, v)$ to be the least number of edges in any X-labeled path in $\text{Cayley}(G, X \cup \mathcal{P})$ having u and v as endpoints. Note that while $\text{Cayley}(G, X \cup \mathcal{P})$ is hyperbolic with respect to the relative metric, it will generally not be hyperbolic with respect to the absolute metric.

A *relative geodesic* is an isometry $p : [0, L] \to (\text{Cayley}(G, X \cup \mathcal{P}), d_{X\cup\mathcal{P}})$, where $[0, L]$ is a closed interval of real numbers. We say that the *endpoints* of p are $p(0)$ and $p(L)$. Since every point $\text{Cayley}(G, X \cup \mathcal{P})$ is a distance at most 1 from some vertex, we will assume that L is an integer and that p maps integers to vertices. For $u, v \in \text{Cayley}(G, X \cup \mathcal{P})$, we denote by $[u, v]_{X\cup\mathcal{P}}$ a relative geodesic with endpoints u and v.

Similarly, an *absolute geodesic* is an isometry $p : [0, L] \to (\text{Cayley}(G, X \cup \mathcal{P}), d_X)$. We denote an absolute geodesic having u and v as endpoints by $[u, v]_X$.

A *relative (absolute) broken geodesic* is a finite concatenation of relative (absolute) geodesics. For a finite collection $\{a_1, \ldots, a_k\}$ of points in $\text{Cayley}(G, X \cup \mathcal{P})$, we will denote by $[a_1, a_2, \ldots, a_k]_{X\cup\mathcal{P}}$ a broken relative geodesic which is the union of relative geodesics $\displaystyle\bigcup_{i=1}^{k-1} [a_i, a_{i+1}]_{X\cup\mathcal{P}}$. Likewise, $[a_1, a_2, \ldots, a_k]_X$ denotes the analogous broken absolute geodesic.

The *length* of a path α in $\text{Cayley}(G, X \cup \mathcal{P})$, denoted $\text{len}(\alpha)$, is the number of edges in the path. Note that $\text{len}([a, b]_{X\cup\mathcal{P}}) = d_{X\cup\mathcal{P}}(a, b)$ and $\text{len}([a, b]_X) = d_X(a, b)$, for instance.

Definition 3.21 (Fellow traveling)**.** Let $p, q : [0, L] \to (\text{Cayley}(G, X \cup \mathcal{P}), d_{X\cup\mathcal{P}})$ be relative geodesics. We say that p and q are *relative (absolute) k-fellow travelers* if $d_{X\cup\mathcal{P}}(p(i), q(i)) \le k$ (resp. $d_X(p(i), q(i)) \le k$) for every integer i in $[0, L]$. We say that p and q *relatively (absolutely) k-fellow travel for a length of L'* if $p|_{[0,L']}$ and $q|_{[0,L']}$ are relative (absolute) k-fellow travelers.

Remark. Our notion of k-fellow traveling is often referred to in the literature as *synchronous k-fellow traveling*, to distinguish it from *asynchronouse k-fellow traveling*, which does not respect the parameterization of the geodesics. We will not require the notion of asynchronous k-fellow traveling here.

Definition 3.22 (Relatively quasiconvex)**.** A subgroup H of $(G, \mathbb{P})$ is called *relatively quasiconvex* if there exists a constant $\epsilon > 0$ such that the following holds. Let $g, h \in H$ and let $[g, h]_{X\cup\mathcal{P}}$ be an arbitrary relative geodesic in $\text{Cayley}(G, X \cup \mathcal{P})$. Then for every vertex $v \in [g, h]_{X\cup\mathcal{P}}$, there exists a vertex $u \in H$ such that

$$d_X(v, u) \le \epsilon.$$

Definition 3.23 (Strongly relatively quasiconvex)**.** A relatively quasiconvex subgroup H of $(G, \mathbb{P})$ is called *strongly relatively quasiconvex* if the intersection $H \cap P^g$ is finite for any $g \in G$ and $P \in \mathbb{P}$.

Osin notes in Proposition 4.10 of [42] that the relative and strong relative quasiconvexity properties are invariant with respect to choice of finite generating set for G.

Proposition 3.24 ([42, 4.19])**.** *Let $(G, \mathbb{P})$ be a relatively hyperbolic group, and let $u \in G$ be a hyperbolic element. Then the centralizer $C_G(u)$ is a strongly relatively quasiconvex subgroup of G.*

Let $\lambda > 0$ and $c \geq 0$. Recall that a map of metric spaces $f : (X, d_X) \to (Y, d_Y)$ is a (λ, c)-*quasi-isometric embedding* if for all $a, b \in X$, we have

$$\frac{1}{\lambda} d_X(a, b) - c \leq d_Y\big(f(a), f(b)\big) \leq \lambda d_X(a, b) + c.$$

Proposition 3.25 ([42])**.** *Every strongly relatively quasiconvex subgroup of $(G, \mathbb{P})$ is quasi-isometrically embedded in* $\mathrm{Cayley}(G, \mathbb{P})$.

Proposition 3.26 ([42])**.** *Let u be a hyperbolic element of $(G, \mathbb{P})$. Then $C_G(u)$ is cyclic.*

Proposition 3.27 ([42])**.** *For any hyperbolic $u \in (G, \mathbb{P})$ generating its own centralizer, there are constants $\lambda_u > 0, c_u \geq 0$ such that*

$$\frac{1}{\lambda_u} |n| - c_u \leq d_{X \cup \mathcal{P}}(1, u^n) \leq \lambda_u |n| + c_u \tag{3.3}$$

for all $n \in \mathbb{Z}$.

3.2 Main Results

3.2.1 Relative hyperbolic geometry

We once again fix a relatively hyperbolic group $(G, \mathbb{P})$ with finite relative generating set X such that the relative Cayley graph $\mathrm{Cayley}(G, X \cup \mathcal{P})$ is δ-hyperbolic.

Lemma 3.28. *Let $u \in G$ be a hyperbolic element generating its own centralizer $U = C_G(u)$. There is a function $B_0 : \mathbb{N} \to \mathbb{N}$ depending only on $(G, \mathbb{P})$, X, and u such that the following holds.*

Let $g \in G - U$. Let $p, q \in U$ and $s, t \in gU$. For any $p', q' \in [p, q]_{X \cup \mathcal{P}}$ and $s', t' \in [s, t]_{X \cup \mathcal{P}}$ such that $[p', q']_{X \cup \mathcal{P}}$ and $[s', t']_{X \cup \mathcal{P}}$ are absolute k-fellow travelers, then

$$d_{X \cup \mathcal{P}}(p', q'),\ d_{X \cup \mathcal{P}}(s', t') \leq B_0(k).$$

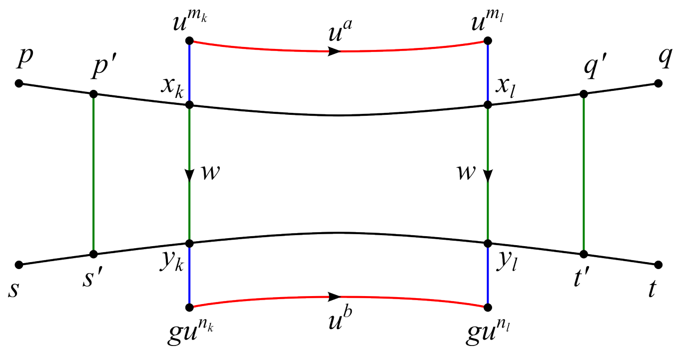

Figure 3.1: Producing the relation $w^{-1}u^a w = u^b$ in the proof of Lemma 3.28

Proof. Set $B_0(k) = (2\epsilon + 1)(2|X|)^{k+2\epsilon}$, and suppose that for some nonnegative integer k, there exist p, p', q, q', s, s', t, and t' which satisfy the hypotheses but such that $d_{X\cup\mathcal{P}}(p', q') > B_0(k)$.

We may find $(2|X|)^{k+2\epsilon}$ vertices, denoted x_i, on $[p', q']_{X\cup\mathcal{P}}$ such that if $i \neq j$ then $d_{X\cup\mathcal{P}}(x_i, x_j) > 2\epsilon$. To each x_i we may associate a $u^{m_i} \in U$ such that $d_X(x_i, y_i) \leq \epsilon$, since U is relatively quasiconvex. Note that if $i \neq j$, then $m_i \neq m_j$; otherwise, we would have $d_{X\cup\mathcal{P}}(x_i, x_j) \leq d_X(x_i, x_j) \leq 2\epsilon$, contradicting the choice of the x_i.

Since $[p', q']_{X\cup\mathcal{P}}$ and $[s', t']_{X\cup\mathcal{P}}$ are absolute k-fellow travelers, for each x_i there is a vertex $y_i \in [s', t']_{X\cup\mathcal{P}}$ such that $d_X(x_i, y_i) \leq k$. Since U is ϵ-quasiconvex, for each y_i there is $gu^{n_i} \in gU$ such that $d_X(y_i, gu^{n_i}) \leq \epsilon$.

To each point x_i, we associate the broken absolute geodesic $[u^{m_i}, x_i, y_i, gu^{n_i}]_X$. The length of such a path is at most $k+2\epsilon$, and there are $(2|X|)^{k+2\epsilon}$ such distinct paths, since no two of these paths have the same endpoint u^{m_i}.

However, there are strictly fewer than $(2|X|)^{k+2\epsilon}$ distinct path labels for paths of length at most $k + 2\epsilon$. Therefore, there are indices k, l such that $[u^{m_k}, x_k, y_k, gu^{n_k}]_X$ and $[u^{m_l}, x_l, y_l, gu^{n_l}]_X$ have the same label, w. As the endpoints of these w-labeled paths differ by elements of U, we obtain a relation of the form $w^{-1}u^a w = u^b$ for some integers a, b.

Since G is relatively hyperbolic, we must have that $a = \pm b$ [42, Corollary 4.21]. Therefore, w^2 commutes with u^a. Since G is a CSA-group and is therefore commutative-transitive (Proposition 3.11), w commutes with u and hence must be a power of u. This contradicts that U and gU are distinct cosets of U. □

Lemma 3.29. *Let $u \in G$ be a hyperbolic element generating a maximal cyclic subgroup U. There is a function $E_0 : \mathbb{N} \to \mathbb{N}$ depending only on $(G, \mathbb{P})$, X, and u such that the following holds.*

For all $m, n \in \mathbb{Z}$ with $m < 0 < n$, the relative geodesics $[1, u^m]_{X\cup\mathcal{P}}$ and $[1, u^n]_{X\cup\mathcal{P}}$ relatively k-fellow travel for a length of at most $E_0(k)$.

Proof. If not, since U is relatively quasiconvex and therefore quasi-isometrically embedded in Cayley$(G, X \cup \mathcal{P})$, there would have to be arbitrarily large powers

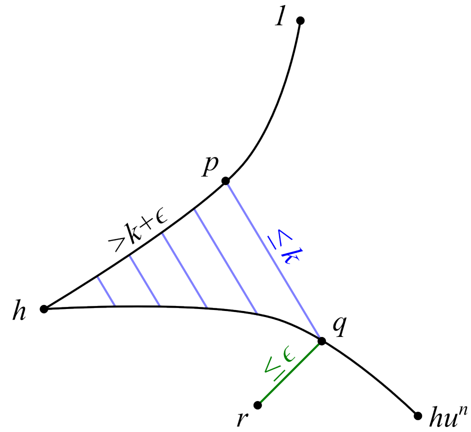

Figure 3.2: Finding a shorter coset representative in Lemma 3.30

of u which have relative length bounded above by a constant. However, this contradicts that U is quasi-isometrically embedded. $\square$

Let S be some set of elements of $(G, \mathbb{P})$. We say that $g \in S$ is an $X \cup \mathcal{P}$-*shortest element of* S if $|g|_{X\cup\mathcal{P}} \leq |h|_{X\cup\mathcal{P}}$ for every $h \in S$.

Lemma 3.30. *Let $u \in G$ generate a cyclic hyperbolic subgroup U. There is a function $C_0 : \mathbb{N} \to \mathbb{N}$ depending only on $(G, \mathbb{P})$, X, and u such that the following holds.*

Let h be an $X \cup \mathcal{P}$-shortest element of hU. Then for any integer n, the geodesics $[h, 1]_{X\cup\mathcal{P}}$ and $[h, hu^n]_{X\cup\mathcal{P}}$ absolutely k-fellow travel for no longer than $C_0(k)$.

Proof. Suppose that for fixed k and n, $[h, 1]_{X\cup\mathcal{P}}$ and $[h, hu^n]_{X\cup\mathcal{P}}$ absolutely k-fellow travel for longer than $k + \epsilon$. Then there is a vertex $p \in [h, 1]_{X\cup\mathcal{P}}$ with $d_{X\cup\mathcal{P}}(h, p) > k+\epsilon$ and such that there exists $w \in [h, hu^n]_{X\cup\mathcal{P}}$ with $d_X(p, q) \leq k$. Since U is relatively quasiconvex with constant ϵ, there is a vertex $r \in hU$ with $d_X(q, r) \leq \epsilon$. Then $[1, p, q, r]_{X\cup\mathcal{P}}$ is a broken relative geodesic of length at most $d_{X\cup\mathcal{P}}(1, p) + k + \epsilon < d_{X\cup\mathcal{P}}(1, h)$, contradicting that h among the $d_{X\cup\mathcal{P}}$-shortest elements of hU. (See Figure 3.2.) $\square$

Remark. The analogous statement holds for elements h which are $X \cup \mathcal{P}$-shortest in the coset Uh. Moreover, also note that if h is $X \cup \mathcal{P}$-shortest in UhU, then h is $X \cup \mathcal{P}$-shortest in both Uh and hU.

Proposition 3.31 ([41])**.** *Let $(G, \mathbb{P})$ be relatively hyperbolic with finite relative generating set X. There exist constants $\rho, \sigma > 0$ having the following property.*

Let Δ be a triangle with vertices x, y, z whose sides $[x, y]_{X\cup\mathcal{P}}$, $[y, z]_{X\cup\mathcal{P}}$, $[x, z]_{X\cup\mathcal{P}}$ are relative geodesics in Cayley$(G, X \cup \mathcal{P})$*. Suppose that u and v are vertices on $[x, y]_{X\cup\mathcal{P}}$ and $[x, z]_{X\cup\mathcal{P}}$ respectively such that*

$$d_{X\cup\mathcal{P}}(x, u) = d_{X\cup\mathcal{P}}(x, v)$$

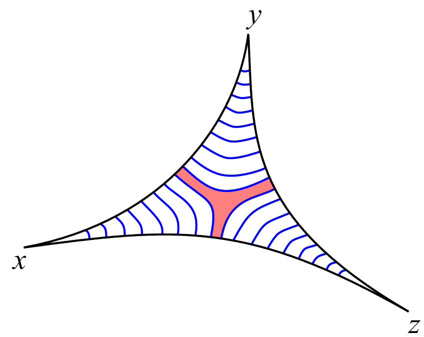

Figure 3.3: A relative geodesic triangle. The shaded lines join pairs of points on the triangle which are ρ-close in the absolute metric. The shaded area represents the region where the absolute ρ-fellow traveling property may fail.

and

$$d_{X\cup\mathcal{P}}(u,y) + d_{X\cup\mathcal{P}}(v,z) \geq d_{X\cup\mathcal{P}}(y,z) + \sigma.$$

Then

$$d_X(u,v) \leq \rho.$$

Recall that if x, y, and z are vertices in $\mathrm{Cayley}(G, X \cup \mathcal{P})$, then the *Gromov inner product* is defined as

$$\langle y|z\rangle_x := \frac{1}{2}(d_{X\cup\mathcal{P}}(x,y) + d_{X\cup\mathcal{P}}(x,z) - d_{X\cup\mathcal{P}}(y,z)).$$

Corollary 3.32. *Let ρ, σ, x, y, z be as in Proposition 3.31. Then adjacent sides $[x,y]_{X\cup\mathcal{P}}$ and $[x,z]_{X\cup\mathcal{P}}$ absolutely ρ-fellow travel for length at least $\langle y|z\rangle_x - \sigma/2$.*

Proof. Let $u \in [x,y]_{X\cup\mathcal{P}}$ and $v \in [x,z]_{X\cup\mathcal{P}}$ be such that $d_{X\cup\mathcal{P}}(x,u) = d_{X\cup\mathcal{P}}(x,v) = \ell$ and $d_{X\cup\mathcal{P}}(u,y) + d_{X\cup\mathcal{P}}(v,z) \geq d_{X\cup\mathcal{P}}(y,z) + \sigma$. We then have

$$d_{X\cup\mathcal{P}}(u,y) + d_{X\cup\mathcal{P}}(v,z) = d_{X\cup\mathcal{P}}(x,y) + d_{X\cup\mathcal{P}}(x,z) - 2\ell.$$

Further,

$$\begin{aligned} d_{X\cup\mathcal{P}}(x,y) + d_{X\cup\mathcal{P}}(x,z) - 2\ell &\geq d_{X\cup\mathcal{P}}(y,z) + \sigma \\ d_{X\cup\mathcal{P}}(x,y) + d_{X\cup\mathcal{P}}(x,z) - d_{X\cup\mathcal{P}}(y,z) - 2\ell &\geq \sigma \\ 2\langle y|z\rangle_x - 2\ell &\geq \sigma \\ \langle y|z\rangle_x - \sigma/2 &\geq \ell. \end{aligned}$$

Therefore, if $\ell \leq \langle y|z\rangle_x - \sigma/2$, then u and v satisfy the hypotheses of Proposition 3.31 and are therefore ρ-close in the absolute metric. $\square$

For a given relative geodesic triangle with vertices x, y, z, the *center* of the side $[x,y]_{X\cup\mathcal{P}}$ is the point $c \in [x,y]_{X\cup\mathcal{P}}$ such that $d_{X\cup\mathcal{P}}(x,c) = \langle y|z\rangle_x$ and $d_{X\cup\mathcal{P}}(y,x) = \langle x|z\rangle_y$.

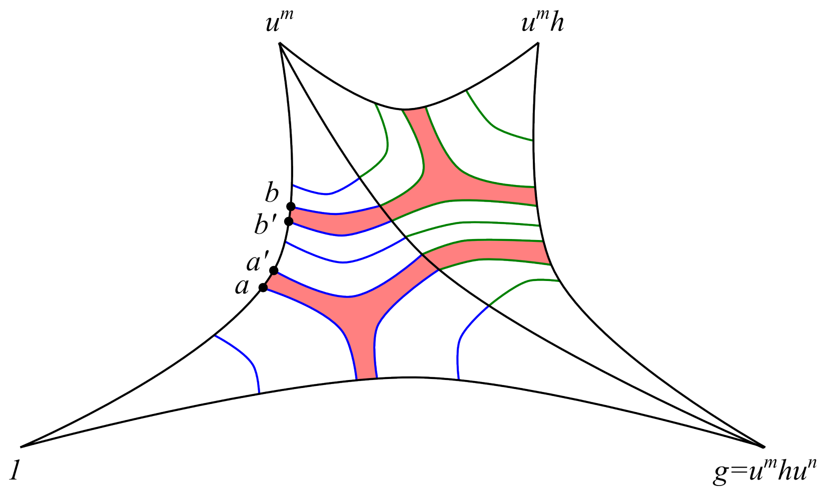

Figure 3.4: A decomposition of Q and one of its sides.

Lemma 3.33. *Let $u \in G$ generate a maximal cyclic hyperbolic subgroup U. Let $g \in G$, and let $h \in G$ be a $X \cup \mathcal{P}$-shortest element of UgU. There is a constant F_0 depending only on $(G, \mathbb{P})$, X, and u such that the following holds.*

Suppose that we have m and n such that $g = u^m h u^n$. Then $[1, u^m]_{X\cup\mathcal{P}}$ and $[u^m h, u^m h u^n]_{X\cup\mathcal{P}}$ each absolutely 2ρ-fellow travel $[1, u^m h u^n]_{X\cup\mathcal{P}}$ from their respective shared endpoints for all but at most F_0 of their length.

Proof. Let Q be the relative geodesic quadrilateral with sides $[1, u^m]_{X\cup\mathcal{P}}$, $[u^m, u^m h]_{X\cup\mathcal{P}}$, $[u^m h, u^m h u^n]_{X\cup\mathcal{P}}$, and $[1, u^m h u^n]_{X\cup\mathcal{P}}$.

By drawing a relative geodesic diagonal for Q, we obtain two relative geodesic triangles. As in Proposition 3.31, every pair of sides in either of these triangles absolutely ρ-fellow travel from their common vertex for a length of at least their Gromov inner product minus $\sigma/2$.

We extend the fellow-traveling property of the sides of these triangles to the sides of Q. (See Figure 3.4 for one configuration of such an extension; the shaded area represents the area near the centers of the triangles where absolute fellow traveling is not guaranteed.) We see that there exist vertices $a, a', b, b' \in [1, u^m]_{X\cup\mathcal{P}}$ such that:

1. The subpath $[1, a]_{X\cup\mathcal{P}}$ and some initial subpath of $[1, u^m h u^n]_{X\cup\mathcal{P}}$ absolutely 2ρ-fellow travel;

2. The subpath $[u^m, b]_{X\cup\mathcal{P}}$ and some initial subpath of $[u^m, u^m h]_{X\cup\mathcal{P}}$ absolutely 2ρ-fellow travel;

3. The subpath $[a', b']_{X\cup\mathcal{P}}$ absolutely 2ρ-fellow travels some subpath of $[u^m h u^n, u^m h]_{X\cup\mathcal{P}}$; and

4. The relative lengths of the subpaths $[a, a']_{X\cup\mathcal{P}}$ and $[b', b]_{X\cup\mathcal{P}}$ do not exceed σ.

We are interested in the total length of the subpath $[a, u^m]_{X\cup\mathcal{P}}$, since, as noted, $[1, a]_{X\cup\mathcal{P}}$ fellow travels with a subpath of $[1, u^m h u^n]_{X\cup\mathcal{P}}$. Observation (2) above implies that the length of $[u^m, b]_{X\cup\mathcal{P}}$ is at most $C_0(2\rho)$, by Lemma 3.30. Observation (3) implies that the length of $[a', b']_{X\cup\mathcal{P}}$ is at most $B_0(2\rho)$, by Lemma 3.28.

Consequently, we have that

$$\operatorname{len}([a, u^m]_{X\cup\mathcal{P}}) \leq B_0(2\rho) + C_0(2\rho) + 2\sigma =: F_0.$$

$\square$

Lemma 3.34. *Let u, g, h, m, and n be as in Lemma 3.33. Then we have*

$$\operatorname{len}([1, u^m, u^m h, u^m h u^n]_{X\cup\mathcal{P}}) \leq 3|g|_{X\cup\mathcal{P}} + 2F_0.$$

Proof. The lengths of the subpaths $[1, u^m]_{X\cup\mathcal{P}}$ and $[u^m h, u^m h u^n]_{X\cup\mathcal{P}}$ are bounded above by $|g|_{X\cup\mathcal{P}} + F_0$ by Lemma 3.33. Since h is a $X \cup \mathcal{P}$-shortest representative of UgU, we have $|h|_{X\cup\mathcal{P}} \leq |g|_{X\cup\mathcal{P}}$, and so the length of $[u^m, u^m h]_{X\cup\mathcal{P}}$ is at most $|g|_{X\cup\mathcal{P}}$. $\square$

Let $\mathbf{r} = (r_0, r_1, \ldots, r_k)$ be a tuple of integers. We define

$$\min(\mathbf{r}) := \min_i |r_i|.$$

Lemma 3.35. *Let $(G, \mathbb{P})$ be a relatively hyperbolic group with finite generating set X, and let U be a subgroup generated by a hyperbolic element $u \in G$. There exists a positive integer N_0 depending only on $(G, \mathbb{P})$, X, and u such that the following holds.*

Let $\mathbf{h} = (h_1, h_2, \ldots, h_k)$ be a tuple of elements of X such that each h_i is $X \cup \mathcal{P}$-shortest in the double coset $Uh_iU \neq U$, and let $\mathbf{r} = (r_0, r_1, \ldots, r_k)$ be a tuple of integers. Define

$$w_{\mathbf{h}}(\mathbf{r}) := u^{r_0} h_1 u^{r_1} h_2 u^{r_2} \cdots u^{r_{k-1}} h_k u^{r_k}.$$

Then $w_{\mathbf{h}}(\mathbf{r}) \neq 1$ in G for all $\mathbf{r}$ such that $\min(\mathbf{r}) > N_0$.

Proof. Let α be a path in $\operatorname{Cayley}(G, X \cup \mathcal{P})$ labeled by

$$(u^{r_0} * h_1 * u^{\lfloor r_1/2 \rfloor}) * (u^{\lceil r_1/2 \rceil} * h_2 * u^{\lfloor r_2/2 \rfloor}) * \cdots * (u^{\lceil r_{k-1}/2 \rceil} * h_k * u^{r_k}),$$

where $*$ denotes concatenation of words (as opposed to concatenation followed by free reduction) and $\lfloor \cdot \rfloor, \lceil \cdot \rceil$ are the usual floor and ceiling functions. Let α_1 be the subpath labeled by $u^{r_0} * h_1 * u^{\lfloor r_1/2 \rfloor}$ and α_k the subpath labeled by $u^{\lceil r_{k-1}/2 \rceil} * h_k * u^{r_k}$, and for each $i = 2, \ldots, k-1$, let α_i be the subpath of α labeled by $u^{\lceil r_{i-1}/2 \rceil} * h_i * u^{\lfloor r_i/2 \rfloor}$. The path α is then the concatenation of the α_i. Further define the vertices v_{i-1} and v_i to be the endpoints of α_i for each i.

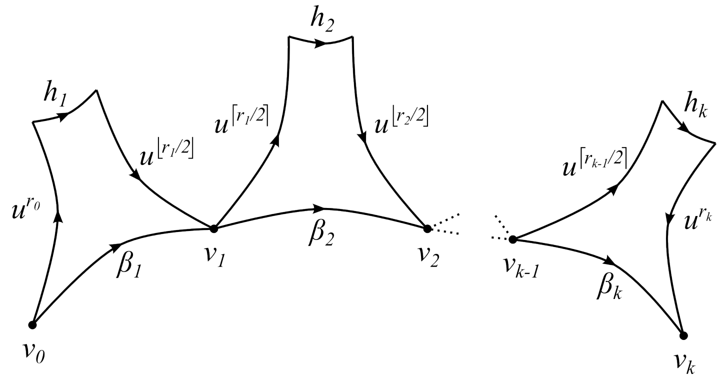

Figure 3.5: The decomposition of α.

Finally, for each i, define β_i to be a relative geodesic $[v_{i-1}, v_i]_{X \cup \mathcal{P}}$, and define β to be the broken relative geodesic which is the concatenation of the β_i. (See Figure 3.5.)

Lemma 3.36. *For each i and n we have*

$$\frac{2}{\lambda_u} \lfloor \min(\mathbf{r})/2 \rfloor - 2c_u - 2F_0 \leq \mathrm{len}(\beta_i). \tag{3.4}$$

Proof. This follows directly from Proposition 3.27 and Lemma 3.33. □

Proposition 3.37. *For all $\mathbf{r}$ with*

$$\lfloor \min(\mathbf{r})/2 \rfloor > \lambda_u(E_0(4\rho + \delta) + F_0 + c_u) \tag{3.5}$$

and $1 \leq i < k$, β_i and β_{i+1} relatively δ-fellow travel for a length of at most $E_0(4\rho + \delta)$ from their common endpoint v_i.

Proof. Suppose there is an $\mathbf{r}$ satisfying (3.5) and i such that β_i and β_{i+1} relatively δ-fellow travel for a length longer than $E_0(4\rho+\delta)$. By construction, there are relative geodesics γ_{i-1} and γ_i starting at v_i labeled by $u^{-\lfloor r_i/2 \rfloor}$ and $u^{\lceil r_i/2 \rceil}$ respectively. These relative geodesics absolutely 2ρ-fellow travel β_i and β_{i+1} for all but at most F_0 of their length. By choice of $\mathbf{r}$ and Corollary 3.27, γ_j and β_j are absolute 2ρ-fellow travelers for a length of at least $E_0(4\rho+\delta)$ for $j = i, i+1$.

However, if β_i and β_{i+1} are relative δ-fellow travelers for longer than $E_0(4\rho+\delta)$, then γ_i and γ_{i+1} are relative $(4\rho+\delta)$-fellow travelers for longer than $E_0(4\rho+\delta)$, contradicting Lemma 3.29. (See Figure 3.6.) □

Note that in a relative geodesic triangle, adjacent sides relatively δ-fellow travel for a length of at least the Gromov inner product. This fellow traveling property allows us to show that the concatenation of relative geodesic segments is a quasi-geodesic with parameters depending on the Gromov inner product.

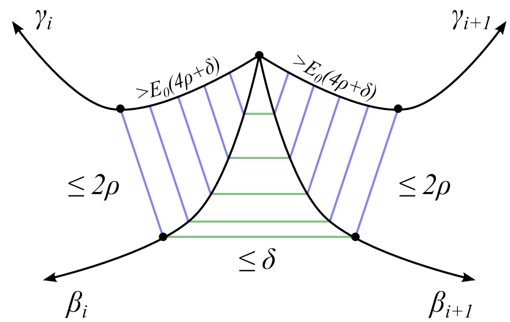

Figure 3.6: β_i and β_{i+1} cannot fellow travel too far without causing γ_i and γ_{i+1} to fellow travel.

Proposition 3.38. *Let $x, y, z \in \mathrm{Cayley}(G, X \cup \mathcal{P})$. Then every subpath of the broken relative geodesic $[x, y, z]_{X\cup\mathcal{P}}$ is a $(1, 2\langle x|z\rangle_y + 2\delta)$-quasigeodesic.*

Proposition 3.38 shows that for $\mathbf{r}$ satisfying (3.5), every adjacent pair of relative geodesics β_i and β_{i+1} is a relative $(1, 2E_0(4\rho + \delta) + 2\delta)$-quasigeodesic.

Proposition 3.39 ([33, Lemma 4.8])**.** *Let Y be a δ-hyperbolic space. Given quasigeodesity constants (λ, c), there exist κ, λ', and c' such that every k-local (λ, c)-quasigeodesic is a (λ', c')-quasigeodesic.*

Proposition 3.40. *Let κ, λ', c' be such that in $\mathrm{Cayley}(G, X \cup \mathcal{P})$, every κ-local $(1, 2E_0(4\rho+\delta)+2\delta)$-quasigeodesic is a (λ', c')-quasigeodesic. Let $\mathbf{r}$ satisfy (3.5) and further assume that*

$$\lfloor \min(\mathbf{r})/2 \rfloor \geq \lambda_u \left(\frac{\kappa}{2} + c_u + F_0 \right). \tag{3.6}$$

Then β is a (λ', c')-quasigeodesic.

Proof. By Proposition 3.38, for every i, the broken geodesic $\beta_i \cup \beta_{i+1}$ is a $(1, 2E_0(4\rho+\delta)+2\delta)$-quasigeodesic. The inequality (3.6) implies that the length of each β_i is larger than κ. Every subpath of β of length at most κ is contained in $\beta_i \cup \beta_{i+1}$ for some i, and is therefore a relative $(1, 2E_0(4\rho+\delta)+2\delta)$-quasigeodesic. The conclusion then follows from applying Proposition . $\square$

Now let $\mathbf{r}$ be such that

$$\lfloor \min(\mathbf{r})/2 \rfloor > \lambda_u \left(\frac{c'}{2} + c_u + F_0 \right). \tag{3.7}$$

Then the length of each β_i is at least c', and so the length of β is at least c'. The broken relative geodesic β, which is also a (λ', c')-quasigeodesic, therefore has necessarily distinct endpoints. Since α has the same endpoints as β and is labeled by $w_{\mathbf{h}}(\mathbf{r})$, we have $w_{\mathbf{h}}(\mathbf{r}) \neq 1$ in G.

Let N_{-1} be an integer larger than the right hand side in the inequalities (3.5), (3.6), and (3.7). Pick an integer N_0 such that $N_0 > 2N_{-1} + 2$. Then for all (r) with $\min(\mathbf{r}) > N_0$, we have that $\lfloor \min(\mathbf{r})/2 \rfloor > N_{-1}$. Thus N_0 is the promised constant. $\square$

Lemma 3.41. *Let $(G, \mathbb{P})$ be a relatively hyperbolic group with finite generating set X, and let U be a subgroup generated by a hyperbolic element $u \in G$. There is a linear function $N_1 : \mathbb{N} \to \mathbb{N}$ such that the following holds.*

Let $\mathbf{g} = (g_1, g_2, \dots, g_k)$ be a tuple of X-words such that $\sum_{i=1}^{k} |g_i|_X \le R$ and $g_i \in G - U$ for all i. For any tuple of integers $\mathbf{r} = (r_0, \dots, r_k)$, define

$$w_{\mathbf{g}}(\mathbf{r}) := u^{r_0} g_1 u^{r_1} g_2 u^{r_2} \cdots u^{r_{k-1}} g_k u^{r_k}. \tag{3.8}$$

Then we have $w_{\mathbf{g}}(\mathbf{r}) \neq 1$ in G for all $\mathbf{r}$ such that $\min(\mathbf{r}) > N_1(R)$.

Proof. Consider a single g_i. We may write $g_i = u^{s_i} h_i u^{t_i}$ with h_i a $X \cup \mathcal{P}$-shortest element of $U g_i U$. By Lemma 3.33, we have

$$|u^{s_i}|_{X \cup \mathcal{P}}, |u^{t_i}|_{X \cup \mathcal{P}} \le |g_i|_X + F_0 \le R + F_0.$$

Using the constants λ_u and c_u from Proposition 3.27, define

$$N_1(R) := N_0 + 2\lambda_u (R + F_0 + c_u),$$

where N_0 is the constant from Theorem 3.35. Note that $\lambda_u (R + F_0 + c_u) > |s_i|, |t_i|$ for all i.

Let $\mathbf{r} = (r_0, r_1, \dots, r_k)$ be a tuple of integers with $\min(\mathbf{r}) > N_1(R)$. Then we have

$$\begin{aligned} w_{\mathbf{g}}(\mathbf{r}) &= u^{r_0} g_1 u^{r_1} g_2 u^{r_2} \cdots u^{r_{k-1}} g_k u^{r_k} \\ &= u^{r_0} (u^{s_1} h_1 u^{t_1}) u^{r_1} (u^{s_2} h_2 u^{t_2}) u^{r_2} \cdots u^{r_{k-1}} (u^{s_k} h_k u^{t_k}) u^{r_k}) \\ &= (u^{r_0 + s_1}) h_1 (u^{t_1 + r_1 + s_2}) h_2 (u^{t_2 + r_2 + s_3}) \cdots (u^{t_{k-1} + r_{k-1} + s_k}) h_k (u^{t_k + r_k}) \end{aligned} \tag{3.9}$$

where every exponent of u appearing in (3.9) has magnitude at least N_0. By Theorem 3.35, $w_{\mathbf{g}}(\mathbf{r})$ is nontrivial in G. $\square$

Lemma 3.42. *Let $(G, \mathbb{P})$ be a relatively hyperbolic group with finite generating set X, and let U be a subgroup generated by a hyperbolic element $u \in G$. There is a linear function $N_2 : \mathbb{N} \to \mathbb{N}$ such that the following holds.*

Let $\mathbf{g} = (g_1, g_2, \dots, g_k)$ be a tuple of X-words, and let g_0, g_{k+1} be X-words such that $\sum_{i=0}^{k+1} |g_i|_X \le R$ and $g_i \in G - U$ for all i. Let $\mathbf{r} = (r_0, \dots, r_k)$ be a tuple of integers and define

$$w_{\mathbf{g}}(\mathbf{r}) = u^{r_0} g_1 u^{r_1} g_2 u^{r_2} \cdots u^{r_{k-1}} g_k u^{r_k}.$$

Then for all $\mathbf{r}$ *such that* $\min(\mathbf{r}) > N_2(R)$*, the elements*

$$\begin{aligned} w_{\mathbf{g}}(\mathbf{r}), \\ g_0 w_{\mathbf{g}}(\mathbf{r}), \\ w_{\mathbf{g}}(\mathbf{r}) g_{k+1}, \ \textit{and} \\ g_0 w_{\mathbf{g}}(\mathbf{r}) g_{k+1} \end{aligned}$$

are all nontrivial in G*.*

Proof. Note that if

$$\min(\mathbf{r}) > 2\lambda_u \left(2\lambda' R + c_u + F_0 + c'\right) + 2,$$

then

$$\lfloor \min(\mathbf{r})/2 \rfloor > \lambda_u \left(2\lambda' R + c_u + F_0 + c'\right)$$

and therefore $|w_{\mathbf{g}}(\mathbf{r})|_{X\cup\mathcal{P}} > 2R$. Define

$$N_2(R) := N_1(R) + 2\lambda_u \left(2\lambda' R + c_u + F_0 + c'\right) + 2,$$

and note that since N_1 is linear in R, so is N_2.

Then for all $\mathbf{r}$ with $\min(\mathbf{r}) > N_2(R)$, we have $|w_{\mathbf{g}}(\mathbf{r})|_{X\cup\mathcal{P}} > 2R \geq |g_0|_{X\cup\mathcal{P}} + |g_{k+1}|_{X\cup\mathcal{P}}$, and so none of the promised words are trivial in G by Lemma 3.41. □

3.2.2 Discriminating complexity

Let H be a finitely generated group, and let G be a fully residually H group. Let X and Y be fixed finite generating sets for G and H, respectively.

Definition 3.43 (Complexity)**.** Let $\phi : G \to H$. The *complexity* of ϕ with respect to the finite generating sets X and Y is

$$|\phi|_X^Y := \max_{x\in X} |\phi(x)|_Y.$$

The following proposition is straightforward to verify.

Lemma 3.44. *Let* $\phi : G \to H$ *and* $\theta : H \to K$ *and let* X*,* Y*, and* Z *be finite generating sets for* $G, H,$ *and* K*, respectively. Then*

$$|\theta \circ \phi|_X^Z \leq |\phi|_X^Y \cdot |\theta|_Y^Z.$$

Remark. Using the above convention, if X' and Y' are alternate finite generating sets for G and H, respectively, we have

$$|\phi|_{X'}^{Y'} \leq |\operatorname{Id}|_{X'}^{X} \cdot |\phi|_X^Y \cdot |\operatorname{Id}|_Y^{Y'}.$$

Since G is fully residually H, for every $R \in \mathbb{N}$, there is a homomorphism ϕ_R which H-discriminates the finite set $B_R(G,X) - 1$.

Definition 3.45 (Discriminating complexity)**.** Define a function $C^{H,Y}_{G,X} : \mathbb{N} \to \mathbb{N}$ via

$$C^{H,Y}_{G,X}(R) := \min\{|\phi|^Y_X \ : \ (\phi : G \to H) \text{ discriminates } (B_R(G,X) - 1)\}.$$

The function $C^{H,Y}_{G,X}$ so defined is called the *H-discriminating complexity of G with respect to finite generating sets X and Y.*

We will be interested in asymptotic classes of the discriminating complexity for a given group. To this end, if $f, g : \mathbb{N} \to \mathbb{N}$, we say that f is *asymptotically dominated by* g, denoted $f \preceq g$, if there is a constant K such that for all n,

$$f(R) \leq Kg(KR) + K.$$

We say that f is *asymptotically equivalent* to g, denoted $f \approx g$, if $f \preceq g$ and $g \preceq f$.

Lemma 3.44 and the remark following it imply the following proposition.

Proposition 3.46. *Let G be a fully residually H group. Let X, X' be finite generating sets for G, and let Y, Y' be finite generating sets for H. Then we have*

$$C^{H,Y}_{G,X} \preceq C^{H,Y'}_{G,X'}.$$

As a result of the above proposition, the asymptotic class of the H-discriminating complexity of G is invariant with respect to choice of finite generating set for both G and H. Therefore, we will omit reference to these generating sets and simply indicate (the asymptotic class of) the H-discriminating complexity of G by C^H_G.

In order to study H-discriminating complexity, we will find it useful to establish some notation for sequences of homomorphisms which discriminate larger and larger balls in a given group.

Definition 3.47 (Discriminating sequence)**.** Let $\Phi = (\phi_R : G \to H)_{R \in \mathbb{N}}$ be a sequence of homomorphisms. If for each $R \in \mathbb{N}$, the set $B_R(G,X) - 1$ is H-discriminated by ϕ_R, we say that Φ is a *H-discriminating sequence* with respect to the finite generating set X.

It is straightforward to see that a finitely generated group G is fully residually H if and only if G admits an H-discriminating sequence with respect to some (every) finite generating set.

We also make the following observation. Let X and X' be finite generating sets for G and let Φ be an H-discriminating sequence for G with respect to X. By passing to an arithmetic subsequence of Φ, we may obtain an H-discriminating sequence with respect to X', and the complexity of this subsequence is equivalent to that of Φ.

Definition 3.48 (Complexity function)**.** Given an H-discriminating sequence Φ, we construct the *H-discriminating complexity function associated to Φ,* the function $C_\Phi : \mathbb{N} \to \mathbb{N}$ defined via:

$$C_\Phi(R) := |\phi_R|_X^Y.$$

We briefly note that complexity functions of discriminating sequences provide an obvious upper bound for discriminating complexity.

Proposition 3.49. *Let G and H be finitely generated groups and let G be fully residually H. Let $\Phi = (\phi_R)_{R\in\mathbb{N}}$ be an H-discriminating sequence for G. Then $C_G^H \preceq C_\Phi$.*

3.2.2.1 Free Abelian groups

We begin by investigating the $\mathbb{Z}$-discriminating complexity of a free Abelian group $\mathbb{Z}^n$.

Proposition 3.50. *The $\mathbb{Z}$-discriminating complexity of $\mathbb{Z}^n$ is asymptotically dominated by a polynomial of degree $n-1$.*

We will consider the elements of $\mathbb{Z}^n$ to be n-tuples of integers. For $R \in \mathbb{N}$, define $[-R,R]^n := \{(t_1,\ldots,t_n) \in \mathbb{Z}^n : |t_i| \leq R,\ 1 \leq i \leq n\}$. Instead of discriminating closed balls in $\mathbb{Z}^n$ with respect to the usual metric, we will construct homomorphisms which are injective on the sets $[-R,R]^n$ for each $R \in \mathbb{N}$.

Lemma 3.51. *For $n, R \in \mathbb{N}$, define the homomorphism $\theta_{n,R} : \mathbb{Z}^n \to \mathbb{Z}$ by*

$$\theta_{n,R}(t_1,\ldots,t_n) = \sum_{i=1}^{n}(2R+1)^{i-1}t_i.$$

Then $\theta_{n,R}$ induces a bijection from $[-R,R]^n$ to the interval

$$I_{n,R} := \left[-\frac{1}{2}\left((2R+1)^n - 1\right), \frac{1}{2}\left((2R+1)^n - 1\right)\right].$$

Proof. We proceed by induction. Since $\theta_{1,R}$ is the identity for all R, we have the promised bijection for $n = 1$.

Fix r and assume that $\theta_{n,R}$ induces a bijection from $[-R,R]^n$ to $I_{n,R}$. Note that that

$$\theta_{n+1,R}(t_1,\ldots,t_{n+1}) = \theta_{n,R}(t_1,\ldots,t_n) + (2R+1)^n t_{n+1}.$$

By the inductive hypothesis, we have

$$\begin{aligned}\left|\theta_{n+1,R}(t_1,\ldots,t_{n+1})\right| &\le \left|\theta_{n,R}(t_1,\ldots,t_n)\right| + (2R+1)^n\left|t_{n+1}\right| \\ &\le \frac{1}{2}\big((2R+1)^n-1\big) + R(2R+1)^n \\ &= \frac{1}{2}(2R+1)^n + \frac{1}{2}2R(2R+1)^n - \frac{1}{2} \\ &= \frac{1}{2}\big((2R+1)^{n+1}-1\big).\end{aligned}$$

Therefore $\theta_{n+1,R}$ maps $[-R,R]^{n+1}$ into the interval $I_{n+1,R}$.

Suppose that there are $(s_1,\ldots,s_n),(t_1,\ldots,t_n)\in[-R,R]^{n+1}$ such that $\theta_{n+1,R}(t) = \theta_{n+1,R}(s)$. We then have

$$\theta_{n,R}(t_1,\ldots,t_n) + (2R+1)^n t_{n+1} = \theta_{n,R}(s_1,\ldots,s_n) + (2R+1)^n s_{n+1}.$$

We must have $t_{n+1}\neq s_{n+1}$ or we contradict the injectivity of $\theta_{n,r}$. However, by using the inductive hypothesis, we have

$$\begin{aligned}(2R+1)^n - 1 &\ge \left|\theta_{n,R}(t_1,\ldots,t_n) - \theta_{n,R}(s_1,\ldots,s_n)\right| \\ &= \left|(2R+1)^n(s_{n+1}-t_{n+1})\right| \\ &\ge (2R+1)^n,\end{aligned}$$

a contradiction.

We have shown that $\theta_{n+1,R}$ maps $[-R,R]^{n+1}$ injectively to $I_{n+1,R}$. Since both sets have the same cardinality, $\theta_{n+1,R}$ is a bijection between $[-R,R]^{n+1}$ and $I_{n+1,R}$. ∎

Proposition 3.50 follows immediately from Lemma 3.51 since each homomorphism $\theta_{n,R}$ is injective on B_R and therefore discriminates $B_R - 1$. Furthermore, the complexity of $\theta_{n,R}$ is $(2R+1)^{n-1}$, as promised.

The following result is well-known from number theory and will help us to establish a lower bound on the $\mathbb{Z}$-discriminating complexity of $\mathbb{Z}^n$.

Siegel's Lemma ([6, 49])**.** *Let A be an $M\times N$ integer matrix with $M>N$ and $A\neq 0$. Let B be a constant such that for every entry a_{ij} of A, we have $|a_{ij}|\le B$. Then there exists a nonzero $N\times 1$ integer matrix X with entries x_i such that $AX=0$ and for each i,*

$$|x_i| \le (NB)^{M/(N-M)}.$$

Corollary 3.52. *The $\mathbb{Z}$-discriminating complexity of $\mathbb{Z}^n$ asymptotically dominates a polynomial of degree $n-1$.*

Proof. Let $\Phi = (\phi_R)_{R\in\mathbb{N}}$ be a $\mathbb{Z}$-discriminating sequence for $\mathbb{Z}^n$. By definition, ϕ_R discriminates the set $B_R - 1$, the closed ball of radius R with respect to (WLOG) the standard basis of $\mathbb{Z}^n$.

Each ϕ_R can be represented by an $n \times 1$ integer matrix whose entries are bounded above in magnitude by $C_\Phi(R)$. By Siegel's lemma, there exists for each ϕ_R an element of the kernel of ϕ_R whose entries are bounded above in magnitude by $(nC_\Phi(R))^{1/(n-1)}$. Since ϕ_R discriminates $B_R - 1$, it also discriminates the set of nontrivial elements whose entries are bounded above in magnitude by $\lfloor R/n \rfloor$. We must then have

$$\begin{aligned}\frac{R}{n} - 1 \leq \left\lfloor \frac{R}{n} \right\rfloor &\leq (nC_\Phi(R))^{1/(n-1)} \\ \frac{(R-n)^{n-1}}{n^{n-1}} &\leq nC_\Phi(R) \\ \frac{(R-n)^{n-1}}{n^n} &\leq C_\Phi(R).\end{aligned}$$

Therefore $C_\Phi(R) \succeq R^{n-1}$.

In particular, taking Φ such that $C_\Phi(R) = C_G^\Gamma(R)$, we have that $C_G^\Gamma(R) \succeq R^{n-1}$. ☐

Theorem 3.53. *The $\mathbb{Z}$-discriminating complexity of $\mathbb{Z}^n$ is asymptotically equivalent to a polynomial of rank $n-1$.*

For $p \in \mathbb{Z}$, define a homomorphism $\theta^p_{n,R} : \mathbb{Z}^n \to \mathbb{Z}$ by

$$\theta^p_{n,R}(t_1, \ldots, t_n) := p\theta_{n,R}(t_1, \ldots, t_n).$$

Note that since $\theta_{n,R}$ discriminates the set $[-R, R]^n - 1$, if $i \in \theta^p_{n,R}([-R, R]^n - 1)$, then $|i| > |p|$. Clearly $\theta^p_{n,R}$ then also discriminates $[-R, R]^n - 1$.

3.2.2.2 Extensions of centralizers

Let Γ be a non-Abelian, torsion-free hyperbolic group. Let G be an iterated extension of centralizers over Γ with finite generating set X, and let $u \in G$ be a hyperbolic element which generates its own centralizer. Let G' be a rank n extension of the centralizer $C(u) = C_G(u)$. Fix elements $T = \{t_1, \ldots, t_n\} \subset G'$ be such that $\{u, t_1, \ldots, t_n\}$ is a basis for the free Abelian group $C_{G'}(u)$.

We define a homomorphism $\Theta^p_{n,R} : G' \to G$ via:

$$\begin{aligned}\Theta^p_{n,R}(g) &:= g \text{ for all } g \in G \\ \Theta^p_{n,R}(t_i) &:= u^{p(2R+1)^{i-1}} \text{ for } i = 1, \ldots, n.\end{aligned}$$

By putting T in bijection with the standard basis for $\mathbb{Z}^n$, it is clear that the homomorphism $\Theta^p_{n,R} \mid_{\langle T \rangle}$ is equivalent to $\theta^p_{n,R}$. Consequently, for all nontrivial $a \in \langle T \rangle$ is such that $|a|_T < R$, then $\Theta^p_{n,R}(a)$ is a power of u of exponent greater than or equal to p in magnitude. We further observe that $\Theta^p_{n,R}$ is a retraction onto G.

Lemma 3.54. *Let w be an element of G' with $|w|_{X \cup T} \leq R$. There is a linear function $N_3 : \mathbb{N} \to \mathbb{N}$ such that $\Theta^{N_3(R)}_{n,R}(w) \neq 1$.*

Proof. Since G' is an amalgamated product, we may write w as a geodesic $X \cup T$-word

$$w = g_0 a_0 g_1 a_1 \cdots g_k a_k g_{k+1} \tag{3.10}$$

where for each i, g_i is an X-word and a_i is a T-word. We may further assume that no g_i or a_i is the empty word, except possibly g_0, g_{k+1}, or both.

First, we may assume that if some g_i is not a power of u, then no g_i is a power of u. To see this, suppose that g_j is some power of u but g_{j-1} is not, and consider the subword $g_{j-1}a_{j-1}g_j a_j$. Since a_{j-1} is a word in the generators T, it represents an element of the centralizer of u. Consequently, we may rewrite this subword as $g_{j-1}g_j a_{j-1}a_j$ without increasing the $X \cup T$-length of the overall word. By replacing $g_{j-1}g_j$ and $a_{j-1}a_j$ with possibly shorter words representing the same elements, we obtain another word representing w in G' of length at most R.

Define

$$N_3(R) := N_2(R) + R + 1$$

and note that, because $N_2(R)$ is linear in R, the function $N_3(R)$ is also linear in R .

Consider the homomorphism $\Theta_{n,R}^{N_3(R)} : G' \to G$. Then

$$\begin{aligned}\Theta_{n,R}^{N_3(R)}(w) &= g_0 u^{r_0} g_1 u^{r_1} g_2 u^{r_2} \cdots g_k u^{r_k} g_{k+1} \\ &= g_0 w_{\mathbf{g}}(\mathbf{r}) g_{k+1},\end{aligned}$$

where $\mathbf{g} = (g_1, \ldots, g_k)$, $\mathbf{r} = (r_0, \ldots, r_k)$, $w_{\mathbf{g}}(\mathbf{r})$ is as in Equation 3.8 possibly g_0 or g_{k+1} or both are trivial. Since $|a_i|_T \leq |w|_{X\cup T} \leq R$, we have $\min(\mathbf{r}) > N_2(R)$ for all i. Since $\sum |g_i|_X \leq R$ and G is relatively hyperbolic with u a hyperbolic element generating its own centralizer, by Theorem 3.42 we have that $\Theta_{n,R}^{N_3(R)}(w) \neq 1$ in G.

Now suppose that w can be written as a geodesic $(X \cup T)$-word

$$w = u^{r_0} a_0,$$

where r_0 is an integer, a_0 is a nonempty T-word, and $|u^{r_0}|_X + |a_0|_T \leq R$. Since $|u|_X \geq 1$, $|r_0| \leq R$. By definition, $\Theta_{n,R}^{N_3(R)}(a) = u^e$ where $|e| > R$, and so $\Theta_{n,R}^{N_3(R)}(w) \neq 1$ in G. □

Theorem 3.55. *Let G be an iterated extension of centralizers over Γ. Let G' be a rank n extension of a cyclic centralizer of G. Then the G-discriminating complexity of G' is asymptotically dominated by a polynomial of degree n.*

Proof. By the previous theorem, the homomorphism $\Theta_{n,R}^{N_3(R)}$ maps all elements of G' with $X \cup T$-length at most R to nontrivial elements of G. Therefore, $\left(\Theta_{n,R}^{N_3(R)}\right)_{R\in\mathbb{N}}$ is a G-discriminating sequence for G'.

To compute the complexity of $\Theta_{n,R}^{N_3(R)}$, we first note that $\Theta_{n,R}^{N_3(R)}$ fixes elements of X. For $t_i \in T$, we have $\Theta_{n,R}^{N_3(R)}(t_i) = u^{(N_3(R))(2R+1)^{i-1}}$. Therefore, as a function of r,

$$|\Theta_{n,R}^{N_3(R)}| \leq |u|_X (N_3(R))(2R+1)^{n-1} \approx R^n,$$

since $N_2(R)$ is linear in R. Thus the complexity of the sequence $\left(\Theta_{n,R}^{N_3(R)}\right)_{R\in\mathbb{N}}$ is asymptotically dominated by R^n. □

3.2.2.3 Iterated extensions of centralizers

Theorem 3.56. *The Γ-discriminating complexity of an iterated extension of centralizers over Γ is asymptotically dominated by a polynomial with degree equal to the product of the ranks of the extensions.*

Proof. Let G be an iterated extension of centralizers over Γ, and let

$$\Gamma = G_0 \leq G_1 \leq \cdots \leq G_k = G$$

be a sequence such that G_i is an extension of a centralizer of G_{i-1} for $i = 1, \ldots, k$.

By Theorem 3.55, each G_i has a G_{i-1}-discriminating family with complexity polynomial of degree equal to the rank of the extension. By composing these families, we obtain a Γ-discriminating sequence for G which is also of polynomial complexity; in particular, the properties of complexity imply that the degree of the polynomial is equal to the product of the ranks of the extensions required to construct G. □

3.2.2.4 Arbitrary Γ-limit groups

Theorem 3.57. *The Γ-discriminating complexity of any Γ-limit group is asymptotically dominated by a polynomial.*

Proof. Let G be a Γ-limit group. By Proposition 3.9, there is a G' which is an iterated extension of centralizers over Γ such that $G \leq G'$. Choose a finite generating set X for G' which includes a finite generating set Y for G. Then for all $R \in \mathbb{N}$, we have $B_R(G, Y) \subseteq B_R(G', X)$, so a Γ-discriminating sequence exists for G' which is also a Γ-discriminating sequence for G. □

Lemma 3.58. *Let G be a Γ-limit group with a free Abelian subgroup of rank $n+1$. Then the Γ-discriminating complexity of G asymptotically dominates a polynomial of degree n.*

Proof. Since the asymptotic class of the complexity of a Γ-discriminating sequence is invariant with respect to choice of finite generating set, we may choose a generating set Y for G with a subset $T \subseteq Y$ such that $\langle T \rangle$ is free Abelian of rank $n+1$. Let $\Phi = (\phi_R)$ be a Γ-discriminating sequence for G. Since Γ is torsion-free hyperbolic, every Abelian subgroup of Γ is isomorphic to $\mathbb{Z}$, and therefore every ϕ_R must map $\langle T\rangle$ to a cyclic subgroup. Since $T \subseteq Y$, restricting Φ to $\langle T\rangle$ gives us a $\mathbb{Z}$-discriminating sequence for $\langle T\rangle \cong \mathbb{Z}^{n+1}$. Therefore, the complexity of Φ must asymptotically dominate a polynomial of degree n by Proposition 3.52. $\square$

References

[1] G. N. Arzhantseva. Generic properties of finitely presented groups and Howson's theorem. *Comm. Algebra*, 26(11):3783–3792, 1998.

[2] G. N. Arzhantseva. A property of subgroups of infinite index in a free group. *Proc. Amer. Math. Soc.*, 128(11):3205–3210, 2000.

[3] G. N. Arzhantseva and A. Yu. Ol′shanskiĭ. Generality of the class of groups in which subgroups with a lesser number of generators are free. *Mat. Zametki*, 59(4):489–496, 638, 1996.

[4] Benjamin Baumslag. Residually free groups. *Proc. London Math. Soc. (3)*, 17:402–418, 1967.

[5] M. Bestvina and M. Feighn. Outer limits (preprint), 1994.

[6] E. Bombieri and J. Vaaler. Addendum to: "On Siegel's lemma". *Invent. Math.*, 75(2):377, 1984.

[7] Francis Bonahon. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math. (2)*, 124(1):71–158, 1986.

[8] Francis Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.

[9] Francis Bonahon. Geodesic currents on negatively curved groups. In *Arboreal group theory (Berkeley, CA, 1988)*, volume 19 of *Math. Sci. Res. Inst. Publ.*, pages 143–168. Springer, New York, 1991.

[10] Alexandre V. Borovik, Alexei G. Myasnikov, and Vladimir Shpilrain. Measuring sets in infinite groups. In *Computational and statistical group theory (Las Vegas, NV/Hoboken, NJ, 2001)*, volume 298 of *Contemp. Math.*, pages 21–42. Amer. Math. Soc., Providence, RI, 2002.

[11] Khalid Bou-Rabee. Quantifying residual finiteness. *J. Algebra*, 323(3):729–737, 2010.

[12] Khalid Bou-Rabee. Approximating a group by its solvable quotients. arXiv:1102.4030v1 [math.GR], 2011.

[13] J. Burillo and E. Ventura. Counting primitive elements in free groups. *Geom. Dedicata*, 93:143–162, 2002.

[14] Christophe Champetier. Petite simplification dans les groupes hyperboliques. *Ann. Fac. Sci. Toulouse Math. (6)*, 3(2):161–221, 1994.

[15] Christophe Champetier. Propriétés statistiques des groupes de présentation finie. *Adv. Math.*, 116(2):197–262, 1995.

[16] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[17] François Dahmani. Combination of convergence groups. *Geom. Topol.*, 7:933–963 (electronic), 2003.

[18] S. M. Gersten. On Whitehead's algorithm. *Bull. Amer. Math. Soc. (N.S.)*, 10(2):281–284, 1984.

[19] Étienne Ghys. Groupes aléatoires (d'après Misha Gromov,...). *Astérisque*, (294):viii, 173–204, 2004.

[20] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[21] V. S. Guba. Conditions under which 2-generated subgroups in small cancellation groups are free. *Izv. Vyssh. Uchebn. Zaved. Mat.*, (7):12–19, 87, 1986.

[22] Vincent Guirardel. Approximations of stable actions on **R**-trees. *Comment. Math. Helv.*, 73(1):89–121, 1998.

[23] G. Christopher Hruska. Relative hyperbolicity and relative quasiconvexity for countable groups. *Algebr. Geom. Topol.*, 10(3):1807–1856, 2010.

[24] Sašo Kalajdžievski. Automorphism group of a free group: centralizers and stabilizers. *J. Algebra*, 150(2):435–502, 1992.

[25] Ilya Kapovich and Martin Lustig. Intersection form, laminations and currents on free groups. *Geom. Funct. Anal.*, 19(5):1426–1467, 2010.

[26] Ilya Kapovich, Alexei Miasnikov, Paul Schupp, and Vladimir Shpilrain. Generic-case complexity, decision problems in group theory, and random walks. *J. Algebra*, 264(2):665–694, 2003.

[27] Ilya Kapovich, Alexei Miasnikov, Paul Schupp, and Vladimir Shpilrain. Average-case complexity and decision problems in group theory. *Adv. Math.*, 190(2):343–359, 2005.

[28] Ilya Kapovich and Paul Schupp. Genericity, the Arzhantseva-Ol′shanskii method and the isomorphism problem for one-relator groups. *Math. Ann.*, 331(1):1–19, 2005.

[29] Ilya Kapovich, Paul Schupp, and Vladimir Shpilrain. Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups. *Pacific J. Math.*, 223(1):113–140, 2006.

[30] Martin Kassabov and Francesco Matucci. Bounding the residual finiteness of free groups. arXiv:0912.2368v1 [math.GR], 2009.

[31] O. Kharlampovich and A. Myasnikov. Limits of relatively hyperbolic groups and lyndon's completions. arXiv:0904.2423v3 [math.GR], 2009.

[32] A. G. Myasnikov and V. N. Remeslennikov. Exponential groups. II. Extensions of centralizers and tensor completion of CSA-groups. *Internat. J. Algebra Comput.*, 6(6):687–711, 1996.

[33] Walter D. Neumann and Michael Shapiro. Automatic structures, rational growth, and geometrically finite hyperbolic groups. *Invent. Math.*, 120(2):259–287, 1995.

[34] Y. Ollivier. Sharp phase transition theorems for hyperbolicity of random groups. *Geom. Funct. Anal.*, 14(3):595–679, 2004.

[35] Yann Ollivier. Critical densities for random quotients of hyperbolic groups. *C. R. Math. Acad. Sci. Paris*, 336(5):391–394, 2003.

[36] Yann Ollivier. Cogrowth and spectral gap of generic groups. *Ann. Inst. Fourier (Grenoble)*, 55(1):289–317, 2005.

[37] Yann Ollivier. Effondrement de quotients aléatoires de groupes hyperboliques avec torsion. *C. R. Math. Acad. Sci. Paris*, 341(3):137–140, 2005.

[38] Yann Ollivier. *A January 2005 invitation to random groups*, volume 10 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro, 2005.

[39] A. Yu. Ol′shanskiĭ. Almost every group is hyperbolic. *Internat. J. Algebra Comput.*, 2(1):1–17, 1992.

[40] A. Yu. Ol′shanskiĭ. On residualing homomorphisms and G-subgroups of hyperbolic groups. *Internat. J. Algebra Comput.*, 3(4):365–409, 1993.

[41] D. Osin. Asymptotic dimension of relatively hyperbolic groups. *Int. Math. Res. Not.*, (35):2143–2161, 2005.

[42] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.

[43] Patrick Reynolds. Dynamics of irriducible endomorphisms of f_n (preprint).

[44] Abdó Roig, Enric Ventura, and Pascal Weil. On the complexity of the Whitehead minimization problem. *Internat. J. Algebra Comput.*, 17(8):1611–1634, 2007.

[45] Zlil Sela. Diophantine geometry over groups. I. Makanin-Razborov diagrams. *Publ. Math. Inst. Hautes Études Sci.*, (93):31–105, 2001.

[46] Zlil Sela. Diophantine geometry over groups. VII. The elementary theory of a hyperbolic group. *Proc. Lond. Math. Soc. (3)*, 99(1):217–273, 2009.

[47] Jean-Pierre Serre. *Trees*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[48] Abe Shenitzer. Decomposition of a group with a single defining relation into a free product. *Proc. Amer. Math. Soc.*, 6:273–279, 1955.

[49] C.L. Siegel. Über einige Anwendungen diophantischer Approximationen. *Abh. der Preuß. Akad. der Wissenschaften. Phys.-Math. Kl.*, Nr. 1, 1929. (=*Ges. Abh.,* I, 209-266).

[50] Brent B. Solie. Genericity of filling elements. arXiv:1007.4022v1 [math.GR], 2010.

[51] J. R. Stallings. Free groups which are free products with cyclic amalgamations. *Notices A.M.S.*, 1(1):49, 1980.

[52] G. A. Swarup. Decompositions of free groups. *J. Pure Appl. Algebra*, 40(1):99–102, 1986.

Index

TS', 28
X-digraph, 8
Γ-limit group, 32
$\mathbb{R}$-tree, 14

Automorphic subgroup problem, 16

Cayley graph, 31
complexity function, 48
complexity of a homomorphism, 46
CSA group, 33
cyclic splitting, 15

discriminating complexity, 47
discriminating homomorphism, 31
discriminating sequence, 47
dual digraph, 9
dual tree, 1

extension of a centralizer, 32

fellow traveling, 36
filling curve, 1
filling element, 14
filling subgroup, 14
fully residually P, 5

genericity, 28
geometric intersection number, 1

hypergraph, 11

iterated extension of centralizers, 32

Property (L), 26
Property (S), 20

quasi-isometric embedding, 37

relative Dehn function, 34
relative generating set, 33
relative isoperimetric function, 34
relative presentation, 33
relatively hyperbolic group, 34
relatively quasiconvex, 36
residual property, 31
residually P, 5

Stallings graph, 10
strongly relatively quasiconvex, 36
surface group, 1

translation length, 1

vertex subgroup
    loop, 16
    segment, 16

Whitehead automorphism, 10
Whitehead hypergraph, 11