ON RESIDUAL PROPERTIES OF GROUPS AND DEHN FUNCTIONS FOR MAPPING TORI OF RIGHT ANGLED ARTIN GROUPS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Kristen Pueschel

August 2016

ON RESIDUAL PROPERTIES OF GROUPS AND DEHN FUNCTIONS FOR MAPPING TORI OF RIGHT ANGLED ARTIN GROUPS

Kristen Pueschel, Ph.D.

Cornell University 2016

Given a group G and an automorphism $\phi : G \to G$, the algebraic mapping torus M_ϕ, is the HNN-extension of G where the stable letter conjugates g to $\phi(g)$. In this thesis, we study mapping tori with base groups G that are right angled Artin groups on few generators. In particular, we examine how the Dehn function of a mapping torus depends on the automorphism ϕ used to generate the group when G is $F_2 \times \mathbb{Z}$ or $\mathbb{Z}^2 * \mathbb{Z}$. We then extend our methods to produce bounds for the Dehn functions of other mapping tori with few-generator base groups.

A group G is residually finite if for every $g \in G - \{1\}$, there is a finite quotient of G in which the image of g is non-trivial. The hydra are a family of groups with fast-growing Dehn functions; the Dehn function of each group is equivalent to an Ackermann function. In this thesis we show that hydra are not residually finite, answering a question of Kharlampovich, Myasnikov, and Sapir.

A group G is residually solvable if for every $g \in G - \{1\}$, there is a solvable quotient of G in which the image of g is non-trivial. In this thesis we introduce and explore properties of the function that measures the smallest derived length of a solvable quotient in which g survives, in terms of the length of g.

BIOGRAPHICAL SKETCH

Kristen Pueschel grew up in Vestal, NY. She attended the University of Pittsburgh where she majored in Mathematics. Kristen did her graduate work at Cornell University, under the advisement of Timothy Riley. She is moving to Fayetteville to be a Visiting Assistant Professor at the University of Arkansas.

To Andrew, my partner.

ACKNOWLEDGEMENTS

Giant thanks to my advisor Tim Riley, for his ever enthusiastic support. I am particularly grateful for the fun of doing math together, and your patience along the way. To the past and present members of the Topology and Geometric Group Theory seminar, thank you. This has been an important community for me while at Cornell. Thank you to my committee members, Ken Brown and Jason Manning. Thank you to Tim and Tara for always hosting, to Kelly and Jason for your help along the way, to Margarita, Teddy, Harry, and Juan, who have been my mathematical family, and to all of the good friends who have come and gone. To my family, your support of my pursuits, academic and otherwise, has been a great gift over the years. Finally, thank you to my partner, Andrew, who has believed in me when I didn't, learned what Dehn functions were so I could talk about work, and agreed to getting Ruby-cat.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Part I

Dehn Functions of Mapping Tori

(with Timothy Riley)

CHAPTER 1

INTRODUCTION

In this chapter we review the history of classifications of geometric properties of mapping tori based on the maps used to define them. We begin with Nielsen–Thurston's classification of surface homeomorphisms, which is the paradigm for this study, and then consider classifications for the Dehn functions of algebraic mapping tori with free abelian and free bases. We introduce the right angled Artin groups, which include both free and free abelian groups, and consider common properties of their outer automorphism groups. Finally, we present classifications for the Dehn functions of mapping tori with 3-generator right angled Artin group bases. We finish with a set of questions about ascending HNN extensions of right angled Artin groups.

1.1 Topological mapping tori and the Nielsen–Thurston classification of surface automorphisms

Definition 1.1.1. Given a topological space X and a homeomorphism $f : X \to X$, the topological mapping torus associated to f is the space

$$M_f = \frac{X \times [0,1]}{(x,1) \sim (f(x),0)}.$$

The Nielsen–Thurston classification of surface homeomorphisms describes the geometry of a mapping torus M_f from properties of f.

Theorem 1 (Thurston [58]). *If Σ is a closed surface, and f is a homeomorphism of Σ, then up to isotopy, f is at least one of the following:*

1. *reducible: f is invariant on a finite set of disjoint simple closed curves. The mapping torus contains an incompressible torus.*

2. *periodic: some power of f is isotopic to the identity map. The mapping torus admits an $\mathbb{H}^2 \times \mathbb{R}$ structure.*

3. *pseudo-Anosov: f admits a pair of transverse measured foliations. Away from a set of isolated singularities, f stretches by λ in the direction of one of the foliations, and compresses by λ in the direction of the other. The mapping torus admits a hyperbolic structure.*

This kind of classification motivates our study. We wish to find such a description for geometric features of an algebraic analogue of mapping tori in terms of properties of the algebraic analogue of the homeomorphism.

1.2 Algebraic mapping tori and their Dehn functions

Let X be a CW-complex with homeomorphism $f : X \to X$, and M_f the topological mapping torus. Then f induces $f^* : \pi_1(X) \to \pi_1(X)$, and the fundamental group of M_f is $\pi_1(M_f) = \langle \pi_1(X),\ t \ \mid x^t = f^*(x)\rangle$. This group, which is an HNN-extension where conjugation sends elements of the base to their image under f^*, is the motivating example which is generalized by the following definition:

Definition 1.2.1. Given a group G and an injective endomorphism $\Phi : G \to G$, the *algebraic mapping torus* associated to Φ is the group

$$M_\Phi = \langle G, t \ \mid \ g^t = \Phi(g),\ \forall g \in G\rangle.$$

If Φ is an automorphism, $M_\Phi = G \rtimes_\Phi \mathbb{Z}$.

In this thesis, we restrict our attention to the mapping tori of automorphisms. The next three classifications describe the Dehn function of an algebraic mapping torus M_Φ, in terms of properties of the automorphism Φ.

1.2.1 Dehn functions of mapping tori with base $\mathbb{Z}^k$

The Dehn functions of mapping tori with base $\mathbb{Z}^k$ are classified as follows:

Theorem 2. *Suppose $\Phi \in \mathrm{Aut}(\mathbb{Z}^k) = \mathrm{GL}(k, \mathbb{Z})$. Then either:*

1. *Not all eigenvalues of Φ are unit. Then the Dehn function of the mapping torus M_Φ is exponential; or*
2. *All eigenvalues of Φ are unit. If the largest Jordan Block of the Jordan Canonical Form associated to Φ is size $c \times c$, then the Dehn function of M_Φ is a polynomial of degree $c + 1$.*

Theorem 2 was proved in two parts. In [30], Bridson and Pittet proved the upper bounds on the Dehn function of M_Φ. In [28], Bridson and Gersten proved the lower bounds.

To establish the upper bounds, Bridson and Pittet construct a combing for $\mathbb{Z}^k \rtimes \mathbb{Z}$ from a combing on $\mathbb{Z}^k$ that approximates straight lines in $\mathbb{R}^k$. This is of the form $t^k u$, where u is a combing on $\mathbb{Z}^k$. Suppose two combing lines end distance one apart. If they are both of the form $t^k u$ and $t^k v$, then the two combing lines never get far apart. When they differ, the combing lines follow lines in $\mathbb{R}^n$. Otherwise, the elements are of the form $t^k u$ and $t^{k\pm 1} v = t^k u t^{\pm 1}$, where $v = \Phi^{\pm 1}(u)$. Since Φ is a linear map, it sends straight lines to straight lines, so each vertex in the path in level k of the Cayley graph is only a bounded distance to the path in level $k+1$, and vice-versa. Thus the combing lines in $\mathbb{Z}^k \rtimes \mathbb{Z}$ asynchronously fellow travel. Bridson and Pittet bound the length of these combing lines, $L(n)$, by $||\Phi^n|| \sim n^c$. Using the following lemma, they obtain the desired upper bound:

Lemma (Bridson [25])**.** *If there exists a combing for the group Γ whose asynchronous width is bounded by a constant and whose length is bounded by the function $L(n)$, then the Dehn function for any finite presentation of Γ is bounded above by $\delta(n) \preceq nL(n)$.*

To establish the lower bounds on the Dehn function of $\mathbb{Z}^n \rtimes_\Phi \mathbb{Z}$, Bridson and Gersten explicitly construct families of words $(w_n)_{n\in\mathbb{N}}$ such that all van Kampen diagrams for w_n have area bounded below by An^{c+1}: In the presentation $\langle x_1, \ldots, x_k, t \;\mid\; x_i^t = \Phi(x_i), [x_i, x_j] = 1\rangle$, if Φ has a non-unit eigenvalue, they consider $w_n = x_{i(n)}^n t^{2n} x_{j(n)}^n t^{-2n} x_{i(n)}^{-n} t^{2n} x_{j(n)}^{-n} t^{-2n}$, where $i(n)$ and $j(n)$ are chosen to maximize $|\Phi^n(x_{i(n)})|$ and $|\Phi^{-n}(x_{j(n)})|$, respectively. Each t-edge in the boundary is connected by a t-corridor to another t-edge in the boundary, but there are few choices for ending edges. The t-edges are clustered together into groups of $2n$. The length of the n-th t-corridor in a segment is approximately $\max|\Phi^n(x_i)| \asymp \lambda^n$. This forces an exponential area. In the case that Φ has only unit eigenvalues, more care is needed. Bridson and Gersten consider a different family of words, and instead of finding lower bounds for the lengths of individual t-corridors, they find lower bounds for the sum of lengths of many t-corridors at a time. In Section 5.2, we adapt their proof to find lower bounds for the Dehn functions of mapping tori with base group $F_k \times F_l$ for $k, l > 1$.

1.2.2 Dehn functions of mapping tori arising from surface homeomorphisms

The Nielsen–Thurston classification for surface homeomorphisms implies that the mapping tori of pseudo-Anosov maps are always hyperbolic, and thus have linear Dehn functions. Epstein and Thurston proved that the fundamental groups of mapping tori of surface homeomorphisms are automatic for surfaces with genus at least two, and hence satisfy a quadratic isoperimetric inequality.

Theorem (Epstein–Thurston [48])**.** *Suppose f^* is induced by a homeomorphism of the closed surface Σ with $\xi(\Sigma) < 0$. If f is psuedo-Anosov, then M_{f^*} has linear Dehn function. Otherwise, M_{f^*} has quadratic Dehn function.*

1.2.3 Dehn functions of mapping tori with base F_k

Let F_k be the rank-k free group. An automorphism Φ of F_k is *atoroidal* when there is no conjugacy class $[w]$ such that $[\Phi^n(w)] = [w]$ for some n. The following theorem is proved using Bestvina–Feighn's Combination Theorem, which gives conditions for HNN-extensions and amalgamated products of hyperbolic groups to be hyperbolic.

Theorem (Bestvina–Feighn, Gersten [12, 13, 54])**.** *Suppose $M_\Phi = F_k \rtimes_\Phi \mathbb{Z}$ is hyperbolic. Then Φ is atoroidal.*

Brinkmann used the train track technology developed by Bestvina and Handel to show the converse statement. With this technology, he was able to find a representative of the same outer automorphism class as the atoroidal automorphism Φ, such that Φ satisfies a flaring condition that makes M_Φ hyperbolic.

Theorem (Brinkmann [31]). *Suppose $\Phi \in \mathrm{Aut}(F_k)$. If Φ is atoroidal then the mapping torus $M_\Phi = F_k \rtimes_\Phi \mathbb{Z}$ is hyperbolic.*

An automorphism $\Phi : F_k \to F_k$ has *polynomial growth* if the function $n \mapsto \max\{|\Phi^n(x_1)|, \ldots, |\Phi^n(x_k)|\}$ is equivalent to a polynomial. In her thesis, Macura established an upper bound on the Dehn functions of mapping tori M_Φ when the automorphism Φ has polynomial growth.

Theorem (Macura [73]). *If $\Phi : F_k \to F_k$ is an automorphism with polynomial growth then M_Φ satisfies a quadratic isoperimetric inequality.*

In a series of three papers, which make up [29], Bridson and Groves completed the classification of Dehn functions of mapping tori.

An automorphism Φ of $F_k = \langle a_1, \ldots, a_k \rangle$ is *positive* if $\Phi(a_i) = u_i$ where u_i is a word in the generators $\{a_1, \ldots a_k\}$, but not their inverses.

Theorem (Bridson–Groves [29]). *If $\Phi \in \mathrm{Aut}(F_k)$ is positive and is not atoroidal, then $M_\Phi = \langle a_1, \ldots a_k, t \mid a_i^t = \Phi(a_i) \rangle$ has quadratic Dehn function.*

Bridson and Groves argued that the Dehn function had a quadratic upper bound with an amazing hands-on analysis of the area of van Kampen diagrams. The van Kampen diagrams of M_Φ are unions of t-corridors. Their elaborate proof bounds the length of t-corridors in terms of the length of the boundary word in a minimal area van Kampen diagram. Using train-track technology, Bridson and Groves were able to generalize their arguments from the case of positive automorphisms to prove the theorem for arbitrary automorphisms.

Theorem (Bridson–Groves [29]). *If $\Phi \in \mathrm{Aut}(F_k)$ is not atoroidal, then $M_\Phi = \langle a_1, \ldots a_k, t \mid a_i^t = \Phi(a_i) \rangle$ has quadratic Dehn function.*

Together these results yield the following classification of the Dehn functions of mapping tori with free group base.

Theorem 3. *Suppose* $\Phi \in \mathrm{Aut}(F_k)$*. Then either:*

1. Φ *is atoroidal and* M_Φ *has linear Dehn function; or*
2. Φ *is not atoroidal and* M_Φ *has quadratic Dehn function.*

1.3 Right angled Artin groups

Free groups and free abelian groups are the most basic examples of the following construction:

Definition 1.3.1. If $\Gamma = (V, E)$ is a finite simplicial graph, the group

$$A(\Gamma) := \langle V \;\mid [x, y] = 1 \text{ for } (x, y) \in E\rangle.$$

The group $A(\Gamma)$ is called a right angled Artin group (RAAG).

The subgroups of RAAGs have long been a source of interesting examples. For example, the Bieri–Stallings group, G_n, for $n \geq 3$, is the kernel of the map from $\oplus_{i=1}^{n} F(a_i, b_i) \to \mathbb{Z}$, given by $a_i, b_i \mapsto 1$ for all i. These groups are of type F_{n-1} but not F_n. Bestvina–Brady generalized this construction from $F_2 \times \cdots \times F_2$ to all RAAGs $A(\Gamma)$. The Bestvina–Brady group for $A(\Gamma)$ is the kernel of the map $A(\Gamma) \to \mathbb{Z}$ that sends each generator to 1. These subgroups were used to establish that the property type F_n, having a $K(G, 1)$ with a finite n-skeleton, is different from the property FP_n, admitting a projective resolution of length n. In fact they found examples that are not finitely presentable, but are FP_n for every n

[11]. Brady and Soroko have recently shown that for every natural number k, there is a RAAG $A(\Gamma_k)$ and a subgroup $H_k \leq A(\Gamma_k)$ such that H_k has Dehn function $\delta_{H_k}(n) \cong n^k$. In work that established the Virtual Haken Conjecture, Agol proved that the fundamental groups of closed hyperbolic 3-manifolds virtually quasi-isometrically embed as subgroups of RAAGs [1].

1.3.1 Commonality in Out$(A(\Gamma))$

The idea that RAAGs interpolate between free and free-abelian groups has been particularly important in the study of $\mathrm{Out}(A(\Gamma))$. Although the mapping class groups, $\mathrm{MCG}(\Sigma)$, are not of the form $\mathrm{Out}(A(\Gamma))$, they have important analogous behavior to both $\mathrm{GL}(n, \mathbb{Z})$ and $\mathrm{Out}(F_n)$. Many of the properties below were first proved for $\mathrm{GL}(n, \mathbb{Z})$, then for $\mathrm{MCG}(\Sigma)$. The $\mathrm{MCG}(\Sigma)$ proofs provided inspiration for proofs in $\mathrm{Out}(F_n)$, and in turn these inspired proofs for $\mathrm{Out}(A(\Gamma))$.

$\mathrm{GL}(n, \mathbb{Z})$, $\mathrm{MCG}(\Sigma)$, $\mathrm{Out}(F_n)$, and recently $\mathrm{Out}(A(\Gamma))$, have been shown to be virtually torsion free and residually finite. All have finite virtual cohomological dimension [59, 40, 37, 38].

The Tits alternative, that every subgroup H is virtually solvable or contains a non-abelian free group, holds for $\mathrm{GL}(n, \mathbb{Z})$ [89] and for $\mathrm{MCG}(\Sigma)$ by Ping-Pong arguments [76]. Bestvina, Feighn, and Handel showed the Tits alternative for $\mathrm{Out}(F_n)$ [14, 15]. Charney and Vogtmann have shown that the Tits alternative for subgroups holds too for $\mathrm{Out}(A_\Gamma)$ for arbitrary RAAG $A(\Gamma)$ [37].

The group $\mathrm{GL}(n, \mathbb{Z})$ acts properly on the symmetric space $\mathrm{SO}(n)\backslash\mathrm{SL}(n, \mathbb{R})$, the space of marked flat tori. The group $\mathrm{MCG}(\Sigma)$ acts properly on Teichmuller space, and $\mathrm{Out}(F_n)$ acts properly on Culler–Vogtmann Outer Space, a space of marked metric graphs. Charney and Vogtmann have constructed an analogue of Outer Space for a subgroup, $\mathrm{Out}^u(A(\Gamma))$ of $\mathrm{Out}(A(\Gamma))$. Each of these spaces is contractible. All of these spaces have a spine onto which the space deformation retracts, which is preserved by the action of the group [4, 61, 40, 36]. The spines in $\mathrm{Out}(F_n)$ and $\mathrm{Out}^u(A(\Gamma))$ can be used to calculate upper bounds for the virtual cohomological dimensions of these groups.

The study of algebraic mapping tori provides a way to study individual elements of $\mathrm{Aut}(A(\Gamma))$. The commonality above suggests that the classifications of Sections 1.2.1 and 1.2.3 might fit into a general classification across all RAAGs.

Open question 1.3.2. Are the Dehn functions of mapping tori of RAAGs always equivalent to polynomial and exponential functions?

Open questions 1.3.3. How are the Dehn functions of mapping tori of RAAGs characterized and classified by the automorphisms Φ? What characteristics of Φ determine the Dehn function?

1.4 Our results

In this thesis we complete the three-generator case of Questions 1.3.2 and 1.3.3. The three-generator RAAGs are $\mathbb{Z}^3$, F_3, $F_2 \times \mathbb{Z}$ and $\mathbb{Z}^2 * \mathbb{Z}$. For $\mathbb{Z}^3$ and F_3 (and for RAAGs on one or two generators), the theorems of Sections 1.2.1 and 1.2.3 classify the Dehn functions of M_Φ. Here are our classifications for $F_2 \times \mathbb{Z}$ and $\mathbb{Z}^2 * \mathbb{Z}$:

Theorem 4. *Suppose $\Phi \in \mathrm{Aut}(F_2 \times \mathbb{Z})$. Let ϕ be the automorphism of F_2 induced by Φ via the map killing the $\mathbb{Z}$ factor. Let $\phi_{ab} \in \mathrm{Aut}(\mathbb{Z}^2)$ be the automorphism induced by the abelianization map $F_2 \to \mathbb{Z}^2$, which sends $g \mapsto g_{ab}$. Let $p : F_2 \times \mathbb{Z} \to \mathbb{Z}$ be projection to the $\mathbb{Z}$ factor. Exactly one of the following holds:*

1. *ϕ_{ab} only has unit eigenvalues and there exists $g \in F_2 \times \{1\}$ and $m \in \mathbb{N}$ such that $\phi_{ab}^m(g_{ab}) = g_{ab}$ and $p(\Phi^m(g)) \neq 0$. Then M_Φ has cubic Dehn function.*

2. *M_Φ has quadratic Dehn function.*

Theorem 5. *Suppose $\Phi \in \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$. Then Φ^2 is conjugate to Ψ which preserves the free abelian factor, and so defines some $\psi \in \mathrm{GL}(2, \mathbb{Z})$. Exactly one of the following holds:*

1. *ψ is finite order. Then M_Φ has quadratic Dehn function.*

2. *ψ has a non-unit eigenvalue. Then M_Φ has exponential Dehn function.*

3. *M_Φ has cubic Dehn function.*

1.5 Other properties of mapping tori

This section is intended as supplementary introductory material on mapping tori of automorphisms and mapping tori of injective automorphisms. We will not discuss this material further but it would be an interesting avenue to pursue in the future.

For G free and free-abelian, for every $\Phi \in \mathrm{Aut}(G)$, M_Φ is residually finite — for every element $g \in M_\Phi$, there is $f : M_\Phi \to Q$ where Q is a finite group, such that $f(g) \neq 1$. This follows from the group splitting: $1 \to G \to M_\Phi \to \mathbb{Z} \to 1$. Mapping tori of RAAGs are also always residually finite for the same reason.

For G free and free-abelian, for every $\Phi \in \mathrm{Aut}(G)$, M_Φ is coherent — every finitely generated $H < G$ is finitely presentable. This is proved by Feighn and Handel in [50]. This is not true in general for RAAGs. For example, if the graph Γ contains an induced subgraph with four vertices that is graph isomorphic to a square, then $A(\Gamma)$ contains an $F_2 \times F_2$ subgroup. The subgroup $G_2 = \ker(F_2 \times F_2 \to \mathbb{Z})$ is a finitely generated but not finitely presentable subgroup of $F_2 \times F_2$. Thus if Γ contains an induced square, for every automorphism $\Phi \in \mathrm{Aut}(A(\Gamma))$, M_Φ will not be coherent. This leads to the question: If the RAAG $A(\Gamma)$ is coherent, will the mapping tori M_Φ of $A(\Gamma)$ be coherent for all $\Phi \in \mathrm{Aut}(A(\Gamma))$?

Relatively little is known about the mapping tori of injective endomorphisms with right angled Artin group bases, even in the case of free group bases. In [50], Feighn and Handel proved that the mapping tori of injective free group endomorphisms of free groups are coherent. In [65], Kapovich shows

Theorem (I. Kapovich)**.** *If $\phi : F_n \to F_n$ is an endomorphism such that $\phi(x)$ begins and ends with x, and $\phi(x) \neq x$ for all generators x, then M_ϕ is either hyperbolic, and has linear Dehn function, or the Dehn function is exponential and M_ϕ contains a Baumslag-Solitar subgroup of the form $BS(1, p)$.*

In [52], Geoghegan, Mihalik, Sapir, and Wise proved that the mapping tori of injective free group endomorphisms are Hopfian, and they conjectured that these groups are actually residually finite. In [17, 18], Borisov and Sapir were able to prove this and more; using algebraic geometry over groups, they showed that if the base group G is linear (for example if G is a RAAG) and $\phi : G \to G$ is an injective endomorphism, then M_ϕ is residually finite. Although they are always residually finite, mapping tori of free group endomorphisms are not always linear. For example, Drutu and Sapir [45] proved that groups of the form $\langle a, b, t \mid a^t = a^k, b^t = b^l \rangle$ are not linear when $k, l \neq \pm 1$, while Button and R. Kropholler recently showed that Sapir's group $\langle a, b, t \mid a^t = ab, b^t = ba \rangle$ is linear [33].

1.6 Outline of Part I

Chapter 2 introduces the basics of Dehn functions, mapping tori, and right angled Artin groups. Chapter 3 is devoted to the proof of Theorem 4. Chapter 4 is devoted to the proof of Theorem 5. In Chapter 5 we examine further examples and some families of RAAG bases that admit few automorphisms.

CHAPTER 2
PRELIMINARIES

In this section, we begin with the basics of the RAAG base groups, focusing on their automorphisms. We introduce the Dehn function, discussing basic properties of this function and giving a few examples. For the rest of the chapter we explain some of the techniques we will use to calculate Dehn functions.

Throughout this thesis, we write $|w|$ to denote the length of a word w. Our conventions are $a^t := t^{-1}at$ and $[a, b] := a^{-1}b^{-1}ab$.

2.1 Right angled Artin group basics

Definition 2.1.1. If $\Gamma = (V, E)$ is a finite simplicial graph, define the group

$$A(\Gamma) := \langle V \ \mid [x, y] = 1, \ \forall \ (x, y) \in E \rangle.$$

The groups defined by this construction are called *right angled Artin groups* (RAAGs), *partially commutative groups* and *graph groups.*

The free group F_n is $A(\Gamma)$ for Γ the graph with n vertices and no edges. The free abelian group $\mathbb{Z}^n$ is $A(\Gamma)$ for Γ the complete graph on n vertices. The group $L = \langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1 \rangle$, which will be discussed further in Part 2, Chapter 1, is also a RAAG.

Droms showed that non-isomorphic graphs produce non-isomorphic RAAGs:

Proposition 2.1.2 (Droms [44])**.** *The group $A(\Gamma_1) \cong A(\Gamma_2)$ if and only if Γ_1 and Γ_2 are graph isomorphic.*

Example 2.1.3. We list graphs with up to four vertices, up to graph isomorphism, and the corresponding right angled Artin groups:

	$\mathbb{Z}$	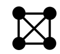	$\mathbb{Z}^4$
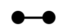	$\mathbb{Z}^2$	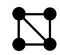	$\mathbb{Z}^3 *_{\mathbb{Z}^2} \mathbb{Z}^3$
	F_2	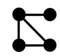	$\mathbb{Z}^3 *_{\mathbb{Z}} \mathbb{Z}^2$
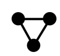	$\mathbb{Z}^3$	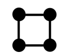	$F_2 \times F_2$
	$\mathbb{Z}^2 * \mathbb{Z}$	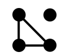	$\mathbb{Z}^3 * \mathbb{Z}$
	$F_2 \times \mathbb{Z}$	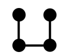	L
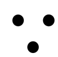	F_3	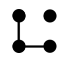	$(F_2 \times \mathbb{Z}) * \mathbb{Z}$
		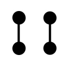	$\mathbb{Z}^2 * \mathbb{Z}^2$
		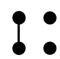	$\mathbb{Z}^2 * F_2$
		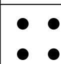	F_4

Table 2.1: Graphs with ≤ 4 vertices and their associated RAAGs

Recall the following graph-theoretic vocabulary:

Definition 2.1.4. Let Γ be a simplicial graph and v a vertex in Γ. Then $\text{link}(v)$ is the set of neighbors of v, i.e. the vertices sharing an edge with v, and $\text{star}(v) := \{v\} \cup \text{link}(v)$. If we assign a length of 1 to each edge, then $\text{star}(v)$ is the closed ball of radius 1 in Γ.

Proposition 2.1.5. *The group $A(\Gamma)$ is isomorphic to $A * B$ for some A, $B \neq 1$, if and only if Γ has multiple components. The group $A(\Gamma)$ is isomorphic to $A \times B$ for some $A, B \neq 1$, if and only if the vertices of Γ can be partitioned into two non-empty sets, V_A and V_B, and every vertex $v \in V_A$ has $V_B \subset \text{link}(v)$ and vice-versa.*

RAAGs are CAT(0) groups. This implies for example that when $A(\Gamma)$ is not a free group, that it has quadratic Dehn function. Further, they are algorithmically well-behaved:

Proposition 2.1.6 (VanWyk [91])**.** *RAAGs are biautomatic.*

If $\Theta \subset \Gamma$ is an *induced subgraph* of Γ, that is, all edges in Γ between vertices of Θ are edges in Θ, then $A(\Theta) \hookrightarrow A(\Gamma)$. Such an $A(\Theta)$ is called a *special subgroup*.

Proposition 2.1.7. *The special subgroups quasi-isometrically embed in $A(\Gamma)$.*

Although we work primarily with 3-generator RAAGs, where it is easy to work out the automorphisms, we state here Laurence and Servatius's generating set for $\mathrm{Aut}(A(\Gamma))$ [70, 87]. The automorphism group $\mathrm{Aut}(A(\Gamma))$ is finitely generated by:

1. Inner automorphisms.
2. Inversions: Send $v \to v^{-1}$ where v is a standard generator. Fix the other generators.
3. Partial conjugations: If $\Gamma - \mathrm{star}(v)$ is disconnected, conjugate the vertices of one of the resulting components by v. Fix the other generators.
4. Twists: If $\mathrm{star}(x) \subset \mathrm{star}(y)$, then $x \mapsto xy = yx$. Fix the other generators.
5. Folds: If $\mathrm{link}(x) \subset \mathrm{link}(y)$, then $x \mapsto xy \neq yx$. Fix the other generators.
6. Graph symmetries.

Twists and folds both satisfy the condition that $\mathrm{link}(x) \subset \mathrm{star}(y)$, that is, for $z \neq x$, if $[x, z] = 1$ then $[y, z] = 1$, and both have the form $x \mapsto xy$, with all other generators fixed. These two kinds of automorphisms are often lumped together and referred to as *transvections*.

Recalling that $\mathrm{Out}(G) = \mathrm{Aut}(G)/\mathrm{Inn}(G)$, we get the following trivial corollary from Servatius's generating set for $\mathrm{Aut}(A(\Gamma))$.

Corollary 2.1.8. *If Γ is such that for all vertices v and w, $\Gamma - \mathrm{star}(v)$ is connected, and $\mathrm{link}(v) \subset \mathrm{star}(w)$ implies that $v = w$, then $\mathrm{Out}(A(\Gamma))$ is finite.*

Example 2.1.9. If C_n is a cycle of length at least 5, then $\mathrm{Out}(A(C_n))$ is finite.

In [41], Day defined a finite family of Whitehead automorphisms for $\mathrm{Aut}(A(\Gamma))$ and used Peak-Reduction like arguments to establish that $\mathrm{Aut}(A(\Gamma))$ is always finitely presentable .

2.2 The Dehn function

The Dehn function was popularized by Gromov — in [56] he showed that hyperbolic groups are characterized by having linear Dehn functions. Dehn functions provide a brute-force strategy for solving the Word Problem. Gersten showed that the Word Problem is solvable for a finitely presentation group $G = \langle X \mid R \rangle$ if and only if the Dehn function is dominated by a recursive function [54]. We give here both an algebraic description of the Dehn function, and the geometric description which will be used throughout this thesis.

2.2.1 Algebraic description

Suppose that $\langle X \mid R \rangle$ is a finite presentation of the group G, with $X = \{x_1, \ldots, x_N\}$ and $R = \{r_1 \ldots r_M\}$. There is a short exact sequence of the form

$$1 \to \langle\langle R \rangle\rangle \to F(X) \to G \to 1.$$

If $w \in F(X)$ represents the identity in G, it is in the kernel of the map to G, so it can be expressed as a product of conjugates of the relators. That is, there are sequences $\{k_i\}$, with $k_i \in \{1 \dots M\}$ and $g_i \in F(X)$ such that

$$w = \prod_{i=1}^{n} r_{k_i}^{g_i}.$$

The smallest n for which such a product exists is called $\mathrm{Area}_1(w)$.

2.2.2 Geometric description

Suppose that $\langle X \mid R \rangle$ is a finite presentation for the group G and $w \in F(X)$ represents the identity in G.

Definition 2.2.1. A *van Kampen diagram* Δ for w is a simply-connected planar 2-complex with edges labeled by elements of X and directed so that the following holds: when traversing $\partial\Delta$ counterclockwise from some base vertex, we read off w, and around the boundary of each 2-cell we read an element of R. If an edge is traversed in the direction of its orientation, the positive generator is implied, and if it is traversed against its orientation, the inverse of the generator is implied.

$\mathrm{Area}_2(\Delta)$ is the number of 2-cells in Δ. $\mathrm{Area}_2(w)$ denotes the minimum area among all van Kampen diagrams with boundary word w. The next lemma asserts that the algebraic and geometric descriptions of the area of a word are equivalent. [64].

Lemma 2.2.2 (van Kampen)**.** *For all words $w \in F(X)$ representing the identity in G,* $\mathrm{Area}_1(w) = \mathrm{Area}_2(w)$.

Proof. We show that $\mathrm{Area}_1(w) \leq \mathrm{Area}_2(w)$ and $\mathrm{Area}_2(w) \leq \mathrm{Area}_1(w)$. If $\mathrm{Area}_1(w) = n$, then w is freely equal to $\prod_{i=1}^{n} r_i^{g_i}$, where $r_i \in \mathcal{R}$. We construct a bouquet of lollipops, where the ith lollipop has a 1-cell path labeled by g_i, that ends in a 2-cell corresponding to r_i. Reading along the boundary of the diagram, we get $\prod_{i=1}^{n} r_i^{g_i}$, which freely reduces to w. We can obtain a diagram with boundary word w by removing hairs and folding together successive edges which represent aa^{-1} or $a^{-1}a$, as in Figure 2.1. This reduces the number of edges in the diagram, and so eventually there will be no more possible reductions. Since each step preserves planarity, we arrive at a van Kampen diagram for w. $\text{Area}_2(w) \leq n$, since this diagram has area no more than n. Thus $\text{Area}_2(w) \leq \text{Area}_1(w)$. We induct on $\text{Area}_2(\Delta)$ to show the other inequality. The base case is clear, so suppose $\text{Area}_1(v) \leq \text{Area}_2(v)$ for all words v such that $\text{Area}_2(v) \leq n-1$. Suppose then that $\text{Area}_2(w) = n$ is witnessed by the minimal area diagram Δ, with basepoint $*$. We want to cut Δ, pulling off a lollipop as in Figure 2.2 to produce a new diagram, Θ, which has two parts: the lollipop and a diagram Δ' such that $\text{Area}_2(\Delta') < \text{Area}_2(\Delta)$. The boundary word of Θ is freely equal to the boundary word of Δ, and by the induction hypothesis, $\text{Area}_1(\Delta') \leq \text{Area}_2(\Delta')$. Therefore, $\text{Area}_1(w) \leq \text{Area}_1(\Delta') + 1 \leq \text{Area}_2(\Delta) = \text{Area}_2(w)$. $\square$

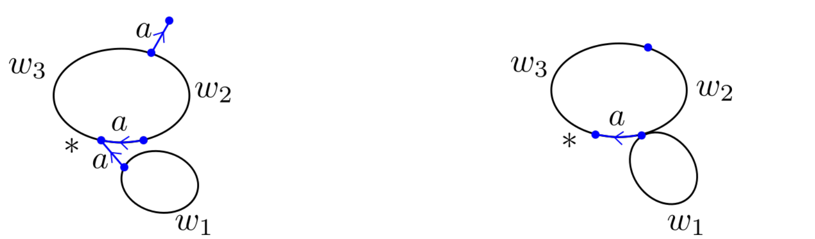

(a) A non-reduced diagram: has ‘hairs’ and admits folding

(b) Removing hairs and folding

Figure 2.1: Reducing a boundary word to w

The *Dehn function* associated to this presentation is the map $\delta : \mathbb{N} \to \mathbb{N}$ such that

$$\delta(n) := \max\{\text{Area}(w) \mid |w| \leq n \text{ and } w = 1 \text{ in } G\}$$

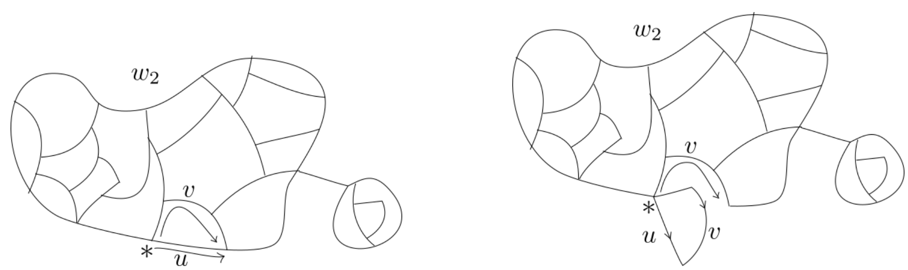

(a) Initial diagram Δ

(b) A lollipop and Δ'. $\mathrm{Area}_2(\Delta') < \mathrm{Area}_2(\Delta)$

Figure 2.2: Van Kampen's Lemma

For $f, g : \mathbb{N} \to \mathbb{N}$, write $f \preceq g$ when there exists $C > 0$ such that $f(n) \leq Cg(Cn + C) + Cn + C$ for all $n \in \mathbb{N}$, and write $f \simeq g$ when $f \preceq g$ and $g \preceq f$. This relation distinguishes for example, between quadratic and cubic Dehn functions, but it doesn't distinguish between constant and linear Dehn functions. Exponential functions are lumped together by this relation; 2^d and c^d are equivalent for all $c > 1$. Up to the equivalence relation, the Dehn function does not depend on the choice of finite presentation for G. Moreover, it is a quasi-isometry invariant among finitely presented groups.

Lemma 2.2.3 (Alonso [2]). *If $H = \langle S \mid T \rangle$ and $G = \langle X \mid R \rangle$ are quasi-isometric finitely presented groups with Dehn functions h and g, respectively, then $h \simeq g$.*

Therefore we may speak of 'the' Dehn function of a group G, by which we mean the equivalence class of the Dehn functions of G. The following is an immediate corollary of Lemma 2.2.3, which we highlight because we use it frequently.

Corollary 2.2.4. *If G is finitely presentable and $H \leq G$ is a finite index subgroup, then H is also finitely presentable and H and G have equivalent Dehn functions.*

We will ocassionally use the following third description of the Dehn function. Suppose $\langle X \mid R \rangle$ is a finite presentation for the group G. It is convenient to ask that R is closed under taking inverses and cyclic permutations. Thus for example, if $xyxy^{-1} \in R$, then so too is the inverse $yx^{-1}y^{-1}x^{-1}$ and the cyclic permutations $y^{-1}xyx, xy^{-1}xy$, and $yxy^{-1}x$. Suppose $w \in F(X)$ represents the identity in G. Then w has a *reduction sequence*: a sequence of words in $F(X)$: $w_0 = w, \ \ldots, \ w_m = 1$, where w_i is obtained from w_{i-i} by applying one of the following three moves:

(a) apply a relator: uav goes to ubv where $ab^{-1} \in \mathcal{R}$.

(b) cancel a pair of inverses: $uxx^{-1}v$ goes to uv.

(c) insert a pair of inverses: uv goes to $ux^{-1}xv$.

For a given sequence, count the number of times a move of type (a) is performed. A third area, $\mathrm{Area}_3(w)$ is the minimal number of moves of type (a) across all reduction sequences for w. Since w represents the identity, w is freely equal to $\prod_{i=1}^{n} r_{k_i}^{g_i}$. Then clearly $\mathrm{Area}_3(w) \leq \mathrm{Area}_1(w)$ — we only need to apply moves of type (a) n times to get a word of the form $\prod_{i=1}^{n} g_i^{-1} g_i$, which is freely equal to the identity. It is also the case that $\mathrm{Area}_1(w) \leq \mathrm{Area}_3(w)$. The application of a relator subsequence

$$w_i = uav \to w_{i+1} = ubv \text{ for } ab^{-1} = r_{k_i}$$

can be replaced with the sequence

$$w_i = uav \to w_{i'} = uab^{-1}bv = ur_{k_i}bv \to w_{i''} = ur_{k_i}u^{-1}ubv \to w_{i+1} = ubv.$$

This new subsequence still only has one type (a) move. Because $w_{i''} = ur_{k_i}u^{-1}uw_{i+1}$, we can gather the discarded prefixes $ur_{k_i}u^{-1}$ in order. This expresses w as a product of conjugates of the relators $\prod_{i=1}^{n} r_{k_i}^{g_i}$. Therefore $\mathrm{Area}_1(w) \leq \mathrm{Area}_3(w)$. We use such sequences in Section 5.2 to find upper bounds on the Dehn functions of mapping tori with bases $F_k \times F_l$.

Some important examples of the Dehn functions of groups include:

Example 2.2.5. The group $F_2 = \langle a, b \rangle$ has linear Dehn function. In this presentation, $\delta_{F_2}(n) = 0$ for all n, since there are no relators.

Example 2.2.6. The group $\mathbb{Z}^2$, given by the presentation $\langle a, b \mid [a, b] = 1 \rangle$ has quadratic Dehn function.

Example 2.2.7. The Heisenberg group $\mathcal{H} = \langle a, b, t \mid a^t = ab, t^b = t, a^b = b \rangle$ has cubic Dehn function.

Example 2.2.8. The Baumslag-Solitar group $BS(1, 2) = \langle a, t \mid a^t = a^2 \rangle$ has exponential Dehn function.

2.2.3 Functions equivalent to Dehn functions

In [56], Gromov showed that hyperbolic groups are characterized by having linear Dehn functions. Moreover, he showed that if $f(n) \prec n^2$, then $f(n) \simeq n$. The isomperimetric spectrum $\mathrm{IP}^1 \subset [1, \infty)$ is the set of all α such that there is a group G with Dehn function equivalent to n^α. Gromov's result implies that there is a gap in the isoperimetric spectrum: $(1, 2) \cap \mathrm{IP}^1 = \emptyset$ and also that there are no groups with Dehn function $n \log(n)$, for example. Since there are only countably many finitely presented groups, IP^1 cannot be all of $\{1\} \cup [2, \infty)$. Theorem 2 implies that $\mathbb{N} \subset \mathrm{IP}^1$. Bridson produced the first examples of groups with fractional Dehn functions in [27]. He constructed groups having subgroups with fractional distortion functions. By doubling along these subgroups, he produced groups with fractional lower bounds on their Dehn function. In [24], Bridson–Brady–Forester–Shankar proved that for every rational number $r \geq 2$ that r is the Dehn function of some finitely presented group. This team constructed 'snowflake groups' for every such $r = p/q$. Example 2.2.8 shows that not all Dehn functions have the form n^α for some α. The iterated Baumslag-Solitar group $\langle x, y, z \mid x^y = x^2, y^z = y^2 \rangle$ has Dehn function the 2-tower of exponentials 2^{2^n}. Gersten developed cohomological techniques to show that the group $\langle x, y \mid x^{x^y} = x^2 \rangle$ has Dehn function growing faster than every tower of exponentials [53]. In Part 2, we will discuss Dison and Riley's hydra groups, Γ_n, where the Dehn function of Γ_n is equivalent to the n-th Ackermann function [43]. In [85], the description of IP^1 was mostly settled — the authors find a criterion which identifies many recursive functions as Dehn functions:

Theorem (Sapir–Birget–Rips). *If $n^4 \preceq f(n)$ and $f(n)$ is a super-additive function for which the binary representation of $f(n)$ is computable in time $O(\sqrt[4]{f(n)})$ then $f(n)$ is the Dehn function of a finitely presented group.*

Now we move to methods and results that are directly used in Chapters 3–5.

2.3 Corridors and partial corridors

Corridors appear in van Kampen diagrams over a presentation $\langle X \mid R \rangle$ when there is some $a \in X$ such that all relators $r \in R$ in which a appears can be expressed as $w_1 a^{\pm 1} w_2 a^{\mp 1} w_3$ where w_1, w_2, and w_3 are words not containing $a^{\pm 1}$. Such presentations naturally arise in HNN-extensions, where the stable letter has the form of a above. Suppose Δ is a van Kampen diagram for a word w over such a presentation and that w contains an a. This edge is either in the *thin part* of that diagram, so that this edge is not in the boundary of any 2-cells, as in Figure 2.3a or it is in the *thick part* and there is a 2-cell in Δ with that edge in its boundary, as in Figure 2.3b. This 2-cell will have exactly one other a-edge. In turn, this other

a-edge either is in $\partial\Delta$ or is common with another 2-cell. This continues likewise and eventually must end elsewhere in the boundary. The resulting collection of 2-cells is an *a-corridor*. The number of 2-cells involved is the *length* of the corridor.

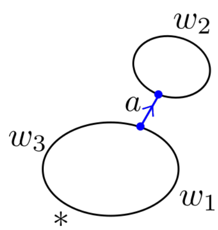

(a) a-edge in the thin part of Δ

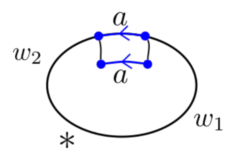

(b) a-edges in the thick part of Δ'

Corridors can also be thought of as the images of maps $\rho : [0,1] \times [0,n]$ into the van Kampen diagram, where n is the length of the a-corridor. The interval $[0,1] \times \{i\}$ maps to an edge labeled by a, and oriented such that $\rho((0,i))$ is the start of the a-edge and $\rho((1,i))$ is the end of the a-edge. The 'bottom' of an a-corridor is the path in the van Kampen diagram corresponding to the restriction of ρ to $\{0\} \times [0,n]$, the path that begins from the start of the boundary a-edge and goes along the corridor until it ends at the start of the other boundary a-edge. Suppose that G is an ascending HNN-extension. The 'top' of an a-corridor is the path in the van Kampen diagram corresponding to the restriction of ρ to $\{1\} \times [0,n]$, as in Figure 2.4. An a-corridor is *reduced* if it contains no two 2-cells sharing an a-edge for which the word around the boundary of their union is freely reducible to the identity in the group. Since G is an ascending HNN extension, this is equivalent to the word along the bottom of the a-corridor being freely reduced.

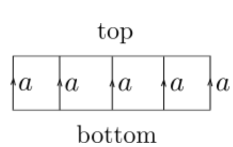

Figure 2.4: An a-corridor has a natural bottom and top, coming from the orientation of edges labeled by a.

Even in a reduced a-corridor, the word along the top of the corridor may not be freely reduced. In a *folded corridor*, 1-cells with cancelling labels (of the form bb^{-1}) are folded together so that the word along the top of the corridor is freely reducible. The words along the top of the folded and unfolded corridors are freely equal to one another. Consider the following example:

Example 2.3.1. $G = \langle a, b, t \mid a^t = ab, b^t = aba \rangle$. A t-corridor with the freely reduced word $a^{-1}b$ along the bottom will have the word $b^{-1}a^{-1}aba$ along the top. The folded corridor will have $a^{-1}b$ along the bottom and a along the top.

(a) An unfolded t-corridor

(b) A folded t-corridor

Figure 2.5: Folding a corridor

An important property of a-corridors is that they do not cross. If two a-corridors crossed, they would have a common 2-cell. This would require a relator with more than two a-edges, but all relators either have zero a-edges or two a-edges. Since a-corridors do not intersect, analysing their lengths can be a powerful tool for estimating the areas of van Kampen diagrams, as for example in [28, 29]. Since a-corridors start and end on the boundary, the number of a-corridors in a diagram for w (where w is freely reduced) is exactly the number of a in w.

Let Δ be a van Kampen diagram with N a-corridors. Since a-corridors cannot cross, removing all a-corridors from Δ leaves $N+1$ connected subdiagrams called *a-complementary regions*. The words read around the perimeters of each of these regions contain no $a^{\pm 1}$. If Δ is a van Kampen diagram with a collection of a-corridors, we can construct a *dual tree* to the set of a-corridors as follows: vertices correspond to a-complementary regions; an edge between two vertices corresponds to an a-corridor which borders the two corresponding a-complementary regions. (There is no vertex corresponding to the outside of the van Kampen diagram.)

Definition 2.3.2. A group $G = \langle X, a \mid R \rangle$, with $a \notin X$ has *a-partial corridors* when all relations in R with both a and a^{-1} have the form of corridor relations: $awa^{-1} = w'$ where $w, w' \in F(X)$. A partial corridor is a maximal set of 2-cells joined by corridor relations as above.

That is, in addition to the usual corridor relations of the form $awa^{-1} = w'$, there can also be defining relators which contain a's or a^{-1}s (but not both). We refer to the corresponding 2-cells as *capping faces*, since they cap off partial corridors. An a-edge in the boundary will either be connected by a full a-corridor to another edge labeled by a in the boundary, or it begins a partial a-corridor ending at one of the capping faces. An a-edge on a capping face is either connected to the boundary via a partial a-corridor (possibly of length zero), or is connected to an a-edge of another capping face via a partial a-corridor.

As with corridors (full corridors), partial corridors cannot cross, but unlike in the case of full corridors, there is no control a priori on the number of partial corridors, since they may begin and end within the diagram.

2.4 Gersten and Riley's electrostatic model

The electrostatic model developed by Gersten and Riley in [55] is a method of constructing van Kampen diagrams for central extensions. The electrostatic model and variants will be a primary tool in bounding Dehn functions in Chapters 3–5.

Suppose Γ is a central extension $1 \to \mathbb{Z} \to \Gamma \to \overline{\Gamma} \to 1$ with kernel $\mathbb{Z} = \langle c \rangle$. If $\overline{\Gamma}$ has presentation $\mathcal{P}_{\overline{\Gamma}} = \langle X \,|\, r_1 = 1, \ldots, r_n = 1 \rangle$, then for some $k_1, \ldots, k_n \in \mathbb{Z}$, Γ has presentation $\mathcal{P}_{\Gamma} = \langle X, c \mid r_1 = c^{k_1}, \ \ldots, \ r_n = c^{k_n}, \ [c, x] = 1, \ \forall x \in X \rangle$.

Suppose $w \in F(X \cup \{c\})$. Since c is central, $w = \overline{w}c^m$ in Γ, where $m \in \mathbb{Z}$ and $\overline{w}$ is w with all $c^{\pm 1}$ removed. If w is a word representing the identity in Γ, the word $\overline{w} \in F(X)$ represents the identity in $\overline{\Gamma}$. The electrostatic model is a way to construct a van Kampen diagram for w with respect to $\mathcal{P}_{\Gamma}$, from a minimal area diagram $\overline{\Delta}$ for $\overline{w}$ with respect to $\mathcal{P}_{\overline{\Gamma}}$. Let $f : \mathbb{N} \to \mathbb{N}$ be the Dehn function of $\overline{\Gamma}$.

Each 2-cell in $\overline{\Delta}$ has boundary corresponding to a defining relator r_i when read clockwise or counterclockwise from an appropriate vertex. We 'charge' 2-cells by inserting $|k_i|$ loops, each labeled by c, oriented in such a way that around the interior of the 2-cell we now read $r_i c^{-k_i}$ (to reflect the relation $r_i = c^{k_i}$), as in Figure 2.6. If $C := \max_i |k_i|$, charging introduces at most $Cf(|w|)$ many 1-cells labeled by c.

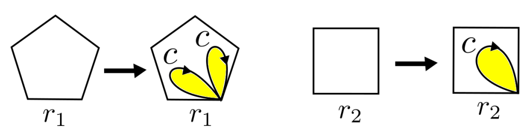

Figure 2.6: 'Charge' the diagram by replacing 2-cells from the presentation for $\overline{\Gamma}$ with 2-cells from the presentation for Γ. The pinching and shading indicate that there is not another 2-cell on the other side of the edge.

To discharge, Gersten and Riley pick a geodesic spanning tree $\mathcal{T}$ in $\overline{\Delta}^{(1)}$, rooted at the basepoint. For each introduced c-edge, e, a partial c-corridor is added which starts at e and follows $\mathcal{T}$ to the base-point of the diagram. Each c-corridor has length bounded above by $\text{Diam}(\overline{\Delta})$, the maximum distance in $\mathcal{T}$ to the base-point. This process produces a diagram Δ' for $\overline{w}c^m$ in Γ with area no more than $f(n) + Cf(n)\text{Diam}(\overline{\Delta})$.

The words w and $\overline{w}c^m$ are equal in M_Φ, so there is a van Kampen diagram Θ such that $\delta\Theta = wc^{-m}\overline{w}^{-1}$. Except for the arrangement of the c's, the words w and $\overline{w}c^m$ are freely equal, so Θ has a filling by c-corridors for which $\text{Area}(wc^{-m}\overline{w}^{-1}) \leq |m||w| \leq |w|^2$. To get a diagram Δ for w, we wrap the diagram Θ around $\overline{\Delta}$ as in Figure 2.7.

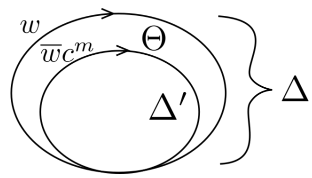

Figure 2.7: Constructing Δ from Δ' and Θ

Totaling the bounds on the number of 2-cells gives us the following theorem.

Theorem 2.4.1 (Gersten–Riley [55])**.** *Suppose Γ is a central extension*

$$1 \to \mathbb{Z} \to \Gamma \to \overline{\Gamma} \to 1$$

of a finitely presented group $\overline{\Gamma}$, and $f, g : \mathbb{N} \to \mathbb{N}$ are simultaneous upper bounds on Area and Diameter. That is, for every word $\overline{w}$ representing the identity in $\overline{\Gamma}$ of length n, there exists a van Kampen diagram $\overline{\Delta}$ such that $\text{Area}(\overline{\Delta}) \leq f(n)$ *and the diameter* $\text{Diam}(\overline{\Delta})$ *of the 1-skeleton of Δ is at most $g(n)$. Then the Dehn function of Γ is bounded above by a constant times $f(n)(g(n)+1)+n^2$.*

To use Theorem 2.4.1, we need simultaneous control of both area and diameter of diagrams. In our setting, we get this thanks to the following theorem of Papasoglu. The radius $r(\Delta)$ of a van Kampen diagram Δ is the minimal N such that for every vertex in Δ there is a path of length at most N in the 1-skeleton of Δ from that vertex to $\partial\Delta$. Since one can travel between any two vertices by concatenating shortest paths to the boundary with a path part way around the boundary,

$$\text{Diam}(\Delta) \;\leq\; 2r(\Delta) + |\partial\Delta|. \tag{2.1}$$

Theorem 2.4.2 (Papasoglu [83]). *For a group G given by a finite presentation in which every relator has length at most three, if Δ is a minimal area van Kampen diagram such that $|\partial\Delta| = n$ and $\mathrm{Area}(\Delta) \leq Mn^2$, then $r(\Delta) \leq 12Mn$.*

Every finitely presentable group has such a presentation, and changing between two finite presentations of a group alters diameter and area by at most a multiplicative constant, so, in light of (2.1), Theorem 2.4.2 gives us:

Corollary 2.4.3. *If a finitely presented group G has Dehn function bounded above by a quadratic function, then there exists $K > 0$ such that for every word of length n representing the identity, there is a van Kampen diagram whose area is at most Kn^2 and whose diameter is at most Kn.*

Corollary 2.4.4. *If Γ is a central extension of $\overline{\Gamma}$ by $\mathbb{Z}$, and $\overline{\Gamma}$ has a quadratic Dehn function, then Γ satisfies a cubic isoperimetric inequality.*

Proof. Corollary 2.4.3 implies that the conditions of Theorem 2.4.1 can be satisfied, for f a quadratic function and g a linear function. Therefore the Dehn function of Γ is bounded above by a cubic function. $\square$

In Chapter 3, we will show that all mapping tori of $F_2 \times \mathbb{Z}$ are quasi-isometric to central extensions of mapping tori of F_2. Therefore, directly applying Gersten and Riley's electrostatic model produces a cubic isoperimetric inequality for the mapping tori of $F_2 \times \mathbb{Z}$. We will refine this method in Sections 3.2 and 3.3 to improve the upper bound to quadratic in the appropriate cases. In Section 4 we generalize the method to some non-central situations.

2.5 Tools for the Dehn functions of mapping tori

We use the following tools throughout Part I. The first section gives conditions that guarantee mapping tori have equivalent Dehn functions. The next two sections give upper and lower bounds on the Dehn functions of mapping tori of RAAGs.

2.5.1 Equivalent Dehn functions

The next lemma, from [16], allow us to specialize to convenient Φ for the purpose of calculating the Dehn functions of mapping tori.

As previously, $G = \langle X \mid R \rangle$ is a finitely presented group, Φ is an automorphism of G, and $M_\Phi := \langle X, t \mid \mathcal{R},\ t^{-1}xt = \Phi(x),\ \forall x \in X \rangle$.

Lemma 2.5.1 (Bogopolski [16]). *The following mapping tori have equivalent Dehn functions:*

1. *M_Φ and M_{Φ^n}, for any $n \in \mathbb{N}$.*
2. *M_Φ and $M_{\Phi^{-1}}$.*
3. *M_{Φ_1} and M_{Φ_2} when $[\Phi_1]$ and $[\Phi_2]$ are conjugate in $\mathrm{Out}(G)$.*

Proof. To see (1), notice that $M_{\Phi^n} = \langle X, \tau \mid R, \tau^{-1}x\tau = \Phi^n(x) \rangle$ is isomorphic to the subgroup $H = \langle X, t^n \mid R, t^{-n}xt^n = \Phi^n(x) \rangle$. We show that H is finite index in G, so that by Proposition 2.2.4, M_Φ and M_{Φ^n} have equivalent Dehn functions. Consider the short-exact sequence

$$\{1\} \to G \hookrightarrow M_\Phi \xrightarrow{p} \mathbb{Z} \to \{1\}.$$

By composing p with $q : \mathbb{Z} \to \mathbb{Z}/n\mathbb{Z}$, we see that the kernel of $q \circ p : M_\Phi \to \mathbb{Z}/n\mathbb{Z}$ is H. In particular, H has index n in M_Φ.

To see (2), for every $x \in X$, we can rewrite the relation $g^t = \Phi(g)$ as $h^{t^{-1}} = \Phi^{-1}(h)$, where $h = \Phi(g)$. Mapping $t \mapsto t^{-1}$ and G to G by the identity gives an isomorphism $M_\Phi \to M_{\Phi^{-1}}$, so M_Φ and $M_{\Phi^{-1}}$ have equivalent Dehn functions.

The final result allows us to freely change automorphisms when calculating Dehn functions, so long as we stay in the same conjugacy class of $\mathrm{Out}(G)$. The proof of 3 follows [16]: If Φ_1 and Φ_2 are conjugate in $\mathrm{Out}(G)$, there exists $\psi \in \mathrm{Aut}(G)$ and $h \in G$ such that $\Phi_2(g) = \psi^{-1}(\Phi_1(\psi(g^h)))$ for all $g \in G$. (Indeed, conjugacy in $\mathrm{Out}(G)$ implies that there are $[\phi] = [\psi] \in \mathrm{Out}(G)$ such that $\Phi_2 = \psi^{-1}(\Phi_1(\phi))$. Since $\psi = \phi \circ \iota^h$ for some $h \in G$, we get the claim.) We will show that M_{Φ_1} and M_{Φ_2} are isomorphic. Consider $\Psi : M_{\Phi_2} \to M_{\Phi_1}$ given by $x \mapsto \psi(x)$ for $x \in G$ and $t \mapsto t\hat{h}$, where $\hat{h} := \Phi_1(\psi(h))$. It is a homomorphism because the relators $(g^{-1})^t\Phi_2(g)$ for $g \in G$ are mapped to the identity in M_{Φ_1}. Indeed,

$$\begin{aligned}\Psi\left((g^{-1})^t\Phi_2(g)\right) &= \Psi\left((g^{-1})^t\psi^{-1}(\Phi_1(\psi(g^h)))\right) = \psi(g^{-1})^{t\hat{h}}\Phi_1(\psi(g^h)) \\ &= (\psi(g)^{-1})^{t\hat{h}}\Phi_1(\psi(g))^{\hat{h}} = \left((\psi(g)^{-1})^t\Phi_1(\psi(g))\right)^{\hat{h}} = 1^{\hat{h}} = 1.\end{aligned}$$

It is certainly onto. This homomorphism has inverse given by $x \mapsto \psi^{-1}(x)$ for $x \in G$ and $t \mapsto t\ \psi^{-1}(\hat{h}^{-1})$, so it is an isomorphism. $\square$

Lemma 2.5.1 has an immediate corollary for the groups $A(\Gamma)$ where $\mathrm{Out}(A(\Gamma))$ is finite, such as $A(C_n)$ for $n > 4$ (see Example 2.1.9). For any $\Phi \in \mathrm{Aut}(A(\Gamma))$, for n large enough, $[\Phi]^n = [\Phi^n] = [\mathrm{Id}]$. Therefore M_Φ has equivalent Dehn function to $M_{\mathrm{Id}} = A(\Gamma) \times \mathbb{Z}$, so M_Φ has quadratic Dehn function.

2.5.2 General bounds on Dehn functions of mapping tori of RAAGs

Lemma 2.5.2. *If G is a non-free RAAG and $\Phi \in \mathrm{Aut}(G)$, the Dehn function of M_Φ satisfies $n^2 \preceq f(n) \preceq 2^n$.*

Proof. Let $A(\Gamma)$ be a non-free RAAG. Then it is built from a graph with at least one edge and so M_Φ will contain a $\mathbb{Z}^2$ subgroup ([5]). This implies that M_Φ is not hyperbolic and therefore $n^2 \preceq f(n)$ (see [26] and the references therein).

Suppose $A(\Gamma) = \langle X \mid R \rangle$ and $\Phi \in \mathrm{Aut}(A(\Gamma))$. If $w \in F(X \cup \{t\})$ represents the identity in M_Φ, then $w = t^{k_0} a_1 t^{k_1} \cdots a_m t^{k_m}$ for some $a_1, \ldots, a_m \in X^{\pm 1}$, and $k_1, \ldots, k_m \in \mathbb{Z}$. Let $C = \max_i \{|\Phi^{\pm}(a_i)|\}$. To get the upperbound, the strategy is to apply relators to shuffle all $t^{\pm 1}$ to the right. Moving $t^{\pm 1}$ to the right past a_i will replace a_i with $\Phi^{\pm 1}(a_i)$, so the length of segments in $F(X)$ may grow. At the end of this shuffling stage we get $u \in F(X)$ that represents the identity in $A(\Gamma)$. A generous overestimate for the number of relator applications it takes to get to u is to assume that every t has to be shuffled through a word of length $|w|C^{|w|}$, and that there are $|w|$ such t's that need to be shuffled. We get $|w|C^{|w|}$ by estimating that under $|w|$ applications of $\Phi^{\pm 1}$, the image of the generator has length no more than $C^{|w|}$, and assuming that all elements in w grow this big. This produces a $|w|^2 C^{|w|}$ upper bound on the number of relators necessary for this stage. All RAAGs are automatic, so the Dehn function of $A(\Gamma)$ is at most quadratic, and there is $c > 0$ such that $\mathrm{Area}(u) \leq c|u|^2 \leq c(|w|C^{|w|})^2$. Thus $\mathrm{Area}(w)$ is at most $|w|^2\, C^{|w|} + (|w|\, C^{|w|})^2$, and therefore $f(n) \preceq 2^n$. $\square$

2.5.3 Abelian subgroups in the base group

The following lemma is the special case of Theorem 4.1 of Bridson and Gersten [28] in which (in their notation) $G = H$ and K is quasi-isometrically embedded:

Lemma 2.5.3 (Bridson–Gersten)**.** *Suppose that $K = \langle k_1, \ldots, k_m \rangle \cong \mathbb{Z}^m$ is a quasi-isometrically embedded infinite abelian subgroup of a finitely generated base group G. If $\Phi \in \operatorname{Aut}(G)$ and Φ restricts to an automorphism of K. Then the Dehn function $\delta_{M_\Phi}(n)$ of $M_\Phi = G \rtimes_\Phi \mathbb{Z}$ satisfies*

$$n^2 \max_{1 \le i \le m} \left| \Phi^{\pm n}(k_i) \right| \;\preceq\; \delta_{M_\Phi}(n).$$

Equivalently, if ϕ is the restriction of Φ to K, exactly one of the following holds:

1. *Not all eigenvalues of ϕ are unit. Then $2^n \preceq \delta_{M_\Phi}(n)$.*
2. *All eigenvalues of ϕ are unit. If the largest Jordan block associated to ϕ is size $c \times c$, then $n^{c+1} \preceq \delta_{M_\Phi}(n)$.*

This will be an important tool in upcoming chapters — it provides lower bounds for the Dehn functions of many of the mapping tori we consider here. The basic idea of the proof was described in Section 1.2.1, and a proof that is closely modeled on Bridson and Gersten's proof appears in Section 5.2.

2.6 Growth and automorphims of $\mathbb{Z}^2$

For mapping tori of $F_2 \times \mathbb{Z}$ and $\mathbb{Z}^2 * \mathbb{Z}$, we will work with automorphisms of $\mathbb{Z}^2$. In this section we present the basics of conjugacy of integer-valued 2×2 matrices, with an eye towards applying Lemma 2.5.1 so that we only need to calculate Dehn functions for mapping tori built from automorphisms with particularly nice forms.

Choose a basis X for $\mathbb{Z}^2$. The *growth* of $\phi \in \mathrm{SL}(2,\mathbb{Z})$ is the function $\mathbb{N} \to \mathbb{N}$ given by $f(n) = \max_{x\in X}\{|\phi^{\pm 1}(x)|\}$. Conjugate matrices A and $P^{-1}AP$ have growth functions that are bi-Lipschitz equivalent. Putting this equivalence relation on the growth functions, we see that there are three types of growth functions: trivial growth, where this function is bounded, linear growth, where it is equivalent to a linear function, $n \mapsto Cn$ and exponential growth, where the function is equivalent to $n \mapsto \lambda^n$ for some $\lambda > 1$. The trivial growth elements of $\mathrm{SL}(2,\mathbb{Z})$ have finite order. There are only finitely many conjugacy classes of finite order elements. Indeed, since $A \in \mathrm{SL}(2,\mathbb{Z})$, the eigenvalues λ satisfy that $\lambda^2 - \mathrm{tr}(A)\lambda + 1 = 0$. Either $\lambda = \pm 1$, or λ is not real, and so the discriminant $\mathrm{tr}(A)^2 - 4 < 0$. Since $\mathrm{tr}(A)$ is integer-valued, $\mathrm{tr}(A) \in \{0, \pm 1\}$. From the trace of the Jordan Canonical Form (JCF), if $a + bi$ and $a - bi$ are the eigenvalues of A, $\mathrm{tr}(A) = 2a$, and thus $a \in \{0, \pm\frac{1}{2}\}$. Since $a^2 + b^2 = 1$, the matrix can only correspond to rotations by $\pm\pi/2$, $\pm 2\pi/3$, or $\pm\pi/3$, which are all finite order. Since there are only finitely many JCFs with finite order, there are only finitely many conjugacy classes of finite order elements. The linear growth elements are those with JCF of $\left(\begin{smallmatrix} \lambda & \mu \\ 0 & \lambda \end{smallmatrix}\right)$ where $\mu \neq 0$ and $|\lambda| = 1$. Exponential growth elements are those with JCF of $\left(\begin{smallmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{smallmatrix}\right)$ where λ is real and $|\lambda| \neq 1$.

The following lemma will allow us to specialize to convenient cases of Φ when analyzing Dehn functions of mapping tori M_Φ of $F_2 \times \mathbb{Z}$ and $\mathbb{Z}^2 * \mathbb{Z}$.

Lemma 2.6.1. *If $A \in \mathrm{SL}(2,\mathbb{Z})$ has linear growth and only the eigenvalue 1, then it is conjugate in $\mathrm{SL}(2,\mathbb{Z})$ to $\left(\begin{smallmatrix} 1 & \alpha \\ 0 & 1 \end{smallmatrix}\right)$.*

Proof. By the JCF, matrices which have linear growth and only the eigenvalue

1 are conjugate in $\mathrm{SL}(2,\mathbb{C})$ to $\left(\begin{smallmatrix}1 & \alpha\\ 0 & 1\end{smallmatrix}\right)$. We show that this conjugacy still holds in $\mathrm{SL}(2,\mathbb{Z})$ when A has integral entries. Let $N = A - I$. Either from the JCF for A, or the fact that by definition, $Nv_1 = 0$ for some eigenvector $v_1 \in \mathbb{C}^2$ and $Nv_2 = v_1$ for some $v_2 \in \mathbb{C}^2$, we can see that $N \neq 0$, $N^2 = 0$ and that the trace of N is zero. Therefore either A is already in JCF, or $N = \left(\begin{smallmatrix}a & \frac{-a^2}{b}\\ b & -a\end{smallmatrix}\right)$ for some $b \neq 0$. In the second case, we require a matrix in $\mathrm{SL}(2,\mathbb{Z})$ that conjugates A to its JCF. If $a \neq 0$, choose $p, q \in \mathbb{Z}$ such that $(p,q) = 1$ such that $p/q = b/a$. If $a = 0$, let $p = 0$ and $q = 1$. By Bezout's identity, integers $\alpha, \beta \in \mathbb{Z}$ exist such that $p\beta - q\alpha = 1$.

With respect to the basis $(p,q)^{\mathsf{T}}$, $(\alpha,\beta)^{\mathsf{T}}$, our linear transformation will have matrix representative of the form $\left(\begin{smallmatrix}1 & *\\ 0 & *\end{smallmatrix}\right)$. Since this matrix has determinant 1, has the form $\left(\begin{smallmatrix}1 & *\\ 0 & 1\end{smallmatrix}\right)$. $\square$

CHAPTER 3

THE BASE GROUP IS $G = F_2 \times \mathbb{Z}$

3.1 Automorphisms of $F_2 \times \mathbb{Z}$

Recall the notation of Theorem 4:

> $\Phi \in \mathrm{Aut}(F_2 \times \mathbb{Z})$ and induces $\phi \in \mathrm{Aut}(F_2)$ via the map $F_2 \times \mathbb{Z} \twoheadrightarrow F_2$ that kills the $\mathbb{Z}$ factor;
>
> $\phi_{ab} \in \mathrm{Aut}(\mathbb{Z}^2)$ is induced by ϕ via the abelianization map $F_2 \to \mathbb{Z}^2$; and
>
> $p : F_2 \times \mathbb{Z} \to \mathbb{Z}$ is projection onto the second factor.

We prove Theorem 4 by addressing the following cases:

(i) ϕ_{ab} has a non-unit eigenvalue.

(ii) ϕ_{ab} has only unit eigenvalues, and

  (a) $\exists\, g \in F_2$ and $m \in \mathbb{N}$ such that $p(\Phi^m(g)) \neq 0$ and $\phi_{ab}^m(g_{ab}) = g_{ab}$,

  (b) for all $g \in F_2$ and all $m \in \mathbb{N}$ such that $\phi_{ab}^m(g_{ab}) = g_{ab}$, we have $p(\Phi^m(g)) = 0$.

In cases (ii)a and (ii)b, either ϕ_{ab} has a real unit eigenvalue, or it corresponds to a finite order rotation, as discussed in Section 2.6. In either case, some power of ϕ_{ab} has a fixed point. In case (ii)a, there is a $g \in F_2$ that abelianizes to a fixed point of some power m of ϕ_{ab}, but $p(\Phi^m(g)) \neq p(g)$. In case (ii)b, there is no such g.

The three cases listed above are comprehensive and mutually exclusive. Let M_Φ be the mapping torus of $F_2 \times \mathbb{Z}$ associated to Φ. In Section 3.2 we prove that the Dehn function of M_Φ is quadratic in case (i); in Section 3.3 we prove it is cubic in case (ii)a and it is quadratic in case (ii)b.

The following theorem of Nielsen is Proposition 5.1 in [72]:

Lemma 3.1.1 (Nielsen)**.** *For all $\theta \in \mathrm{Aut}(F(a,b))$, θ^2 sends $[a,b]$ to a conjugate. Thus there is $\psi \in \mathrm{Aut}(F(a,b))$ such that $[\theta^2] = [\psi]$ in $\mathrm{Out}(F(a,b))$ and $\psi([a,b]) = [a.b]$.*

Proof. We show that for every $\theta \in \mathrm{Aut}(F(a,b))$ that θ maps $[a,b]$ to a conjugate of $[a,b]$ or $[a,b]^{-1}$. Nielsen transformations generate the automorphism group of a free group. In the case of $F(a,b)$ there are five Nielsen transformations: $a \mapsto a^{-1}$, $b \mapsto b$; $a \mapsto a$, $b \mapsto b^{-1}$; $a \mapsto b$, $b \mapsto a$; $a \mapsto ab$, $b \mapsto b$; and $a \mapsto a$, $b \mapsto ba$. Each maps $[a,b]$ to a conjugate of $[a,b]$ or $[a,b]^{-1}$, so all compositions also send $[a,b]$ to a conjugate of $[a,b]$ or $[a,b]^{-1}$. If $\theta^2([a,b]) = [a,b]^g$, then $\psi := \iota_{g^{-1}} \circ \theta^2$. □

Lemma 3.1.2. *If Φ is an automorphism of $G = F_2 \times \mathbb{Z} = \langle a,b \rangle \times \langle c \rangle$, then*

$$M_{\Phi^2} \;=\; \langle a,b,c,t \mid a^t = \phi(a)c^{k_a},\ b^t = \phi(b)c^{k_b},\ c^t = c,\ [a,c]=1,\ [b,c]=1 \rangle$$

for some $k_a, k_b \in \mathbb{Z}$. In particular, M_{Φ^2} is a central extension of the free-by-cyclic group

$$\langle a,b,t \mid a^t = \phi(a),\ b^t = \phi(b) \rangle.$$

Moreover, if $\Phi^2([a,b]) = [a,b]^g$ then $[\Phi^2] = [\Psi] \in \mathrm{Out}(G)$ where $\Psi = \iota_{g^{-1}} \circ \Phi^2$. Then

$$M_\Psi \;=\; \langle a,b,c,t \mid a^t = \psi(a)c^{k_a},\ b^t = \psi(b)c^{k_b},\ c^t = c, \rangle$$

where the k_a and k_b are the same, and $\psi \in \mathrm{Aut}(F(a,b))$ fixes $[a,b]$.

Proof. The center $\langle c \rangle$ of G, being characteristic, is preserved by Φ. So Φ maps c to c or c^{-1}, and Φ^2 maps c to c. On killing c, Φ induces some $\theta \in \mathrm{Aut}(F(a,b))$.

Therefore, the mapping torus $M_{\Phi^2} = G \rtimes_{\Phi^2} \mathbb{Z}$ is presented by

$$\langle a, b, c, t \mid a^t = \theta^2(a)c^{k_a},\ b^t = \theta^2(b)c^{k_b},\ c^t = c,\ [a, c] = 1,\ [b, c] = 1\rangle$$

for some $k_a, k_b \in \mathbb{Z}$. By Lemma 3.1.1, $\theta^2([a, b]) = [a, b]^g$ for some $g \in F(a, b)$. Defining $\Psi = \iota_{g^{-1}} \circ \Phi^2$ as above, since c is central, the indices k_a and k_b will not change, but the image of $[a, b]$ will be $[a, b]$, as desired. □

By Lemma 2.5.1, the Dehn functions of M_Φ and M_{Φ^2} and M_Ψ are equivalent, so when analyzing their Dehn functions, it will suffice to study mapping tori of $F_2 \times \mathbb{Z}$ that are presented as in Lemma 3.1.2.

Moreover, conditions (i), (ii)a, and (ii)b hold for the map ϕ_{ab}^2 exactly if they hold for ϕ_{ab}. The map ϕ_{ab} has a non-unit eigenvalue if and only if ϕ_{ab}^2 has a non-unit eigenvalue, and ϕ_{ab} has only unit eigenvalues if and only if ϕ_{ab}^2 has only unit eigenvalues. If there exists $g \in F_2$ and $m \in \mathbb{N}$ with $\phi_{ab}^m(g_{ab}) = g_{ab}$, it is also the case that $(\phi_{ab}^2)^m(g_{ab}) = g_{ab}$. Moreover, $p((\Phi^2)^m(g)) = 2p(\Phi^m(g))$, so $p((\Phi^2)^m(g))$ is non-zero or zero if and only if $p(\Phi^m(g))$ is non-zero or zero, respectively. Thus, we can use this classification for Φ^2 instead of Φ. A similar argument implies that these properties hold for Ψ if and only if they hold for Φ^2, so we can use this classification for Ψ instead of Φ if convenient.

3.2 The quadratic Dehn function when ϕ_{ab} has non-unit eigenvalues

The primary tool for this section is *relative hyperbolicity*. This concept was first studied by Gromov, and later developed by Farb, Bowditch, Osin, and others [49, 23, 82].

3.2.1 Relative hyperbolicity

Let H be a subgroup of the finitely generated group G, generated by X. Then G is the quotient of the free product with H:

$$1 \to N \to H * F(X) \to G \to 1.$$

Definition 3.2.1. A presentation for G *relative to the subgroup* H is

$$\mathcal{P} = \langle X, H \mid \mathcal{S}, \mathcal{R} \rangle,$$

where $\mathcal{S}$ is the set of all words in H that represent the identity in H, and $\langle\langle \mathcal{R} \rangle\rangle = N$. If X and $\mathcal{R}$ can be chosen to be finite, this is called a *finite relative presentation.*

Definition 3.2.2 (Osin)**.** Let w be a word in X and the alphabet $H - \{1\}$, ie $w \in H * F(X)$. Then $\mathrm{Area}^{rel}(w)$ is the smallest m such that $w =_{H*F(X)} \prod_{i=1}^{m} r_i^{a_i}$, where $r_i \in \mathcal{R}$, and $a_i \in H * F(X)$.

If it is well-defined, the *relative Dehn function* of G is the maximum of $\mathrm{Area}^{rel}(w)$ over words $|w| \leq n$ in $H * F(X)$. A group G is *strongly hyperbolic relative to the subgroup* H if $\delta^{rel}(n)$ is equivalent to a linear function. As long as G is countable, this is equivalent to other standard forms of strong relative hyperbolicity such as Farb's relative hyperbolicity with Bounded Coset Penetration and Bowditch's relative hyperbolicity with fineness [82].

Van Kampen's Lemma still holds, so if w is a word in X and $H - \{1\}$, we can construct a van Kampen diagram Δ for w over this presentation. The 2-cells can be divided into $\mathcal{R}$-cells, which represent relators in $\mathcal{R}$, and $\mathcal{S}$-cells, which represent relators in $\mathcal{S}$. Δ is of *minimal type* if under lexicographic ordering, it minimizes

$$(N_{\mathcal{R}} = \# \text{ of } \mathcal{R}\text{-faces}, \quad N_{\mathcal{S}} = \# \text{ of } \mathcal{S} \text{ faces}, \quad E = \text{total } \# \text{ of edges}),$$

Call an edge *internal to the $\mathcal{S}$-faces* when it has $\mathcal{S}$-faces (or a single $\mathcal{S}$-face) on both sides. Call an edge *external*, or in the *thin part* of the diagram, if it is not a face of any 2-cell.

Lemma 3.2.3 (Osin, Lemma 2.15 in [82])**.** *If we have a finite relative presentation and Δ is a van Kampen diagram of minimal type for w, then Δ has no edges which are internal to $\mathcal{S}$-faces. Moreover, letting $M = \max_{R\in\mathcal{R}} |R|$, the total number of edges in $\mathcal{S}$-faces in Δ is at most $|w| + MN_{\mathcal{R}}(\Delta)$.*

A consequence of Lemma 3.2.3 is that if Δ is a diagram of minimal type, $\mathcal{S}$-faces form disjoint *islands* in Δ and there are no $\mathcal{R}$-faces enclosed within these islands.

Lemma 3.2.4 (Osin [82])**.** *If $G = \langle X \mid R\rangle$ is strongly hyperbolic relative to a subgroup H, there exists $C > 0$ such that every word w in X of length n representing the identity has a van Kampen diagram Δ with respect to $\langle X, H \mid \mathcal{R}, \mathcal{S}\rangle$ with the following properties:*

1. *The number of $\mathcal{R}$-faces is at most Cn.*
2. *There is a spanning tree $\mathcal{T}$ in the 1-skeleton of the union of the $\mathcal{R}$-faces and external edges, which satisfies $\mathrm{Diam}(\mathcal{T}) \leq (MC + 1)n$.*

The first claim comes from the definition of relative hyperbolicity. By Lemma 3.2.3, every edge in Δ either belongs to the 1-skeleton of the union of the $\mathcal{R}$-faces (there are no more than MCn such edges) or does not have a 2-cell on either side (there are no more than n such edges). The number of edges in the 1-skeleton of Δ, $MCn + n$, is clearly an upperbound on the diameter of any spanning tree.

3.2.2 A quadratic upperbound on the Dehn function

Suppose that M_ϕ is presented by $\mathcal{P}_1 := \langle a, b, t \mid a^t = \phi(a),\ b^t = \phi(b) \rangle$ where $\phi \in \mathrm{Aut}(F(a, b))$ such that $\phi([a, b]) = [a, b]$. (See Lemma 3.1.2.)

Lemma 3.2.5. *If ϕ_{ab} has non-unit eigenvalues, then M_ϕ is strongly hyperbolic relative to the subgroup $H = \langle [a, b],\ t \rangle$, where $H \cong \mathbb{Z}^2$.*

Proof. M_ϕ is the fundamental group of a finite-volume hyperbolic once-punctured torus bundle. In Theorem 4.11 of [49], Farb showed that such groups are strongly hyperbolic relative to their cusp subgroups. In our case, that is the subgroup $\langle [a, b],\ t \rangle$. (See also Section 4 of [33] for a survey of when mapping tori of free groups are relatively hyperbolic and acylindrically hyperbolic). □

It is convenient to use the following presentation for M_ϕ:

$$\mathcal{P}_2 \ := \ \langle a, b, z, t \mid a^t = \phi(a),\ b^t = \phi(b),\ z = [a, b],\ z^t = z \rangle,$$

which is obtained from $\mathcal{P}_1$ by adding an extra generator z, an extra relation which declares that z equals $[a, b]$ in the group, and a further extra relation which declares that z commutes with t, which was arranged in Lemma 3.1.2. Thus $\langle t, z \rangle \cong \mathbb{Z}^2$ is the subgroup of Lemma 3.2.5. Refer to 2-cells of a van Kampen diagram over $\mathcal{P}_2$ as $\mathbb{Z}^2$-*faces* when they correspond to the relation $z^t = z$, and refer to the remaining 2-cells as $\mathcal{R}$-*faces.*

Lemma 3.2.6. *There exists $C > 0$ such that every word w on $\{a, b, z, t\}^{\pm 1}$ of length n representing the identity has a van Kampen diagram Δ with respect to $\mathcal{P}_2$ with the following properties.*

1. *The number of $\mathcal{R}$-faces is at most Cn.*

2. *From every vertex of Δ on the perimeter of a $\mathcal{S}$-face, there is a path to $\partial\Delta$ of length at most Cn in the 1-skeleton of the union of the $\mathcal{R}$-faces and the thin part of the van Kampen diagram.*

3. *The number of $\mathbb{Z}^2$-faces in Δ is at most Cn^2.*

Proof. Let $H = \{h_{ij} \mid i, j \in \mathbb{Z}\}$ be the subgroup of M_ϕ generated by t and z in M_ϕ, where h_{ij} is the element represented by $t^i z^j$ and $h_{10} := t$ and $h_{01} := z$. Let $\mathcal{S}$ denote the set of words in the alphabet $H - \{1\}$ that represent the identity in H. Consider the presentation

$$\mathcal{P}_3 \ := \ \langle a, b, H \mid a^t = \phi(a),\ b^t = \phi(b),\ z = [a, b],\ \mathcal{S}\rangle.$$

The elements t and z appear in H and removing the defining relation $z^t = z$ is just a technical convenience: it is an element of $\mathcal{S}$. Since
$\mathcal{R} = \{a^t\phi(a)^{-1},\ b^t\phi(b)^{-1},\ z[a, b]^{-1}\}$ and $X = \{a, b\}$ are finite, this is a finite relative presentation for M_ϕ with respect to H. Moreover, van Kampen diagrams over $\mathcal{P}_2$ and $\mathcal{P}_3$ have the same set of 2-cells that are called $\mathcal{R}$-faces. For every word w in $\{a, b, z, t\}^{\pm 1}$, we find a van Kampen diagram of minimal type with respect to $\mathcal{P}_3$. Since M_ϕ is hyperbolic relative to the subgroup H, Lemma 3.2.3 implies that the number of *frontier edges*, which have an $\mathcal{R}$-face on one side and an $\mathcal{S}$-face on the other side, is bounded above by Cn for some $C > 0$. Our presentation implies that the labels of the frontier edges will be elements of $\{t, z\}^{\pm 1}$. Around each H-island is a word h in $\{t, z\}^{\pm 1}$ which represents the identity in H. We can 'blow-up' the island by replacing it with a minimal area diagram for h over $\mathcal{P}_2$. In particular, it will be filled with $\mathbb{Z}^2$-faces. After blowing up all H-islands in this way, we get a van Kampen diagram Δ for w over the presentation $\mathcal{P}_2$. The properties 1 and 2 of Lemma 3.2.3 translate directly into 1

and 2. Moreover, since the total number of frontier edges is linear in n, the number of $\mathbb{Z}^2$-faces is quadratic in n, giving 3. $\square$

Proof of Theorem 4 in Case (i). In the presentation

$$\langle a, b, c, t \mid a^t = \phi(a)c^{k_a},\ b^t = \phi(b)c^{k_b},\ c^t = c,\ [a, c] = 1,\ [b, c] = 1\rangle,$$

letting $z = [a, b]$, we see that $z^t = z$, since c is central:

$$z^t = [a, b]^t = [\phi(a)c^{k_a}, \phi(b)c^{k_b}] = [\phi(a), \phi(b)] = [a, b] = z.$$

Thus $\mathcal{Q}$, a central extension of $\mathcal{P}_2$, is a presentation for M_Φ:

$$\begin{aligned} \mathcal{Q} \ := \ \langle a, b, c, t, z \mid a^t = \phi(a)c^{k_a},\ b^t = \phi(b)c^{k_b},\ c^t = c, \\ [a, c] = 1,\ [b, c] = 1,\ z = [a, b],\ z^t = z\rangle. \end{aligned}$$

Suppose w is a word of length n in $\{a, b, c, t\}^{\pm 1}$ representing the identity in $\mathcal{Q}$. Let $\overline{w}$ be w with all $c^{\pm 1}$ deleted. Then $\overline{w}$ represents the identity in $\mathcal{P}_2$.

Let $\overline{\Delta}$ be a van Kampen diagram for $\overline{w}$ as per Lemma 3.2.6. Given (2) of that lemma, there is a tree $\mathcal{T}$ in the 1-skeleton of the union of the $\mathcal{R}$-faces and external edges in $\overline{\Delta}$, such that $\mathcal{T}$ is rooted in the boundary and has diameter no more than Cn, for some $C > 0$ that does not depend on $\overline{w}$.

Charge $\overline{\Delta}$. Given that the defining relation $z^t = z$ is unchanged on lifting to the central extension, the Cn^2 $\mathbb{Z}^2$-faces of Lemma 3.2.6(3), are unchanged. Let $m = \max\{|k_a|, |k_b|\}$. The remaining Cn $\mathcal{R}$-faces of Lemma 3.2.6(1), each acquire at most m charges. These are discharged by adding partial c-corridors that follow $\mathcal{T}$ to the boundary. These partial c-corridors have length at most Cn, by Lemma 3.2.6 (2), and so contribute at most C^2mn^2 many 2-cells to the diagram. The result is a diagram over $\mathcal{Q}$ of area at most $(C^2m + C)n^2$ for a word $\overline{w}c^k$, with length at most $|w|$. By adding in an annular region to rearrange $\overline{w}c^k$ to w, as in Section 2.4, it follows that w has a diagram over $\mathcal{Q}$ of area at most $(C^2m + C + 1)n^2$. $\square$

3.3 The Dehn function when all eigenvalues of ϕ_{ab} are unit

We begin by arguing that it suffices for the purpose of determining Dehn functions to consider certain special automorphisms. By Lemma 3.1.2 it suffices to assume that $\Phi(a) = \phi(a)c^{k_a}, \Phi(b) = \phi(b)c^{k_b}, \Phi(c) = c$, for some $\phi \in \mathrm{Aut}(F_2)$.

Since two elements $f,\ g \in \mathrm{Aut}(F_2)$ in the same outer-automorphism class induce isomorphic mapping tori M_f and M_g, the classical result of Nielsen that $\mathrm{Out}(F_2) \cong \mathrm{GL}(2,\mathbb{Z})$ implies that we may choose convenient representatives of $[\phi_{ab}] \in \mathrm{GL}(2,\mathbb{Z})$. As per Lemma 2.5.1, we assume that $\phi_{ab} \in \mathrm{SL}(2,\mathbb{Z})$, else we can replace Φ by Φ^2. If ϕ_{ab} has only unit eigenvalues, then some power ϕ_{ab}^m will have a fixed point. Therefore, since M_Φ and M_{Φ^m} have equivalent Dehn functions, we may assume that ϕ_{ab} has a non-trivial fixed point, w_{ab}.

Lemma 3.3.1. *If ϕ_{ab} has non-trivial fixed point w_{ab}, there is $\Psi \in \mathrm{Aut}(F_2 \times \mathbb{Z})$ such that $[\Phi]$ and $[\Psi]$ are conjugate in $\mathrm{Out}(F_2 \times \mathbb{Z})$ and $\psi(b) = b$, where $\psi = \Psi \restriction_{F(a,b)}$. Moreover,*

1. *If $\phi_{ab} \neq \mathrm{id}$, then $p(\Psi(b)) \neq 0$ if and only if $p(\Phi(w)) \neq 0$.*
2. *If $\phi_{ab} = \mathrm{id}$, then for every $x \in F_2$, x_{ab} is a fixed point of ϕ_{ab}, so the single choice of w may not give us the full picture.*
    - (a) *if $p(\Phi(x)) = 0$ for all $x \in \langle a, b\rangle$, then $\Phi \in \mathrm{Inn}(F_2 \times \mathbb{Z})$ and $\Psi = \mathrm{Id}$.*
    - (b) *if $p(\Phi(x)) \neq 0$ for some x, then $p(\Psi(b)) \neq 0$.*

Lemmas 2.5.1 and 3.3.1 immediately imply that for the purpose of calculating Dehn functions, we can restrict our attention to a much smaller family of automorphisms, Ψ. The automorphisms Ψ and Φ fall in the same slots in the classification of Theorem 4.

Corollary 3.3.2. *The mapping tori M_Φ and M_Ψ above have equivalent Dehn functions. The automorphism Ψ has the form*

$$\Psi : \quad a \mapsto ab^\beta c^{k_a}, \quad b \mapsto bc^{k_b} \quad c \mapsto c,$$

where $k_b \neq 0$ if and only if Φ satisfied Case (ii)a and $k_b = 0$ if and only if Φ satisfied Case (ii)b.

Proof of Lemma 3.3.1. By Lemma 2.6.1, ϕ_{ab} is conjugate in $\mathrm{SL}(2,\mathbb{Z})$ to $\left(\begin{smallmatrix} 1 & \beta \\ 0 & 1 \end{smallmatrix}\right)$, for some $\beta \in \mathbb{Z}$. From the isomorphism between $\mathrm{Out}(F_2)$ and $\mathrm{GL}(2,\mathbb{Z})$, we get that $[\phi]$ is conjugate in $\mathrm{Out}(F_2)$ to the automorphism $[\psi]$ such that $\psi(a) = ab^\beta$ and $\psi(b) = b$. In particular, $\psi = f^{-1} \circ \phi \circ f \circ \iota_g$ for some $f \in \mathrm{Aut}(F_2)$ and $\iota_g \in \mathrm{Inn}(F_2)$. We extend f, ψ, and ι_g from $\mathrm{Aut}(F_2)$ up to $\mathrm{Aut}(F_2 \times \mathbb{Z})$ by defining $F \in \mathrm{Aut}(F_2 \times \mathbb{Z})$ by $F(gc^k) = f(g)c^k$ for $g \in F_2$ and $c \in \mathbb{Z}$, and $\Psi := F^{-1} \circ \Phi \circ F \circ \iota_g$, where $\iota_g \in \mathrm{Inn}(F_2 \times \mathbb{Z})$. Because c is central, $\iota_g(c) = c$.

To prove the second part, we consider $w' = f(b)$:

$$\phi(w') = \phi(f(b)) = f\circ\psi\circ\iota_{g^{-1}}\circ f^{-1}(f(b)) = f(\psi(b^{g^{-1}})) = f(b)^{f(\psi(g^{-1}))} = (w')^{f(\psi(g^{-1}))}.$$

1. If $\phi_{ab} \neq \mathrm{id}$, then since both w_{ab} and w'_{ab} are eigenvectors of ϕ_{ab}, w_{ab} and w'_{ab} have the same span in $\mathbb{R}^2$. Therefore $p(\Phi(w)) \neq 0$ if and only if $p(\Phi(w')) \neq 0$, and $p(\Phi(w')) = p(\Psi(b))$, since

$$\Phi(w') = F\circ\Psi\circ\iota_{g^{-1}}\circ F^{-1}(f(b)) = F(\Psi(b^{g^{-1}})) = F(b^{g^{-1}}c^{p(\Psi(b))}) = f(b)^{g^{-1}}c^{p(\Psi(b))}.$$

2. If $\phi_{ab} = \mathrm{id}$, then $[\Phi]$ is conjugate in $\mathrm{Out}(F_2 \times \mathbb{Z})$ to either the identity map or the map $a \mapsto ac^{k_a},\ b \mapsto bc^{k_b}$.

$\square$

Proof of Theorem 4 in Case (ii)a. For the purpose of calculating the Dehn function, in Case (ii)a, we may assume our group has the form

$$M_\Phi = \langle a, b, c, t \mid a^t = ab^\beta c^{k_a}, b^t = bc^{k_b}, c^t = c,\ [a, c] = 1, [b, c] = 1\rangle$$

for $k_b \neq 0$. Consider the subgroup $K \cong \langle a, c \mid [a, c]\rangle \cong \mathbb{Z}^2$. Since K quasi-isometrically embeds in $F_2 \times \langle c\rangle$, and Φ preserves K, by Lemma 2.5.3 $n \mapsto n^2 \max\{|\Phi^n(b)|, |\Phi^n(c)|\} = n^2(k_b n + 1)$ is a lower bound for the Dehn function of M_Φ. This cubic lower bound matches the upper bound from Gersten and Riley's electrostatic model (Corollary 2.4.4). Thus the claim is established. $\square$

We now turn to Case (ii)b. The methods of Case 1 do not apply for ϕ of linear growth. (For us, this corresponds to automorphisms such that the abelianization has a fixed point, but in general it means that the function $n \mapsto \max_{x \in X}\{|\Phi^{\pm 1}(x)|\}$ is dominated by a linear function.) Although the subgroup $H = \langle a, b^{-1}ab\rangle$ is fixed by Φ, Button and R. Kropholler [33] have shown that for ϕ linearly growing, M_ϕ is not strongly hyperbolic relative to any finitely generated proper subgroup. This implies that in general, van Kampen diagrams over M_ϕ will not decompose into $\langle H, t\rangle$ islands with linear area complement. Therefore a new method is required.

From Corollary 3.3.2, to prove Case(ii)b it suffices to prove that M_Φ, given by the presentation

$$\mathcal{P}_{\beta,k_a} = \langle a, b, t, c \mid a^t = ab^\beta c^{k_a},\ b^t = b,\ ac = ca,\ bc = cb,\ ct = tc\rangle$$

has quadratic Dehn function when β and k_a are both non-zero.

We use Gersten and Riley's electrostatic model to get this result, but add the innovation that the diagram will be discharged along *partial corridors* (see Definition 2.3.2). Van Kampen diagrams over $\mathcal{P}_{\beta,k_a}$ have both partial b-corridors and partial c-corridors. Van Kampen diagrams over the presentation $\mathcal{P}_\beta = \langle a, b, t \mid a^t = ab^\beta, b^t = b \rangle$ for M_ϕ have partial b-corridors.

Proof of Theorem 4 in Case(ii)b. Suppose w is a word of length n with respect to $\mathcal{P}_{\beta,k_a}$, representing the identity in M_Φ. Let $\overline{w}$ be w with all $c^{\pm 1}$ removed. Then $w = \overline{w}c^m$ in M_Φ for some $m \in \mathbb{Z}$, where $|\overline{w}| \leq |w|$. Since M_ϕ has a quadratic Dehn function, there exists a minimal area diagram $\overline{\Delta}$ for $\overline{w}$, over the presentation $\mathcal{P}_\beta$, such that $\text{Area}(\overline{\Delta}) \leq C|\overline{w}|^2$. We charge $\overline{\Delta}$ by replacing 2-cells in $\overline{\Delta}$ with 2-cells labeled by the defining relators from $\mathcal{P}_{\beta,k_a}$. What follows is a scheme for adding in 2-cells to discharge the diagram.

The main idea is that if we can pair off oppositely-oriented capping faces that are joined by partial b-corridors, then we can add in partial c-corridors following the b-corridors, as in Figure 3.1, in order to discharge the c-edges in our diagram. A consistent pairing would guarantee that we could replicate the picture in Figure 3.1 throughout the diagram.

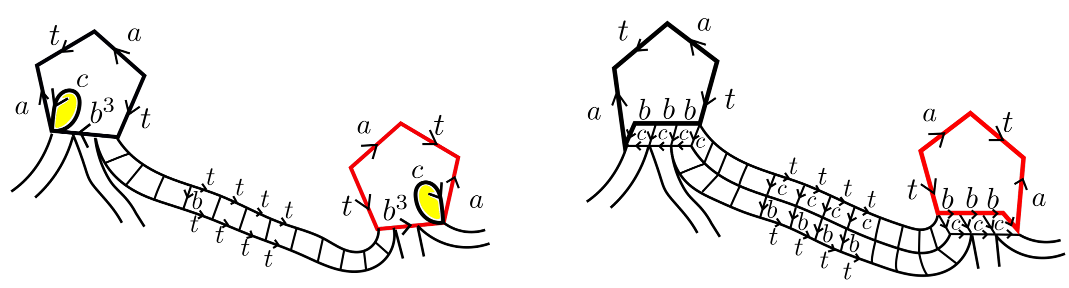

Figure 3.1: If a partial b-corridor joins two capping faces in $\overline{\Delta}$, their c-charges can be discharged by adding partial c-corridors that ‘follow’ the partial b-corridor.

I. Modeling the diagram with a graph. Construct a planar (non-simplicial) graph Γ from $\overline{\Delta}$. Add a vertex for each capping face in $\overline{\Delta}$. If two capping faces are connected by a partial b-corridor, possibly of length zero, add an edge between the corresponding vertices. Two vertices may share multiple edges. We also add an edge and vertex for each partial b-corridor that goes to the boundary. We will want to distinguish vertices that represent an escape to the boundary, so we color them white. Every black vertex in the graph Γ is degree $|\beta|$, and every white vertex is degree 1. This is illustrated in Figure 3.2.

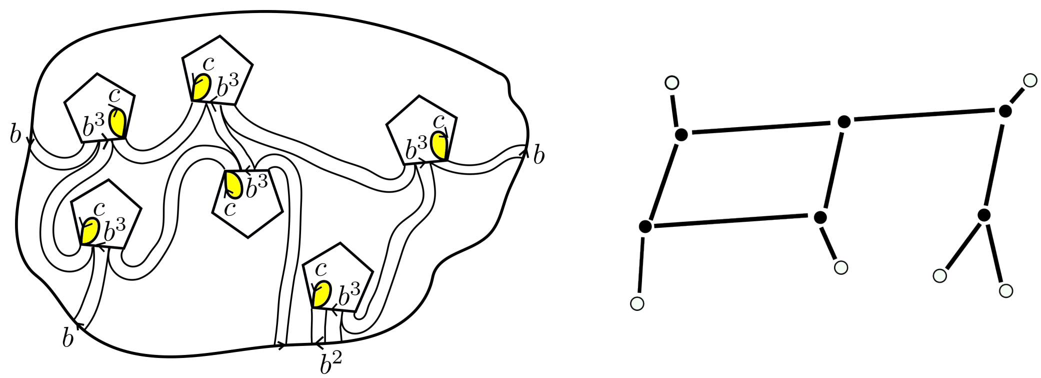

Figure 3.2: From capping faces and partial-corridors in $\overline{\Delta}$, construct a graph Γ. Black vertices correspond to capping faces, white vertices correspond to 1-cells labeled b in $\partial\overline{\Delta}$, and the edges correspond to partial b-corridors.

The graph Γ is naturally bipartite (but not generally black-white bipartite, as you can see in the example above). Partition the vertices of Γ into two types: C, clockwise, and A, anticlockwise vertices. The set of clockwise vertices is made up of those vertices corresponding to capping faces in which the b-edges are oriented clockwise, together with vertices corresponding to clockwise oriented b-edges in $\partial\overline{\Delta}$. The anticlockwise vertex set is defined in the same manner. Partial b-corridors connect b-edges arising in a clockwise oriented 2-cell to b-edges arising in an anticlockwise oriented 2-cell.

A subgraph S_T of a graph T is a *1-factor for T*, if S_T contains all vertices of T, and each vertex is contained in precisely one edge of S_T. That is, S_T is a union of disjoint edges, and every vertex of T is contained in an edge. We recall Hall's Marriage Condition for bipartite graphs, a classic theorem of graph theory. This theorem roughly says that if we have a bipartite graph with vertices partitioned into sets A and B, and we wish to pair every element of A with a neighbor (an element of B), the only obstruction to finding a consistent pairing for all of A and B is if there is some subset $S \subset A$ in which there are fewer candidate partners then there are elements of S. See [42] for a proof:

Theorem 3.3.3 (Hall's Marriage Theorem)**.** *If a graph Γ is bipartite with respect to the vertex partition (A, B), then Γ contains a 1-factor if and only if for any $S \subset A$, $|S| \leq |\text{link}(S)|$, and vice-versa for subsets of B.*

Corollary 3.3.4. *A k-regular bipartite graph with $k \geq 1$ has a 1-factor.*

Proof. It is enough to verify the Marriage Condition. Suppose $S \subset A$. This classic proof uses a clever edge count to establish that $|S| \leq |\text{link}(S)|$. There are $k|S|$ many edges from S to $\text{link}(S)$, since Γ is bipartite and k-regular. These $k|S|$ edges are a subset of all of the $k|\text{link}(S)|$ edges with a vertex in $\text{link}(S)$. Therefore $|S| \leq |\text{link}(S)|$. Likewise, if $T \subset B$, then $|T| \leq |\text{link}(T)|$. $\square$

An obstacle to applying Corollary 3.3.4 is that our graph Γ may not be regular. Therefore we will construct a regular graph $\hat{\Gamma}$ which has Γ as a subgraph.

II. Building a regular bipartite graph. Take $|\beta|$ many copies of Γ, and identify the white vertices in each of the copies. That is,

$$\hat{\Gamma} = \left(\bigsqcup_{i=1}^{|\beta|} \Gamma \times \{i\} \right) / \sim,$$

where $(v, i) \sim (v, j)$ for all i, j when v is a white vertex. This is illustrated in Figure 3.3a. White vertices are degree one, so the identification of $|\beta|$ copies of Γ forces $\hat{\Gamma}$ to be a $|\beta|$-regular graph. Using the same partitions of C and A for each copy of Γ, we see that $\hat{\Gamma}$ is bipartite too, with partitions $\cup_{i=1}^{|\beta|} C_i$ and $\cup_{i=1}^{|\beta|} A_i$.

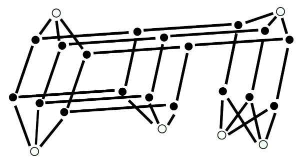

(a) The regular bipartite graph $\hat{\Gamma}$: $|\beta|$ copies of Γ glued together at degree-1 vertices

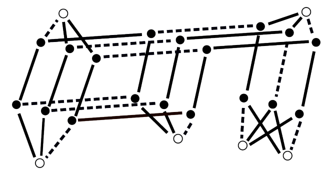

(b) Partners for the vertices of $\hat{\Gamma}$

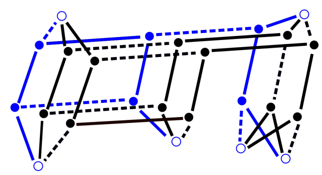

(c) The copy of Γ corresponding to the image of $\Gamma \times \{1\}$ gives a partnering for the vertices of Γ

Figure 3.3: Finding neighbor partners for Γ via Hall's Marriage Theorem

III. Finding pairing partners for b- and c-corridors. Corollary 3.3.4 implies that $\hat{\Gamma}$ has a 1-factor. This assigns to every vertex $v \in \hat{\Gamma}$ a unique partner $v' \in \operatorname{link}(v)$. Noting that the image I of $\Gamma \times \{1\}$ in $\hat{\Gamma}$ is isomorphic to Γ, we consider the partnering on I. As long as $v \in I$ is a black vertex, its partner is also an element of I, since $\operatorname{link}(v) \subset I$. If $v \in I$ is a white vertex, it may not have a partner in I, as $\operatorname{link}(v)$ contains an element from each copy of Γ. The image I is highlighted in blue in Figure 3.3.

IV. Completing to a van Kampen diagram. If v and v' are partnered black vertices in Γ then the corresponding capping faces are connected by at least one partial b-corridor (possibly of length zero). In $\overline{\Delta}$, the capping faces corresponding to v and v' have $|k_a|$ many oppositely oriented c 1-cells. We will make these

1-cells the ends of $|k_a|$ many partial c-corridors, as in Figure 3.1. Since there is at least one partial b-corridor connecting the two capping faces, choose a partial b-corridor joining them. As c is central in M_Φ, partial c-corridors can be run alongside the partial b-corridor, and the word along the edge of the c-corridors will be the same as the word along the edge of the b-corridor, namely some power of t. If one of the paired vertices is white, then the c-corridors follow the b-corridor to the boundary. At the ends of the partial b-corridors it may be necessary to insert rectangles in which the b-corridors and c-corridors cross, as in Figure 3.4, but this requires no more than $|k_a||\beta|\text{Area}(\overline{\Delta})$ additional 2-cells. The total number of 2-cells added to $\overline{\Delta}$ in this process is no more than $(1+|\beta|)|k_a|\text{Area}(\overline{\Delta})$.

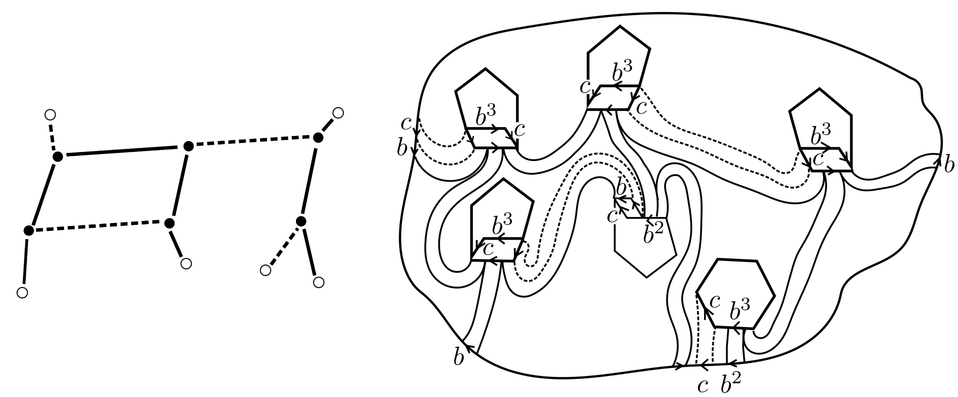

Figure 3.4: Partnering in Γ gives a consistent way to discharge c-charges

V. Correcting the boundary. Partial c-corridors follow partial b-corridors to the boundary in groups of $|k_a|$. Therefore the new boundary has length between $|\overline{w}|$ and $(|k_a|+1)|\overline{w}|$. At this point we have constructed the van Kampen diagram over M_Φ for some word w' in the pre-image of $\overline{w}$. Deleting the c's from w' produces $\overline{w}$, but the arrangement of c's in w' may differ from that in w. As was described in Section 2.4 we glue around the outside of this diagram an annular diagram with the word w' along the inner boundary component and the word w along the outer boundary component. Together, they form Δ, a van

Kampen diagram for w over M_Φ. The annular diagram has area no more than $(|k_a|+1)^2|\overline{w}|^2$, and summing our area estimates, Δ has area no more than $(1+|k_a|)(|\beta|+1)\text{Area}(\overline{\Delta})+(|k_a|+1)^2|\overline{w}|^2$. Since $\text{Area}(\overline{\Delta}) \leq C|\overline{w}|^2$, it follows that there is constant $A > 0$ such that for any given word w in the generators of M_Φ that represents the identity, this construction produces a van Kampen diagram of area less than $A|w|^2$. $\square$

CHAPTER 4

THE BASE GROUP IS $G = \mathbb{Z}^2 * \mathbb{Z}$

4.1 Automorphisms of $\mathbb{Z}^2 * \mathbb{Z}$

The automorphism group of

$$\mathbb{Z}^2 * \mathbb{Z} \;=\; \langle a, b \mid [a,b] = 1\rangle * \langle c\rangle,$$

is generated by inner automorphisms, inversions, graph isomorphisms, and the four transvections

$$\begin{aligned}
\tau_a :&\quad a \mapsto ab, \quad b \mapsto b, \quad\; c \mapsto c,\\
\tau_b :&\quad a \mapsto a, \quad\; b \mapsto ba, \quad c \mapsto c,\\
\psi_a :&\quad a \mapsto a, \quad\; b \mapsto b, \quad\; c \mapsto ca,\\
\psi_b :&\quad a \mapsto a, \quad\; b \mapsto b, \quad\; c \mapsto cb.
\end{aligned}$$

Proposition 4.1.1. *For every $\Psi \in \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$, there is $\Phi \in \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$ of the form*

$$\Phi : \quad a \mapsto \phi(a), \quad b \mapsto \phi(b), \quad c \mapsto cz,$$

where $\phi \in \mathrm{Aut}(\mathbb{Z}^2)$, $z \in \mathbb{Z}^2$, such that M_Ψ and M_Φ have equivalent Dehn functions. Moreover, Φ and Ψ are classified in the same way by Theorem 5.

Proof. Since $\mathrm{Inn}(\mathbb{Z}^2 * \mathbb{Z}) \trianglelefteq \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$, all automorphisms $\Psi \in \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$ can be written as $\iota_g \circ \Phi'$ where $g \in \mathbb{Z}^2 * \mathbb{Z}$ and Φ' is a product of inversions, transvections, and graph isometries in the generating set above. Each of these preserves the subgroup $\langle a, b \mid [a,b] = 1\rangle$, and thus
$\Phi' : a \mapsto \phi(a),\ b \mapsto \phi(b), c \mapsto wc^{\pm 1}x$, for some $\phi \in \mathrm{Aut}(\mathbb{Z}^2)$ and $w, x \in \mathbb{Z}^2$.

Let $\Phi = (\iota_w \circ \Phi')^2$. By Lemma 2.5.1, M_Φ and M_Ψ have equivalent Dehn functions. If Ψ induces $\psi \in \mathrm{GL}(2,\mathbb{Z})$, and the Jordan Canonical Form (JCF) of ψ

is A, then the JCF of ϕ will be either A or A^2. We leave it to the reader to make the easy check that A is finite order if and only if A^2 is and that A has a non-unit eigenvalue if and only if A^2 has one too. Thus Ψ and Φ are classified in the same way by Theorem 5. $\square$

Since $\mathbb{Z}^2$ is abelian, note that $\mathrm{Aut}(\mathbb{Z}^2) = \mathrm{Out}(\mathbb{Z}^2)$. For the purposes of calculating the Dehn function, we may take ϕ to be the most convenient representative of its conjugacy class in $\mathrm{Out}(\mathbb{Z}^2)$:

Proposition 4.1.2. *Suppose Φ has the form of Proposition 4.1.1 above. If ϕ and ψ are conjugate in $\mathrm{Out}(\mathbb{Z}^2)$, there is $\Psi \in \mathrm{Aut}(\mathbb{Z}^2 * \mathbb{Z})$ such that $[\Psi]$ and $[\Phi]$ are conjugate in $\mathrm{Out}(\mathbb{Z}^2 * \mathbb{Z})$, where Ψ also has the form of 4.1.1 and the restriction of Ψ to $\langle a, b\rangle$ is ψ.*

Proof. If $[\phi]$ and $[\psi]$ are conjugate in $\mathrm{Out}(\mathbb{Z}^2) = \mathrm{Aut}(\mathbb{Z}^2)$, then for some $f \in \mathrm{Aut}(\mathbb{Z}^2)$, $\psi = f^{-1} \circ \phi \circ f$. Define $\Psi := \xi^{-1} \circ \Phi \circ \xi$, where ξ restricts to f on $\mathbb{Z}^2$ and $c \mapsto c$. By Lemma 2.5.1, M_Φ and M_Ψ are isomorphic, and so have equivalent Dehn functions. The map Ψ has the appropriate form:

$$\Psi(c) = \xi^{-1} \circ \Phi \circ \xi(c) = \xi^{-1}(\Phi(c)) = \xi^{-1}(cy) = cf^{-1}(y).$$

$\square$

In summary, for the purpose of calculating Dehn functions, it suffices to consider Φ of the form $\Phi(a) = \phi(a),\ \Phi(b) = \phi(b)$, and $\Phi(c) = cz$ where $\phi \in \mathrm{Aut}(\mathbb{Z}^2)$ and $z \in \mathbb{Z}^2$. Further, by Proposition 4.1.2, ϕ may be chosen to be any convenient representative of the conjugacy class of ϕ in $\mathrm{Out}(\mathbb{Z}^2)$. The value z will change to some other $z' \in \mathbb{Z}^2$, but the Dehn function will not depend on the value of z.

4.2 Corridors in reduced van Kampen diagrams over M_Φ

The defining relations of $\langle a, b, c, t \mid [a, b] = 1, a^t = \phi(a), b^t = \phi(b), c^t = cz\rangle$ imply that the van Kampen diagrams of M_Φ can have both c- and t-corridors. All of the 2-cells which make up c-corridors also appear in t-corridors, because the defining relators which include the letter c also include the letter t.

Definition 4.2.1. Suppose that τ is a t-corridor and η is a c-corridor. Let $\hat{\tau} \subset \tau$ and $\hat{\eta} \subset \eta$ be subcorridors. Then $\hat{\tau}$ and $\hat{\eta}$ form a *bigon* when their only common 2-cells are their first and last ones.

Lemma 4.2.2. *Let τ be a t-corridor and η be a c-corridor. If τ and η intersect more than once, then either η is not reduced or there are subcorridors $\hat{\tau} \subset \tau$ and $\hat{\eta} \subset \eta$ forming a bigon.*

Proof. Suppose η and τ cross at least twice. τ is covered by a family of 2-cell arcs that intersect η only at the first and last 2-cell in the arc. These arcs live 'above' η and cross the top of the corridor as in Figure 4.1, or live 'below' η and cross the bottom. If there are at least 2 intersections of τ and η, there must be at least one arc connecting a t and a t^{-1} edge in η.

Consider a τ-arc T_1 that lies above η. Either it forms a bigon with η and we are done, or between the ends of T_1, η contains a 2-cell from τ. Successive arcs containing this 2-cell live on opposite sides of η. Let T_2 be the arc above η. T_2 cannot cross T_1. Either T_2 forms a bigon with a subcorridor of η, or there is another bigon nested between T_2 and η, and so on. Nesting must eventually end (after all there are only finitely many 2-cells in η between the ends of T_1). We eventually find either a bigon, or a pair of non-reduced 2-cells. This is illustrated in Figure 4.1. $\square$

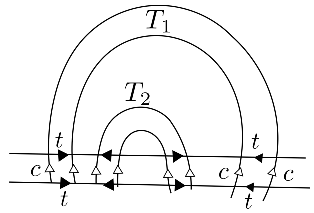

Figure 4.1: An innermost crossing is a bigon for τ and η.

Lemma 4.2.3. *In a van Kampen diagram where c-corridors are reduced, if a t-corridor, τ, intersects a c-corridor, η, it will do so only once.*

Proof. Since c-corridors are made up of a single kind of 2-cell (arising from the defining relation $c^t = cz$), all 2-cells in a reduced c-corridor are identical with respect to the corridor. Let us assume for the contradiction that η is reduced and that τ and η intersect at least twice.

By Proposition 4.2.2, there exist subcorridors $\hat{\tau}$ and $\hat{\eta}$ that form a bigon, with precisely the first and final 2-cells, E_1 and E_2, in common. The orientation of the edges labeled by t in E_1 fixes an orientation of t 1-cells in $\hat{\eta}$ since $\hat{\eta}$ is reduced. It also fixes an orientation of t 1-cells in $\hat{\tau}$. Both $\hat{\eta}$ and $\hat{\tau}$ then specify an orientation for the t 1-cells in E_2, but these two specifications are inconsistent. In Figure 4.2, E_1 is on the left and E_2 is on the right. The t-edges in E_1 force the t-edges in η to point towards E_2. For η to be reduced, the t 1-cells in E_2 must point down. However, from E_1, the t-edges of τ always point up. $\square$

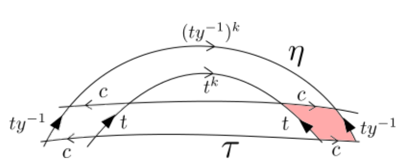

Figure 4.2: If a t-corridor and a c-corridor cross at least twice, the c-corridor cannot be reduced.

We will use the same argument for alternating corridors and c-corridors in Lemma 4.4.8(1) and for α- and t- partial corridors in Lemma 4.4.8(4).

Corollary 4.2.4. *In a van Kampen diagram with reduced c-corridors, there are no c-annuli, and t-annuli do not intersect c-corridors.*

Proof. The word around the outside of a c-annulus contains t's, so it would have to intersect once (and therefore intersect at least twice) with a t-corridor, which is impossible by Lemma 4.2.3. Similarly, if a c-corridor intersected a t-annulus, it would have at least two common cells. □

The following important corollary tells us the length of c-corridors in a diagram Δ from the boundary word $\partial\Delta$ and the way the c-edges are paired up (the so-called *c-corridor pairing*, see Definition 4.4.2).

Corollary 4.2.5. *In a diagram Δ with reduced c-corridors, suppose that the boundary word is $w_1 c^{\mp 1} w_2 c^{\pm 1}$, for some words w_1, w_2 in the generators of M_Φ and that a c-corridor η begins and ends on the two distinguished c-edges. Then the length of η is given by the absolute value of the index sum of t in w_1 (equivalently, in w_2).*

Proof. All t-corridors intersecting η have the same orientation with respect to η. In particular, the word along one side of η is t^k for some k, without any free reductions. Thus the t-corridors starting at t-edges in w_2 that are oppositely oriented to the t's in η cannot cross it, and so must be partnered on the same side of η. This leaves exactly the absolute value of the index-sum of t in w_1 many t-corridors which have no partners on the same side of η, and so must cross it. By Lemma 4.2.3, each of these t-corridors can cross η exactly once. □

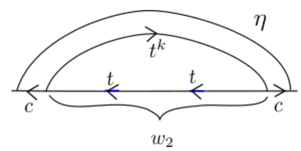

Figure 4.3: The t-corridors of oppositely oriented t-edges in w_2 cannot cross η.

4.3 The exponential Dehn function when ϕ has a non-unit eigenvalue

The map ϕ has a non-unit eigenvalue if and only if ϕ has exponential growth. (See Section 2.6 for a discussion of growth in free abelian groups.)

Proof of Theorem 5(1). From Lemma 2.5.3, as $\langle a, b \rangle = \mathbb{Z}^2$ quasi-isometrically embeds in $\mathbb{Z}^2 * \mathbb{Z}$, when ϕ has exponential growth on $\langle a, b \rangle$, the Dehn function of M_Φ is bounded below by an exponential function. From Lemma 2.5.2, the Dehn functions of mapping tori of RAAGs are always bounded above by exponential functions. Therefore M_Φ has an exponential Dehn function. □

4.4 The cubic Dehn function when ϕ has linear growth

Recall from Section 4.1 that for the purpose of calculating Dehn functions, it suffices to consider Φ of the form $\Phi(a) = \phi(a)$, $\Phi(b) = \phi(b)$, and $\Phi(c) = cz$ where $\phi \in \mathrm{Aut}(\mathbb{Z}^2)$ and $z \in \mathbb{Z}^2$. Further, ϕ may be chosen to be any convenient representative of the conjugacy class of $[\phi]$ in $\mathrm{Aut}(\mathbb{Z}^2)$. Given that $\phi \in \mathrm{Aut}(\mathbb{Z}^2)$ is of linear growth, Lemma 2.6.1 implies that ϕ is conjugate in $\mathrm{Aut}(\mathbb{Z}^2)$ to the automorphism $a \to ab^k$ and $b \to b$, for some $k \neq 0$. Therefore for the purposes of determining the Dehn function of M_Φ, we assume that $M_\Phi = M_{k,l,m}$ has the presentation

$$\mathcal{P}_{k,l,m} = \langle a, b, c, t \mid a^t = ab^k,\ b^t = b,\ c^t = ca^l b^m,\ [a, b] = 1\rangle$$

for some $k, l, m \in \mathbb{Z}$ with $k \neq 0$.

For the group N_l presented by $\mathcal{P}_l = \langle a, c, t \mid a^t = a, c^t = ca^l\rangle$, there is a short exact sequence:

$$1 \rightarrow \langle\langle b\rangle\rangle \rightarrow M_{k,l,m} \rightarrow N_l \rightarrow 1$$

This is not a central extension, as b does not commute with c in $M_{k,l,m}$, but we will use a variant of the electrostatic model (see Section 2.4) to promote upper bounds on the areas of diagrams over $\mathcal{P}_l$ to upper bounds on the Dehn function of $M_{k,l,m}$. The following definitions will be useful:

Definition 4.4.1. For w a word representing the identity in M_Φ, there are equal numbers of c's and c^{-1}'s in w. A *c-pairing* is any pairing of c's and c^{-1}'s in w.

Definition 4.4.2. The set of c-corridors of a diagram for w induces a c-pairing: a c and a c^{-1} are paired if and only if they are joined by a c-corridor in Δ. We say that a pairing P is a *valid c-pairing* if there exists a van Kampen diagram for w in which the set of c-corridors induce P. Not all pairings are valid c-pairings. Further, a word may have multiple valid c-pairings.

Lemma 4.4.3. *Let w be a word in a, b, c, and t that represents the identity in $M_{k,l,m}$, and let Δ be a minimal area diagram for w. Let $\overline{w}$ be w with all occurrences of b and b^{-1} removed. Then there exists a van Kampen diagram Θ for $\overline{w}$ with respect to the presentation $\mathcal{P}_l$ such that Θ and Δ induce the same c-pairing.*

Proof. To construct a van Kampen diagram for $\overline{w}$ with respect to the presentation $\mathcal{P}_l$, take a polygonal loop in the plane labeled by $\overline{w}$. Following the

c-corridor pairing of w, insert c-corridors: every 2-cell with boundary label $t^{-1}c^{-1}tca^lb^m$ in Δ is inserted into the new diagram as a 2-cell with boundary label $t^{-1}c^{-1}tca^l$. The empty regions between these corridors, *c-complementary regions*, have boundary words $\overline{w_i}$ on a and t. They represent the identity in N_l (and so in $\langle a, t \mid [a, t] = 1 \rangle$) since there are corresponding loops in Δ labeled by words w_i such that $\overline{w_i}$ is w_i with all $b^{\pm 1}$ removed. The resulting van Kampen diagram Θ for $\overline{w}$ over $\mathcal{P}_l$ induces the same c-pairing as Δ. $\square$

It will be possible to charge and discharge Θ to make a new van Kampen diagram for w which has only cubic area, as long as we can show that Θ has quadratic area. Unfortunately, we cannot just appeal to the quadratic Dehn function of N_l, as the corridor pairing of Θ may differ from the corridor pairing in a minimal area diagram for $\overline{w}$. It is important to preserve the c-pairing of Δ in Θ, because when we charge the diagram we will only be able to move b's within c-complementary regions: b does not commute with c in $M_{k,l,m}$.

We first establish the quadratic area of Θ in the special case of $M_{1,1,0}$.

Lemma 4.4.4. *If w represents the identity in $M_{1,1,0}$, and Δ is a minimal area van Kampen diagram for w, then the associated diagram Θ, (from the proof of Lemma 4.4.3) for $\overline{w}$, has area less than $A_1|w|^2$ for some constant $A_1 > 0$.*

While Θ is a van Kampen diagram for $\overline{w}$ with respect to the presentation $\mathcal{P}_1 = \langle a, c, t \mid a^t = a, c^t = ca \rangle$, we will find it easier to work with

$$\mathcal{P}_1' = \langle \alpha, t, c \mid \alpha^t = \alpha, t^c = \alpha \rangle.$$

$\mathcal{P}_1'$ presents the same group as $\mathcal{P}_1$ via $\alpha = ta^{-1}$. We begin by considering corridor and corridor-like behavior in van Kampen diagrams over $\mathcal{P}_1'$.

Notation 4.4.5. A 2-cell in $\mathcal{P}_1'$ that has an edge labeled by c or c^{-1} will be referred to as a c-face. The labels around the boundaries of c-faces are cyclic permutations of $c^{-1}tc\alpha^{-1}$ and $c^{-1}\alpha ct^{-1}$.

Van Kampen diagrams over $\mathcal{P}_1'$ have c-corridors as well as partial α- and t-corridors (see Definition 2.3.2). The partial α- and t-corridors fit together in an alternating way: when a partial α-corridor ends at a cell labeled $t^c = \alpha$ in the interior of a diagram, a partial t-corridor will begin, and when this ends, another partial α-corridor will begin. The alternation of partial corridors continues either to the boundary, or until the *alternating corridor* closes up in an annulus.

Definition 4.4.6. An *alternating corridor* in a van Kampen diagram with respect to $\mathcal{P}_1'$ is the maximal union of partial α-partial corridors, t-partial corridors and the c-faces between them, described above. See Figure 4.4.

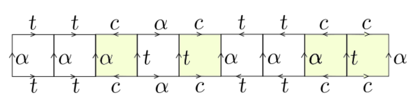

Figure 4.4: An alternating corridor

Alternating corridors differ from standard corridors in that they can intersect one another. Like standard corridors, alternating corridors have a fixed orientation: the bottoms of all partial corridors in an alternating corridor all lie on the same side of the corridor, which we will call the bottom of alternating corridor. This implies that the methods of Lemma 4.2.3 also apply to the intersections of alternating corridors.

Lemma 4.4.7. *Any van Kampen diagram $\mathcal{D}$ over $\langle \alpha, c, t \mid \alpha^t = \alpha, t^c = \alpha \rangle$ is a union of alternating corridors.*

Proof. Every 1-cell in $\mathcal{D}$ labeled by α or t is contained in an alternating corridor.

Since every 2-cell associated to our relator set contains 1-cells labeled by α and t, every 2-cell is contained in an alternating corridor. $\square$

Lemma 4.4.8. *Suppose $\mathcal{D}$ is a van Kampen diagram over $\langle \alpha, c, t \mid \alpha^t = \alpha, t^c = \alpha \rangle$ in which all c-corridors are reduced, and all α- and t-partial corridors are reduced (See Figures 4.5, 4.6). Then in $\mathcal{D}$:*

1. *A c-corridor η and an alternating corridor τ can cross at most once.*
2. *Alternating corridors do not form annuli.*
3. *A single alternating corridor can never cross itself.*
4. *Two alternating corridors cannot cross more than once.*

Figure 4.5: Non-reduced subdiagrams occurring in $\mathcal{D}$

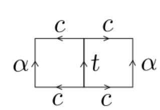

Figure 4.6: Non-reduced subdiagrams not occurring in $\mathcal{D}$

(a) c-corridors are reduced in $\mathcal{D}$

(b) α- and t- partial corridors are reduced in $\mathcal{D}$

Proof. For (1) it suffices (see Lemma 4.2.2) to prove that it is impossible to have a bigon of an alternating corridor τ and a c-corridor η in $\mathcal{D}$. Since c-corridors in $\mathcal{D}$ are reduced, the top of the c-corridor is labeled by a power of α without any free reduction. As in our proof of Lemma 4.2.3, τ and η specify inconsistent orientations for the t edge in the second common 2-cell, as in Figure 4.7a.

For (2), suppose that there is such an annulus. It cannot contain any c-faces, as this would force a c-corridor to cross the alternating annulus twice. If our annulus contains no c-faces, then it is either a t- or α-annulus. The word along the top of the annulus is a power of α or t, respectively. Such an annulus would imply t or α have finite order, but both are infinite order elements of N_1.

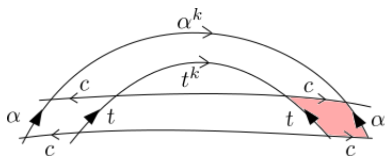

(a) c-corridors cannot double-cross alternating corridors

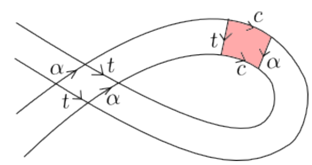

(b) Alternating corridors do not self-intersect.

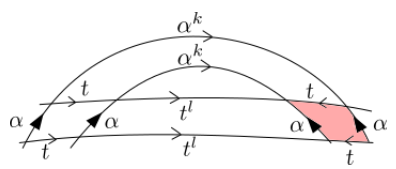

(c) α- and t- partial corridors cannot cross more than once.

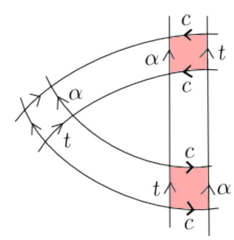

(d) Distinct alternating corridors cannot cross more than once.

Figure 4.7: Impossible behavior for alternating corridors.

For (3), suppose for a contradiction that an alternating corridor η has a self-intersection. An alternating corridor can only have a self-intersection at a 2-cell corresponding to the relation $[\alpha, t] = 1$. Let $\hat{\eta} \subset \eta$ be a subcorridor of η that begins and ends at the self-intersection. Call this first and final 2-cell E.

The 2-cell E is part of both t- and α- partial corridors in $\hat{\eta}$; therefore $\hat{\eta}$ contains at least one c-face (in particular, an odd number of c-faces in order to get both an α- and t-segment at the intersection). Each c-face in $\hat{\eta}$ is part of a c-corridor. By (1), c-corridors can only cross $\hat{\eta}$ once, but each c-corridor must cross $\hat{\eta}$ at least twice, since $\hat{\eta}$ is an annulus.

For (4), assume for a contradiction, that two alternating corridors cross at least twice. Again, we can find a bigon of alternating corridors. There are two cases. In one, no c-corridors intersect the bigon. In this case, one of the alternating corridors is a partial t-corridor, and the other is a partial α-corridor. An argument like Lemma 4.2.3 shows that this kind of double intersection is impossible when t- and α partial corridors are reduced (see Figure 4.7c). In the other case, at least one c-corridor intersects the bigon. We look at the triangle formed by the two bigons and the first c-corridor to cross them. Since it is the first such c-corridor, we have an α- and t-partial corridor that both need to end on the same side of a c-corridor. However, c-corridors always have t's along the bottom and α's along the top — there cannot be both α's and t's on the same side of the c-corridor. Figure 4.7d illustrates this contradiction. Therefore neither case happens. □

Proof of Lemma 4.4.4. We rewrite the boundary word $\overline{w}(a, c, t)$ in terms of $\mathcal{P}_1'$ by substituting in $\alpha^{-1}t$ for each occurrence of a in $\overline{w}$. Add in the c-corridors from Θ using the new c-faces, and then fill the complementary regions with minimal area subdiagrams over $\langle \alpha, t \mid \alpha^t = \alpha \rangle$. Call this diagram Θ'. By construction, the c-corridors are reduced. Since the complementary regions are filled with minimal area subdiagrams, both α- and t- partial-corridors are reduced.

Lemma 4.4.8 implies that the length of an alternating corridor $\mathcal{A}$ in our diagram is bounded above by the number of c-corridors and alternating-corridors that intersect $\mathcal{A}$. There are fewer than $|\overline{w}|/2$ many such corridors. Similarly, the length of each c-corridor is no more than $|w|/2$, by Lemma 4.2.5,and there are no more than $|w|/2$ c-corridors. Thus altogether, $\text{Area}(\Theta') \leq \frac{|\overline{w}|^2+|w|^2}{4}$. The length $|\overline{w}|$ is taken in $\mathcal{P}_1'$, and this is no more than twice the length of $|\overline{w}|$ in $\mathcal{P}_1$, which is less than $|w|$. Since $[t, ta^{-1}]$, which represents the defining relator of $\mathcal{P}_1'$, $[t, \alpha]$, can be filled using a single 2-cell in $\mathcal{P}_1$, the area bound we get for the van Kampen diagram for $\mathcal{P}_1'$ is an upperbound on the area for a minimal area van Kampen diagram for $\mathcal{P}_1$. Therefore, it follows that the area of Θ is less than $A_1|w|^2$ for some constant $0 < A_1 < 2$. $\square$

Lemma 4.4.9. *Suppose $k, l, m \in \mathbb{Z}$ and $k \neq 0$. If w represents the identity in $M_{k,l,m}$, and Δ is a minimal area van Kampen diagram for w, then the associated diagram for $\overline{w}$, Θ_l, (of our proof of Lemma 4.4.3) has area less than $A_l|w|^2$ for some $A_l > 0$ that depends only on l.*

For notational convenience, until Lemma 4.4.9 is proved, N_1 will have the presentation $\mathcal{P}_1^\tau = \langle a, c, \tau \,|\, a^\tau = a, c^\tau = ca\rangle$ and N_l will have the usual presentation $\mathcal{P}_l = \langle a, c, t \mid a^t = a, \ c^t = ca^l\rangle$. By identifying $t = \tau^l$, we get an isomorphism of N_l with the index l subgroup of N_1 generated by a, c, and τ^l.

The basic outline of our proof is as follows: We rewrite the word $\overline{w}(a, c, t)$ in $\mathcal{P}_l$ as a word $\overline{w}(a, c, \tau^l)$ in $\mathcal{P}_1^\tau$. We will fill a planar polygonal path with boundary labeled by $\overline{w}(a, c, \tau^l)$ to construct a van Kampen diagram over $\mathcal{P}_1^\tau$. We first copy in the c-corridors from Δ. This leaves a family of complementary regions to fill. From Lemma 4.4.4, we get a filling for each of these regions with respect to $\mathcal{P}_1^\tau$. Based on this diagram over $\mathcal{P}_1^\tau$, we then construct a diagram over $\mathcal{P}_l$.

Before we prove Proposition 4.4.9, we examine how to build a filling over $\mathcal{P}_l$ from one over $\mathcal{P}_1^\tau$.

Lemma 4.4.10. *If $w(a, \tau^l)$ represents the identity in N_1 and has an area N van Kampen diagram $\mathcal{D}$ over $\langle a, \tau \mid a^\tau = a\rangle$, then $w(a, t)$ has a van Kampen diagram with respect to $\langle a, t \mid a^t = a\rangle$, with area at most N.*

Proof. A τ*-segment* will be a grouping of l consecutive τ's in the boundary that comes from the natural grouping of writing w as a word in a and τ^l. Such segments have a natural orientation that agrees with the orientation of the component τ's. We will find another filling of w for which the induced τ-pairing respects τ-segments, that is, this new filling induces a τ^l-pairing.

For τ and τ^{-1} that are paired by a τ-corridor, let $\widehat{w}$ be the subword of w between them, as in Figure 4.8a. If a first $\tau^{\pm 1}$ in a τ-segment is paired by a corridor to a $\tau^{\mp 1}$ in position i, then the τ-index sum of $\widehat{w}$ will be a multiple of τ only if $i = 1$. The τ-index sum of $\widehat{w}$ must be 0; since τ-corridors cannnot cross, all τ and τ^{-1} in $\widehat{w}$ must be paired within $\widehat{w}$. Thus in any filling, the first τ in a τ-segment must end on the first τ in an oppositely oriented τ-segment.

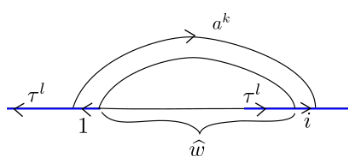

(a) A τ in position 1 in a τ-segment can only pair with another initial τ

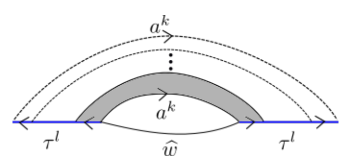

(b) Initial τ corridors provide a guide for joining the rest of the segment

Figure 4.8: In a word w on a and τ^l, there is a valid τ-pairing that pairs whole τ-segments.

Begin with a planar loop labeled by $w(a, \tau^l)$. Add in copies of the τ corridors from $\mathcal{D}$ for the first τ's in each segment. This divides the diagram into regions that remain to be filled. In each region, the index sums of a and τ are both zero, since they are zero in $\mathcal{D}$. If two initial τ's are joined by a τ-corridor, then they both have $l-1$ aligned τ's. We add τ-corridors into the diagram that match the ith τ's in each segment and follow the initial τ-corridor, as in Figure 4.8b

After adding all the τ-corridors, no further 2-cells are required. The area of this diagram will be l times the sum of the initial τ-corridor contributions, and so in particular, the area of the new diagram is at most lN. Moreover, this filling can be used to produce a filling for the equivalent word in $\mathcal{P}_l$. By construction, every τ-corridor has the same top and bottom as its initial τ-corridor. This allows us to change to the relator $[a, t] = [a, \tau^l]$ from $[a, \tau]$. The area of this new diagram is at most N. $\square$

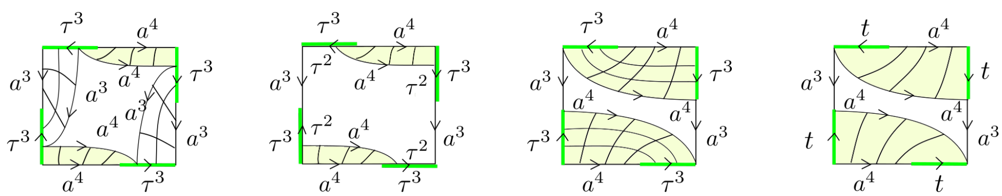

(a) Initial filling $\mathcal{D}$

(b) Refilling all initial τ-corridors

(c) Filling following the initial τ-corridors

(d) Grouping τ^l's to get a filling in $\mathcal{P}_l$

Figure 4.9: A toy example of the procedure of Lemma 4.4.10

Proof of Lemma 4.4.9. The word $\overline{w}(a, c, t)$, which we get from removing all $b^{\pm 1}$ from w, represents the identity in N_l. Define $\overline{v} = \overline{w}(a, c, \tau^l)$ — that is, substitute a $(\tau^l)^{\pm 1}$ for every $t^{\pm 1}$ in $\overline{w}$. Lemma 4.4.4 implies that the van Kampen diagram Θ_1 for $\overline{v}$ in the relators of $\mathcal{P}_1^\tau$ has area less than $A_1|\overline{v}|^2 \leq A_1 l^2 |\overline{w}|^2$. In this proof, we will show that guided by Θ_1, we can construct a van Kampen diagram Θ_l for $\overline{w}$ over $\mathcal{P}_l$ which has comparable area.

First we show that the c-corridors in the diagram over $\mathcal{P}_1^\tau$ are consistent with a diagram over $\mathcal{P}_l$. The diagram Θ_1 comes with a rooted c-corridor dual tree (see Section 2.3). To every complementary region R we can associate a corridor — in particular the corridor shared between R and its parent complementary region.

This matches regions with a distinguished c-corridor. We show that the lengths of c-corridors are multiples of l, starting with corridors associated to the leaves of the tree and working up the tree.

If C is a c-corridor for a leaf complementary region, recall that the length of C is given by the index-sum of τ in the boundary between the paired c-edges. Since in $\overline{v}$, τ only appears in multiples of l, C will have length a multiple of τ. This implies that the leaf complementary region can be filled using the relators of $\mathcal{P}_l$.

We begin building Θ_l by putting in the leaf c-corridors. If there are nl many 2-cells in Θ_1, there will be n 2-cells in Θ_l. This argument extends up the tree, so that we can fill in all of the c-corridors in Θ_l, giving Θ_l the same c-corridor pairing as Θ_1. Corresponding c-corridors in the two diagrams differ in length by exactly the factor l. We will promote the fillings of complementary regions of the c-corridors in the relators of $\mathcal{P}_1^\tau$ to fillings in the relators of $\mathcal{P}_l$.

We consider a single c-complementary region. In Θ_1, the word around the perimeter of this region is of the form $x_0((\tau a^{\epsilon_1})^l)^{k_1} x_1 \ldots x_n((\tau a^{\epsilon_n})^l)^{k_n}$, where $\epsilon_i \in \{0, -1\}$, for all i, and x_i is a subword of $\overline{v}$. We have a filling in $\mathcal{P}_1^\tau$ for this word. The corresponding perimeter in Θ_l, for which we want a $\mathcal{P}_l$-filling, is labeled by $x_0(ta^{\epsilon_1 l})^{k_1} x_1 \ldots x_n(ta^{\epsilon_n l})^{k_n}$. In the generators a, b, τ, this perimeter is $x_0(\tau^l a^{\epsilon_1 l})^{k_1} x_1 \ldots x_n(\tau^l a^{\epsilon_n l})^{k_n}$. These regions are easy to fill.

- Let L be a minimal area diagram for the word $(\tau a)^{-l}\tau^l a^l$, as in Figure 4.11b. Area$(L) < l^2$.
- We extend the fillings of complementary regions in Θ_1, by gluing copies of L along the perimeter, as in Figure 4.11c, which replaces occurrences of $(\tau a)^l$ (along the tops of c-corridors) with $\tau^l a^l$. Across all complementary

regions, this adds no more than $l^2|\overline{w}|^2$ area.

- Apply Lemma 4.4.10 to this new diagram to produce a van Kampen diagram over $\mathcal{P}_l$, which has the same perimeter word as the complementary region in Θ_l that we want to fill.

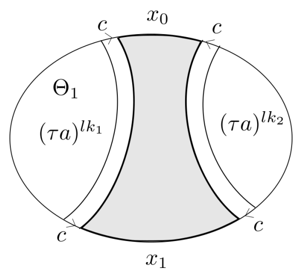

Figure 4.10: The initial diagram Θ_1

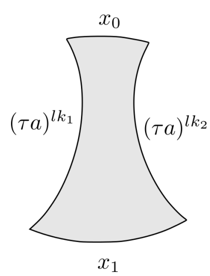

(a) One c-complementary region in Θ_1

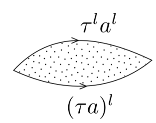

(b) The diagram L. Area$(L) \sim l^2$

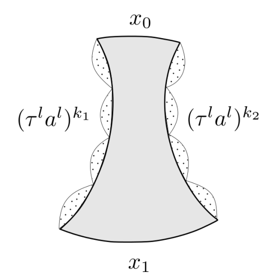

(c) Gluing copies of L along edges of the tops of c-corridors in Θ_1

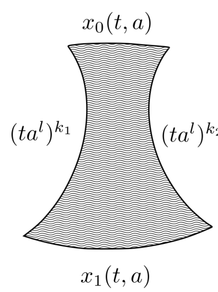

(d) A new filling from Lemma 4.4.10, in terms of t

Figure 4.11: Converting a filling Θ_1 in $\mathcal{P}_1^\tau$ to one in Θ_l

After we get appropriate fillings for each of the complementary regions in Θ_l, we glue them in. The diagram remains planar, so it is a van Kampen diagram for w with area less than or equal to $(A_1 + |l|^2)|w|^2$. Let $A_l = A_1 + |l|^2$. $\square$

Proof of Theorem 5(3). Recall that the c-corridor complementary regions, which have perimeters labeled by words in the commuting element a and t were filled with minimal area diagrams. The boundary word around the perimeter of each complementary region is no longer than $(|l| + |m| + 2)|w|$, since the perimeter is made up of c-corridors and the boundary. The c-corridors in total have length no more than $|w|$, from Lemma 4.2.5, so contribute at most length $(|l| + |m| + 1)|w|$ to the perimeter word. Papasoglu's theorem, Theorem 2.4.2, implies that for each complementary region we can find a maximal tree in the 1-skeleton such that the distance in the tree from any vertex to a root on the boundary is less than $B|w|$ for some constant $B > 0$.

We apply the electrostatic model of Section 2.4. For each complementary region, we replace all of the 2-cells coming from $\mathcal{P}_l$ with the corresponding 2-cells coming from $\mathcal{P}_{k,l,m}$. By Lemma 4.4.9, there are only quadratically many new 1-cells labeled by b's added in total, across all of the complementary regions. Since there are at most $A_l|w|^2$ many b-charges and each one requires the addition of a b-partial corridor of length no more than $B|w|$ to connect to the root in their region, it takes only area $A_l B|w|^3 + A_l|w|^2$ to make each complementary region a van Kampen diagram over $\mathcal{P}_{k,l,m}$. Annular regions are added to the outside of each of the complementary regions to rearrange the b-cells to match the word along the perimeter of the c-corridor complement in Δ. We excise the c-corridor complementary regions of Δ and replace them with these new fillings. The new diagram has area less than $A_l B|w|^3 + (A_l + |l| + |m| + 1)|w|^2$, so Δ is a cubic area diagram, and $M_{k,l,m}$ has a cubic Dehn function. □

4.5 The quadratic Dehn function when ϕ has finite order

We explain here how to apply the methods of Section 4.4 to groups of the form

$$M_{k,l} = \langle a, b, c, t \mid [a, b] = [a, t] = [b, t] = 1,\ c^t = ca^k b^l \rangle.$$

We first establish the quadratic Dehn function for groups $M_{k,0}$, and then show that every $M_{k,l}$ is a finite-index subgroup of some $M_{\hat{k},0}$. This implies Theorem 5(1), that these groups have quadratic Dehn function. Another possible mode of attack is to prove the stronger statement that the groups $M_{k,l}$ are CAT(0), but we do not currently have a proof of this.

Proposition 4.5.1. *The family of groups $M_{k,0}$ all have quadratic Dehn function.*

Proof. Let w be a word representing the identity in $M_{k,0}$, with a minimal area diagram Δ. Let $\overline{w}$ be w with all $b^{\pm 1}$'s removed. By Lemma 4.4.3, there is a diagram Θ over $\langle a, c, t \,|\, [a, t] = 1, c^t = ca^k \rangle$ which induces the same c-pairing as Δ. Lemma 4.4.9 implies that Θ will be a diagram with area at most $A_k|w|^2$. Using Θ, we construct a diagram for w with quadratic area.

Next we apply the electrostatic model, and notice that no b's occur in any of the 2-cells of $M_{k,0} = \langle a, b, c, t \mid [a, b] = [a, t] = [b, t] = 1,\ c^t = ca^k \rangle$. Still, the word w may contain some $b^{\pm 1}$'s, we just have that the index sum of b around the perimeter of each c-complementary region is zero. Further, b's only appear in the perimeter along $\partial\Delta$. Note that the total perimeter of the i-th complementary region has length no more than $(|k| + 1)|w|$. The perimeter consists of a length $L_i > 0$ portion, which is part of $\partial\Delta$, the boundary of the full van Kampen diagram, and the rest is made up of sides of c-corridors. In total there are no more than $\sum_i L_i < |w|$ many such b's. Gluing in an annulus around each complementary region, which creates canceling pairs of b's, and rearranges them to agree with $\delta\Delta$ only adds $\sum_i L_i|w|(|k|+1) \leq |w|^2(|k|+1)$ much area to the diagram, giving a total area of $(A_k + |k| + 1)|w|^2$. $\square$

Proof of Theorem 5(1). If k and l are relatively prime, then there is a pair of integers (x, y) such that $ky - lx = 1$, by Bezout's identity. This implies that there is a basis A, B with $A = a^k b^l$ and $B = a^x b^y$, for which our group has the presentation $\langle A, B, c, t \mid [A, B] = 1, A^t = A,\ B^t = B,\ c^t = cA\rangle$, which is equivalent to the presentation $M_{1,0}$. Otherwise, $\gcd(k, l) = n > 1$, and $M_{k,l} = M_\Phi$ is a subgroup of index n of $M_\Psi = M_{\frac{k}{n},\frac{l}{n}}$, since $\Psi^n = \Phi$. Now M_Ψ has a quadratic Dehn function. By Lemma 2.5.1, $M_{k,l}$ is isomorphic to a finite index subgroup of $M_{\frac{k}{n},\frac{l}{n}}$, so $M_{k,l}$ also has a quadratic Dehn function. $\square$

CHAPTER 5

FURTHER ESTIMATES AND FUTURE WORK

In this chapter, we present some partial results on Dehn functions of mapping tori of RAAGs. In the first section, we consider how our arguments extend from $\mathbb{Z}^2 * \mathbb{Z}$ to $\mathbb{Z}^n * \mathbb{Z}$. We fall short of having an inductive argument, but to demonstrate that our techniques apply more broadly than to just $\mathbb{Z}^2 * \mathbb{Z}$, we classify the Dehn functions for mapping tori with base group $\mathbb{Z}^3 * \mathbb{Z}$. In Section 5.2 we calculate the Dehn functions of mapping tori with base group $F_k \times F_l$, when $k, l > 1$, and we consider the challenges for the case when $k > 1$ and $l = 1$, which generalizes $F_2 \times \mathbb{Z}$. We finish with further questions for this project.

5.1 Base group $G = \mathbb{Z}^n * \mathbb{Z}$

The automorphism group of $\mathbb{Z}^n * \mathbb{Z} \;=\; \langle x_1, \ldots, x_n \mid x_i x_j = x_j x_i \;\forall\; i, j \rangle * \langle c \rangle$, is generated by inner automorphisms, inversions, graph isomorphisms which permute the generators of the $\mathbb{Z}^n$ factor, transvections within the $\mathbb{Z}^n$ factor of the form $x_i \mapsto x_i x_j$ with all other generators fixed, and the transvections on the $\mathbb{Z}$ factor, of the form $c \mapsto c x_i$, with all other generators fixed. From this generating set, we see that Propositions 4.1.1 and 4.1.2 generalize to:

Proposition 5.1.1. *If $\Phi \in \operatorname{Aut}(\mathbb{Z}^n * \mathbb{Z})$, there is $\Psi \in \operatorname{Aut}(\mathbb{Z}^n * \mathbb{Z})$ such that $[\Phi^2] = [\Psi] \in \operatorname{Out}(\mathbb{Z}^n * \mathbb{Z})$ where Ψ restricts to $\psi : \mathbb{Z}^n \to \mathbb{Z}^n$ on the first factor and which maps $c \mapsto c x_1^{k_1} \ldots\ x_n^{k_n}$, for $k_i \in \mathbb{Z}$.*

Proposition 5.1.2. *Suppose Φ has the form of Proposition 5.1.1 above. If $[\phi]$ and $[\psi]$ are conjugate in $\operatorname{Out}(\mathbb{Z}^n)$, there is Ψ such that $[\Psi]$ and $[\Phi]$ are conjugate in $\operatorname{Out}(\mathbb{Z}^n * \mathbb{Z})$, where Ψ has the same form as Φ and the restriction of Ψ to*

$\langle x_1, \ldots, x_n \rangle$ *is* ψ.

This implies that for the purpose of calculating the Dehn function, we may assume that $\Phi \in \mathrm{Aut}(\mathbb{Z}^n * \mathbb{Z})$ restricts to an automorphism ϕ on the $\mathbb{Z}^n$ factor and maps $c \mapsto c x_1^{k_1} \ldots x_n^{k_n}$. Further, we may take ϕ to be the most convenient conjugacy class representative of $[\phi] \in \mathrm{Out}(\mathbb{Z}^n)$. So

$$M_\Phi \ = \ \langle x_1, \ldots, x_n, c, t \mid x_1^t = \phi(x_1), \ldots, \ x_n^t = \phi(x_n), \ c^t = c x_1^{\alpha_1} \ldots x_n^{\alpha_n}, \ [x_i, x_j] = 1 \rangle.$$

If ϕ has a non-unit eigenvalue, then per Lemma 2.5.3, the Dehn function of M_Φ is exponential. We therefore restrict attention to ϕ having only unit eigenvalues.

Bridson–Gersten showed in [28], that these eigenvalues are roots of unity, and thus for a sufficiently large power, ϕ^m has only the eigenvalue 1, and thus the matrix representing this map has characteristic polynomial of the form $(t-1)^n$.

The following classical result (see [80]) implies that if $\phi \in \mathrm{SL}(n, \mathbb{Z})$ has only unit eigenvalues, then ϕ is conjugate to ψ such that for some power m, $\psi^m(x_i)$ has the form $x_i x_{i+1}^{k_{i+1,i}} \cdots x_n^{k_{n,i}}$ for all $i < n$ and $\psi^m(x_m) = x_m$.

Proposition 5.1.3. *If ϕ has only unit eigenvalues, then for m large enough, the matrix representing ϕ^m is conjugate in* $\mathrm{SL}(n, \mathbb{Z})$ *to*

$$A_{\psi^m} = \begin{pmatrix} A_{11} & A_{12} & \ldots & A_{1r} \\ 0 & A_{22} & \ldots & A_{2r} \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \ldots & 0 & A_{rr} \end{pmatrix}$$

where A_{ii} is square, and the characteristic polynomial of A_{ii} is irreducible over $\mathbb{Q}$.

In particular, the characteristic polynomial of our matrix is $p_{A_{\psi^m}}(t) = (t-1)^n$ and $p_{A_{\psi^m}}(t) = p_{A_{11}}(t) \cdots p_{A_{rr}}(t)$. Since each polynomial $p_{A_{ii}}(t)$ is irreducible over

$\mathbb{Q}$, and $(t-1)^l$ is irreducible exactly when $l = 1$, then $r = n$, and A_{ψ^m} is an upper-triangular matrix.

As previously, the Dehn function of M_ϕ is a lower bound on the Dehn function of M_Φ. Take words w_k and minimal area diagrams Δ_k that witness the Dehn function of M_ϕ, $\delta_{M_\phi}(k)$. There are no c-corridors in Δ_k, since w_k contain no c's. Since c-corridors do not form annuli, this implies that the filling in the subgroup M_ϕ is best possible.
An upper bound on the Dehn function of M_Φ is $k\delta_{M_\phi}(k)$. Indeed, in a minimal area van Kampen diagram for a word w in the generators of M_Φ, c-corridors divide the diagram into at most n complementary regions, each with perimeter no worse than $(|\alpha_1| + \cdots + |\alpha_k|)|w|$. Filling the complementary regions minimally produces an upper bound of $n^2 + n\delta_{M_\phi}((|\alpha_1| + \cdots + |\alpha_k|)n)$. Together these give:

$$\delta_{M_{\phi_n}}(k) \preceq \delta_{M_{\Phi_n}}(k) \preceq k\delta_{M_{\phi_n}}(k).$$

We conjecture that in general, the Dehn function agrees with the lower bound. We would like to use an inductive strategy to show this, but so far, we can use this strategy only in the case when the super-diagonal of the associated matrix in Jordan Canonical Form has only non-zero entries, which corresponds to having a single Jordan block.

To illustrate this strategy, let $\Phi = \Phi_n$ have the form of Proposition 5.1.3, and let ϕ_n be the restriction of Φ_n to $\{x_1, \ldots, x_n\}$:

$$\Phi_n : \quad x_1 \mapsto x_1 x_2^{k_{2,1}} \cdots x_n^{k_{n,1}}, \quad \ldots \quad x_{n-1} \mapsto x_{n-1} x_n^{k_{n,n-1}}, \quad x_n \mapsto x_n, \quad c \mapsto c x_1^{\alpha_1} \ldots x_n^{\alpha_n}.$$

Let Φ_{n-1} be the automorphism of $\langle x_1, \ldots x_{n-1}, c \mid [x_i, x_j] = 1\rangle$ that we get by killing x_n, and let ϕ_{n-1} be the restriction of Φ_{n-1} to $\{x_1, \dots, x_{n-1}\}$:

$$\Phi_{n-1} : x_1 \mapsto x_1 x_2^{k_{2,1}} \cdots x_{n-1}^{k_{n-1,1}}, \ x_2 \mapsto x_2 x_3^{k_{3,2}} \cdots x_{n-1}^{k_{n-1,2}} \dots$$
$$x_{n-2} \mapsto x_{n-2} x_{n-1}^{k_{n-1,n-2}} x_{n-1} \mapsto x_{n-1}, \quad c \mapsto c x_1^{\alpha_1} \dots x_{n-1}^{\alpha_{n-1}}.$$

If the largest Jordan block for the matrix associated to ϕ_n is size $c \times c$, then the largest Jordan block for ϕ_{n-1} is either size $c \times c$ or $(c-1) \times (c-1)$. That is, $\delta_{M_{\phi_n}}(k) \approx \delta_{M_{\phi_{n-1}}}(k)$ or $\delta_{M_{\phi_n}}(k) \approx k\delta_{M_{\phi_{n-1}}}(k)$. Indeed, we are killing off x_n. Either x_n is an element of a Jordan chain that is strictly larger than all other Jordan chains, and so the longest Jordan chain of Φ_{n-1} is length $c-1$, in which case $\delta_{M_{\phi_{n-1}}}(k) \approx k^c$ and $\delta_{M_{\phi_n}}(k) \approx k^{c+1}$ or there are other Jordan chains that are still length c, so $\delta_{M_{\phi_{n-1}}}(k) \approx \delta_{M_{\phi_n}}(k) \approx k^{c+1}$.

Proposition 5.1.4. *If $[\Phi] \in \mathrm{Out}(F_n \times \mathbb{Z})$ is conjugate to $[\Phi_n]$, where Φ_n has the form above, and $k_{i+1,i} \neq 0$ for every i, that is, if there is only a single Jordan block of size $n \times n$, then $\delta_{M_\Phi}(k) \approx k^{n+1}$.*

In the proof of this proposition we make use of the following:

Proposition 5.1.5 (Bridson–Gersten [28])**.** *For every k and every $\Psi \in \mathrm{Aut}(\mathbb{Z}^n)$, there exists A and B such that if w represents the identity in $\mathbb{Z}^n \rtimes_\Psi \mathbb{Z}$, and the Dehn function of $\mathbb{Z}^n \rtimes_\Psi \mathbb{Z}$ is of degree $c+1$, there is a van Kampen diagram Δ for w, with a spanning tree, $\mathcal{T} \subset \Delta$, such that $\mathrm{Area}(\Delta) \leq A|w|^{c+1}$ and Diam$(\mathcal{T}) \leq B|w|$.*

Proof of Proposition 5.1.4. The base case was completed in Section 4.4. Let w be a word representing the identity in M_{Φ_n} and Δ be a van Kampen diagram for w with reduced c-corridors. Let $\overline{w}$ be w with all instances of x_n and x_n^{-1} removed. We build Θ, a van Kampen diagram for $\overline{w}$ with the same c-corridor pairing pattern as Δ, that is filled with respect to the presentation for $M_{\Phi_{n-1}}$. To build Θ, begin by labeling a planar loop with $\overline{w}$. We add c-corridors, following the c-pairing pattern of Δ, by replacing each c-face for the presentation of M_{Φ_n} with one for the presentation of $M_{\Phi_{n-1}}$. Since Δ had reduced c-corridors, Θ will have reduced c-corridors. By the inductive hypothesis, such a diagram has area bounded above by $A|w|^n$ for some $A > 0$. Next we charge the complementary regions with x_n 1-cells. There are at most $A|w|^n$ many new x_n-labeled 1-cells. By adding in x_n partial corridors in c-complementary regions that follow the maximal tree of that region to the boundary. If the diameter is bounded by $B|w|$, as per Proposition 5.1.5, this adds $AB|w|^{n+1}$ area to the diagram. Rearranging the word around the perimeter of each complementary region, we get that the area of Δ is less than or equal to $AB|w|^{n+1} + C|w|^3$, so we get the desired upper-bound on the Dehn function. $\square$

To build a more general inductive argument, we will need to be able to handle the case that M_{ϕ_n} and $M_{\phi_{n-1}}$ have the same Dehn function. In some cases, such as when few relators receive charges, we can still calculate the Dehn function via case-by-case arguments, but we lack a general argument.

5.1.1 Base $\mathbb{Z}^3 * \mathbb{Z}$

Making case-by-case arguments, we apply the methods of Sections 4.4 and 4.5 to classify the Dehn functions of mapping tori with base group $\mathbb{Z}^3 * \mathbb{Z}$. If the base group is $\mathbb{Z}^3 * \mathbb{Z}$, Proposition 5.1.3 implies that it is sufficient to consider automorphisms of the form:

$$\Phi_3: \; x_1 \mapsto x_1 x_2^a x_3^b, \quad x_2 \mapsto x_2 x_3^d, \quad x_3 \mapsto x_3 \quad c \mapsto c x_1^{\alpha_1} x_2^{\alpha_2} x_3^{\alpha_3}$$

$$\Phi_2: \; x_1 \mapsto x_1 x_2^a, \quad x_2 \mapsto x_2, \quad c \mapsto c x_1^{\alpha_1} x_2^{\alpha_2}$$

The methods of the previous section imply that the Dehn functions of M_{Φ_2} are equivalent to the area of the diagrams Θ.

Case 1: $a \neq 0, \; d \neq 0$, then M_{Φ_3} has quartic Dehn function, as per Proposition 5.1.4.

Case 2: $a \neq 0, \; d = 0$, then M_{Φ_3} has cubic Dehn function. The largest Jordan block is a 2×2 block. Therefore $n^3 \preceq \delta_{M_{\Phi_3}}(n) \preceq n^4$. The diagram Θ in the generators of M_{Φ_2} has a cubic upper bound on area — but that's not a guarantee that all 2-cells in Θ are charged. The only 2-cells that are charged with x_3's are the c-faces and the faces corresponding to the relator $t^{-1}x_1t = x_1x_2^a$, the intersection of a t-corridor and an x_1-corridor. In a reduced diagram, an x_1- and t-corridor intersect at most once, and so the total number of such cells is no more than n^2, the number of x_1-corridors multiplied by the number of t-corridors. Similarly, for c-corridors there is a quadratic upper bound on the number of c-faces. Since the quadratically many x_3-charges only have linear distance to go, although the area of Θ over M_{Φ_2} is at most cubic, the area of Δ over M_{Φ_3} is also at most cubic.

Case 3: $a = 0, \; d \neq 0, \; b \neq 0$, then M_{Φ_3} has cubic Dehn function. The largest Jordan block has size 2×2, so M_{Φ_3} has a cubic lower bound on the Dehn function. The group M_{Φ_2} has quadratic Dehn function, which produces a cubic upper bound on the Dehn function of M_{Φ_3}.

Case 4: $a = 0, \; d = 0, \; b \neq 0$, then M_{Φ_3} has cubic Dehn function. Again, the largest Jordan block has size 2×2, so there is a cubic lower bound on the Dehn function. The diagram Θ in the generators of the group M_{Φ_2} has quadratic area. Charging the diagram with x_3 cells produces a diagram with at most a cubic upper bound on the area.

Case 5: $a = 0,\ d = 0,\ b = 0$, then M_Φ has quadratic Dehn function, by the same argument as in the Proof of Theorem 5(1).

5.2 Base group $G = F_k \times F_l$

In this section we consider the case that the graph is a join of l and k vertices. We are able to classify the Dehn functions of mapping tori when $k, l > 1$, and we consider some of the challenges for the case $F_k \times \mathbb{Z}$. We begin by examining a generating set for $\mathrm{Aut}(F_k \times F_l)$.

Lemma 5.2.1. *For a bipartite graph with vertex sets X and Y, with $|X| = k$ and $|Y| = l$, consider the following generating set of $\mathrm{Aut}(F_k \times F_l)$:*

1. *Inner automorphisms*
2. *Inversions*
3. *Partial conjugations: If $x \in X$ then $\Gamma - \mathrm{star}(x) = X - \{x\}$, and therefore the automorphism restricts to an automorphism of $F(X)$.*
4. *Transvections: Suppose $y \in Y$, so $\mathrm{star}(y) = \{y\} \cup X$.*
   - *(a) If $w \in Y$, then $\mathrm{link}(w) = X$, and so $\mathrm{link}(w) \subset \mathrm{star}(y)$.*
   - *(b) If $w \in X$, then since $\mathrm{link}(w) = Y$, $\mathrm{link}(w) \subset \mathrm{star}(y)$ exactly when $Y = \{y\}$.*

Thus when $k, l > 1$, there are no transvections between $F_k = F(X)$ and $F_l = F(Y)$. When $l = 1$, then $Y = \{y\}$ and the element y is central, so there are transvections of the form $w \mapsto wy$ for every $w \in X$.

5. *Graph Symmetries: If $k \neq l$, then all graph symmetries are permutations of the vertices within X and within Y. In the case that $k = l$, all symmetries can be expressed as swapping $F(X)$ and $F(Y)$ and composition with permutations within X or Y. These are central, finite order automorphisms.*

5.2.1 Base group $G = F_k \times F_l$ for $k, l > 1$

Above we saw that when $k, l > 1$, transvections preserve the subgroups F_k and F_l. All other automorphisms except possibly the graph symmetries and inner automorphisms also preserve the subgroups $F(X)$ and $F(Y)$. In the case that $k, l > 1$, then for p large enough, up to an inner automorphism, Φ^p is $\phi_1 \times \phi_2$, for $\phi_1 \in \mathrm{Aut}(F_k)$ and $\phi_2 \in \mathrm{Aut}(F_l)$. By Lemma 2.5.1, for the purpose of calculating Dehn functions, we can consider $\Phi = \phi_1 \times \phi_2$. Both M_{ϕ_1} and M_{ϕ_2} are isomorphic to subgroups of M_Φ.

Definition 5.2.2. Let $F_k = \langle x_1, \dots, x_k \rangle$. For ϕ a free group automorphism, the *growth function of* ϕ is $g(n) := \max_i \{|\phi^n(x_i)|, |\phi^{-n}(x_i)|\}$. Let $f(n) := \max_{k \leq n} g(k)$. (I know of no name for f.)

We say that functions $g_1 \asymp g_2$ if $g_1 \preceq g_2$ and $g_2 \preceq g_1$, where $g_1 \preceq g_2$ if there is $C > 0$ such that for all n, $g_1(n) \leq C g_2(Cn)$.

Lemma 5.2.3. *Changing the presentation for F_k produces equivalent g and f.*

Proof. Suppose F_k admits the presentations

$$\mathcal{P}_1 = \langle x_1, \ldots, x_l \mid R\rangle \quad \text{and} \quad \mathcal{P}_2 = \langle y_1, \ldots, y_m \mid S\rangle.$$

In $\mathcal{P}_1$, let w_i be a shortest word representing the same element as y_i. Define $C = \max_i\{|w_i|_{\mathcal{P}_1}\}$. Thus for all r, $B_{\mathcal{P}_2}(1, r) \subset B_{\mathcal{P}_1}(1, Cr)$. Similarly, there is a shortest word v_j in $\mathcal{P}_2$ representing x_j. Define $D = \max_j\{|v_j|_{\mathcal{P}_2}\}$. For all r $B_{\mathcal{P}_1}(1, r) \subset B_{\mathcal{P}_2}(1, Dr)$. If $\phi \in \mathrm{Aut}(F_k)$, then for each $\mathcal{P}_i$, there is a growth function g_i for ϕ. The g_i tell us that

$$\phi^{\pm n}(B_{\mathcal{P}_2}(1, r)) \subset \phi^{\pm n}(B_{\mathcal{P}_1}(1, Cr)) \subset B_{\mathcal{P}_1}(1, Crg_1(n)) \subset B_{\mathcal{P}_2}(1, CDrg_1(n)).$$

Thus for all n, $g_2(n) \leq CDg_1(n)$, and so $g_2 \preceq g_1$. Further,

$$f_2(n) = \max_{k\leq n}\{g_2(n)\} = g_2(j) \leq CDg_1(j) \leq CD \max_{k\leq n}\{g_1(k)\} = CDf_1(n),$$

so $f_2 \preceq f_1$. Switching $\mathcal{P}_1$ and $\mathcal{P}_2$, we see that $g_1 \asymp g_2$ and $f_1 \asymp f_2$. $\square$

Proposition 5.2.4. *If $\Psi \in \mathrm{Aut}(F_k \times F_l)$, then $[\Psi^2] = [\Phi]$ in $\mathrm{Out}(F_k \times F_l)$ such that $\Phi = \phi_1 \times \phi_2$ where $\phi_1 \in \mathrm{Aut}(F_k)$ and $\phi_2 \in \mathrm{Aut}(F_l)$.*

The Dehn function of M_Ψ is equivalent to $n^2 \min\{f_1(n), f_2(n)\}$, where f_i is the growth function for ϕ_i. When ϕ_1 and ϕ_2 both have exponential growth, this implies that M_Ψ has exponential Dehn function.

Proof. We will show this in the case that ϕ_i has polynomial growth. A similar argument establishes the lower bound (and thus the result) in the case of exponential growth. If ϕ is a positive automorphism, it is clear that $g_1 = f_1$. In [14], Bestvina, Feighn, and Handel developed the technology of relative train track maps. A consequence of their theory is that if g is dominated by a polynomial, it is actually equivalent to a polynomial. Thus we may use $g(n)$ to calculate lower bounds and $f(n)$ to calculate upper bounds, and they are equivalent. We will show that

$$n^2 \min\{g_1(n), g_2(n)\} \preceq \delta(n) \preceq n^2 \min\{f_1(n), f_2(n)\}.$$

The proof of the lower bound is a minor adaptation of the proof of Bridson and Gersten [28] of the Dehn function of $\mathbb{Z}^n \rtimes_\phi \mathbb{Z}$. For the purpose of our proof, let us assume that $f_1 \preceq f_2$. We may assume that $k \ll n$. This implies that in the set of words $\{\phi_1^r(a_j) \mid 1 \leq j \leq k, n \leq r \leq 2n\}$, there is some i such that for at least $\frac{1}{k}$ of the values of r, $|\phi_1^r(a_i)| = \max_j\{|\phi_1^r(a_j)|\}$ attains the maximum. Let y be the element a_i that achieves the maximum most frequently, and let
$\mathcal{A} = \{r \mid g(r) = |\phi_1^r(y)| \text{ with } n \leq r \leq 2n\}$, be the subset of $\{n, \ldots, 2n\}$ on which y witnesses the value of g. Let x most frequently achieve the maximum of $\phi^{-r}(b_i)$, across generators of F_l from $n \leq r \leq 2n$.

Consider the word $w_n = t^{-4n}y^n t^{4n} x^n t^{-4n} y^{-n} t^{4n} x^{-n}$, illustrated in Figure 5.1a:

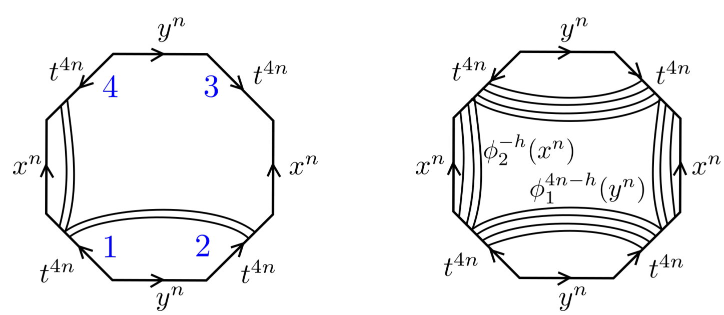

(a) Few choices for t-corridor patterns

(b) A single t-corridor pairing determines all of them.

Figure 5.1: Fillings in the mapping torus with base $F_k \times F_l$ and automorphism $\phi_1 \times \phi_2$

A t-corridor beginning on side 1 can only end on sides 2 and 4. Moreover, since t-corridors cannot cross, there is some value h at which they switch from ending on side 2 to ending on side 4. This switching point determines the diagram, as seen in Figure 5.1b. If $h \neq 0$, there is a rectangular region to fill, as pictured, with word around the outside of $[\phi_1^{4n-h}(y^{-n}),\ \phi_2^h(x^{-n})]$. No matter what the value of h, there is a guaranteed belt of at least $2n$ t-corridors that start and end in the same segment. We do the calculation in the case that sides 1 and 2 have this pairing, since $f_1 \preceq f_2$. If C_r is the rth such corridor, then

$$Area(\Delta) \geq \sum_{r=n}^{r=2n} |C_r| \geq \sum_{r=n}^{2n} |\phi_1^r(y^n)| \geq \sum_{r\in\mathcal{A}} g_1(r)n.$$

Since g is equivalent to an increasing function, this gives us that $\sum_{r\in\mathcal{A}} f_1(r)n \geq \frac{n^2}{k} f_1(n)$, and therefore $\delta_{M_\Phi}(20n) \succeq \frac{n^2}{k} f_1(n)$.

To get an upper bound on the Dehn function, we look at how many relators need to be applied to a word w that represents the identity in $M_{\phi_1\times\phi_2}$, in order to reduce it to the trivial word. We are also allowed to add inverse pairs of generators $x_i^{-1}x_i$, and to cancel away inverse pairs of generators, but only the application of relators will count towards the area. The strategy will be as follows: recall that $\phi_1 : F_k \to F_k$ and $\phi_2 : F_l \to F_l$ and $\phi_1 \preceq \phi_2$. If $F_k = \langle a_1, \ldots, a_k\rangle$ and $F_l = \langle b_1, \ldots, b_l\rangle$, then words have the form

$$w = w_1 t^{c_1} w_2 t^{c_2} \cdots w_n t^{c_n}$$

for $w_i \in \langle a_1, \ldots, a_k, b_1, \ldots b_l\rangle$. By applying fewer than $|w|^2$ relators, w can be rewritten to

$$w' = u_1 v_1 t^{c_1} \ldots u_n v_n t^{c_n}$$

for $u_i \in \langle a_1, \ldots, a_k\rangle$ and $v_i \in \langle b_1, \ldots, b_l\rangle$. Apply relators to shuffle all generators of F_k to the right. For instance, u_n must shuffle past v_n and past t^{c_n}, and u_k and its iterates under ϕ_1 must shuffle past v_i and t_i^c for all $i \geq k$. After shuffling these generators to the right, what remains is a word w'' representing the identity in $F_l \rtimes_{\phi_2} \mathbb{Z}$. By [29], reducing w'' will take no more than $|w''|^2$ relations.

We estimate shuffling by (over)estimating the length that the u_i grow to, and letting them pass the v_i and t^{c_i}. The number of relators required for elements of $\langle a_1, \ldots, a_k \rangle$ to cross v_i and t^{c_i} is bounded above by:

$$(|v_i| + |c_i|) \left(\sum_{j=1}^{i-1} f_1(|c_j| + \cdots + |c_{i-1}|)|u_j| \right) \leq (|v_i| + |c_i|) \left(\sum_{j=1}^{i-1} f_1(|w|)|u_j| \right).$$

In total then, we require no more than

$$\sum_{i=1}^{n} (|v_i| + |c_i|) f_1(|w|) \left(\sum_{j=1}^{i-1} |u_j| \right) \leq \sum_{i=1}^{n} (|v_i| + |c_i|) f_1(|w|)|w| \leq |w|^2 f_1(|w|)$$

relations to rewrite w as a word in M_{ϕ_2}. Since M_{ϕ_2} has only quadratic Dehn function, this implies that $\delta_{M_\Phi}(n) \approx n^2 \min\{f_1(n), f_2(n)\}$. □

5.2.2 Strategies and challenges for $F_k \times \mathbb{Z}$

We have not yet computed the Dehn functions for mapping tori which have base group $F_k \times \mathbb{Z}$, which would generalize the work of Chapter 3. Recall from Lemma 5.2.1 that there are additional transvections in the generating set for $\mathrm{Aut}(F_k \times \mathbb{Z}) = \mathrm{Aut}(\langle x_1, \ldots, x_k \rangle \times \langle c \rangle)$. Indeed, since $\mathrm{link}(x_i) = \{c\} \subset \mathrm{star}(c)$, for every i, there is a transvection $x_i \mapsto x_i c$. It is clear from the generating set that for any $\Psi \in \mathrm{Aut}(F_k \times \mathbb{Z})$, there is $\Phi \in \mathrm{Aut}(F_k \times \mathbb{Z})$ such that $[\Psi^2] = [\Phi]$ in $\mathrm{Out}(F_k \times \mathbb{Z})$ and

$$\Phi : x_i \mapsto \phi(x_i) c^{k_i}, \quad c \mapsto c.$$

For the purpose of calculating Dehn functions, we may assume our automorphism has the form of Φ. We begin with a trivial observation.

Proposition 5.2.5. *If $\phi \in \mathrm{Aut}(F_k)$ is atoroidal, then M_Φ has quadratic Dehn function.*

Proof. We can directly apply Gersten and Riley's electrostatic model. M_Φ is a central extension of M_ϕ, which has a linear Dehn function, and admits simultaneous linear upper bounds on Diameter and Area. Theorem 2.4.1 implies that M_Φ has quadratic area. $\square$

Proposition 5.2.6. *If there is $w \in F_k$ such that $\Phi(w) = wc^k$, then M_Φ has cubic Dehn function.*

Unlike in the case of $k = 2$, homological arguments are not sufficient for Proposition 5.2.6— it is not enough to know that w abelianizes to a fixed point of ϕ_{ab}. As when $k = 2$, the remaining cases are those in which ϕ is not atoroidal, and fixed points of ϕ are fixed points of Φ.

As usual, since M_Φ is a central extension of M_ϕ, we begin by considering M_ϕ:

Theorem 6 (Gautero–Lustig [51]). *For every automorphism ϕ, there is a family of subgroups $\mathcal{H} = (H_1, \ldots, H_r)$ with $H_i \leq F_n$ and $\mathcal{H}_\phi = (H_1^\phi, \ldots, H_r^\phi)$ where $H_i^\phi := \langle H_i, t^{m_i} g_i^{-1} \rangle \leq M_\phi$, such that $H_i^{t^{m_i}} = \phi^{m_i}(H_i) = H_i^{g_i}$. M_ϕ is hyperbolic relative to the family $\mathcal{H}_\phi$.*

Gautero and Lustig define the *growth of ϕ on $w \in F_k$* to be polynomial if there are $C > 0$ and $d \geq 0$ such that $||\phi^l(w)|| \leq Cl^d$ for all $l \in \mathbb{Z}$, where $|| \cdot ||$ denotes the cyclically reduced length. The growth of ϕ on a subgroup is said to be polynomial if it is polynomial for each element of the subgroup. This definition is designed to study automorphisms of exponential growth— ϕ has polynomial growth if and only if the growth of ϕ on every generator is polynomial. Gautero–Lustig show that the subgroups H_i can be chosen so that ϕ has polynomial growth on H_i for each i, and the subset $\cup_{i=1}^r \cup_{g \in M_\phi} H_i^g$ is exactly the set of elements with polynomial growth. Moreover, they can choose the family $\mathcal{H}$

to be closed (up to conjugation) under ϕ: if $h \in H_i$, then $\phi(h) \in H_j^g$ for some $j \in \{1, \ldots, r\}$ and $g \in M_\phi$.

To complete the classification of Dehn functions for mapping tori with base group $F_l \times \mathbb{Z}$, we first need to classify the Dehn functions of mapping tori built from polynomially growing automorphisms. This might be possible to show directly, perhaps by following Macura [73]. We then hope to use the work of Gautero–Lustig to make arguments similar to those of Section 3.2. That is, we find diagrams of minimal type with respect to $\mathcal{H}_\phi$, we charge the diagram, and apply the electrostatic model. The regions of the diagram that are complementary to the peripheral subgroup will discharge just as in Section 3.2. One of the challenges of this strategy, which might make it untenable, is that the peripheral subgroups may receive charges. Although $w = \overline{w}c^k$ satisfies that $|k| < |w|$, the total number of charges added to the diagram may be quadratic in $|w|$. It is not clear when, if ever, this forces the Dehn function to be cubic.

5.2.3 Applications

We would like to complete this study in order to get lower bounds for the Dehn functions of mapping tori when the base groups are 2-dimensional irreducible RAAGs, that is, for RAAGs built from graphs which are connected and contain no triplets of vertices for which the induced subgraph is a triangle. Indeed, in Proposition 3.2 of [35], Charney, Crisp, and Vogtmann show that if Γ is a graph which is connected and triangle-free, with at least two edges, then joins in Γ play an important role in determining the automorphisms of $A(\Gamma)$.

The subgroup $\mathrm{Aut}^0(A(\Gamma))$ generated just by inner automorphisms, inversions, partial conjugations, and transvections is called the *pure automorphism group* and is finite index in $\mathrm{Aut}(A(\Gamma))$. Thus if $\Phi \in \mathrm{Aut}(A(\Gamma))$, there is a $k \in \mathbb{N}$ such that $\Phi^k \in \mathrm{Aut}^0(A(\Gamma))$, and M_{Φ^k} and M_Φ have the same Dehn function. For the purpose of calculating Dehn functions, it suffices to consider compositions of inversions, partial conjugations, and transvections.

Proposition 5.2.7 (Charney, Crisp, Vogtmann [35])**.** *Suppose that $J = U * W \subset \Gamma$ is a maximal join in Γ, a connected, triangle-free graph and $\Psi \in \mathrm{Aut}^{(0)}(A(\Gamma))$. Then Ψ maps $A(J)$ to a conjugate of $A(J)$.*

Proposition 5.2.7 implies there is always Ψ' so that $A(J) = F_k \times F_l$ maps to $A(J)$. Further, if Γ is connected and triangle-free, then Γ contains many joins — every vertex v in Γ is part of a join. Indeed, let $v \in \Gamma$. Then $\mathrm{link}(v) * \{v\}$ is a join in $A(\Gamma)$. There are no edges between two elements of $\mathrm{link}(v)$, as this would form a triangle. J_v is defined to be the maximal join containing $\mathrm{link}(v) * \{v\}$. By completing the classification of Dehn functions for mapping tori with base $F_k \times \mathbb{Z}$, we will get lower bounds for the Dehn functions of mapping tori of 2-dimensional irreducible RAAGs.

Proposition 5.2.8. *If $\Psi \in \mathrm{Aut}(A(\Gamma))$ for Γ a connected, triangle-free graph, then for some n, $\Psi^n \in \mathrm{Aut}^{(0)}(A(\Gamma))$. For every maximal join $J = U * W$ in Γ, with $|U|, |W| \geq 2$, Ψ^n is conjugate to Φ_J such that Φ_J restricts to an automorphism $\phi_J = \phi_1 \times \phi_2$ on $F(U) \times F(W)$. Then $\max\{\delta_{M_{\phi_J}} \mid J \text{is a maximal join, but not a star in } A(\Gamma)\} \preceq \delta_{M_\Psi}$.*

Proof. We can take words of the same form as previously:

$$w_n = t^{-4n} y^n t^{4n} x^n t^{-4n} y^{-n} t^{4n} x^{-n}$$

where y is a generator of $A(U)$ and x is a generator of $A(W)$ such that $|\phi_1^n(y)|$ is maximized and $|\phi_2^{-n}(x)|$ is maximized in the same way as previously. Since the alphabet of M_{ϕ_J} is contained in the alphabet of M_{Φ_J}, this gives us a word of length $20n$. We have a filling for it using a subset of the relators of M_{Φ_J}. In order to get our lower bound, we need that there is no better filling using the full set of relators that define M_{Φ_J}. Again, we rely on the t-corridor pattern to provide a lower bound: there is some turning point at which t-corridors starting on segment 1 end on segment 4 instead of segment 2. Now, we wonder if we have the same t-corridors as before, or new, shorter ones? Note that since M_{Φ_J} is an automorphism on $F(U) \times F(W)$ and $A(\Gamma)$, there are no elements in $A(\Gamma) - F(U) \times F(W)$ that are mapped by Φ_J into $F(U) \times F(W)$. This implies that the words along both the tops and the bottoms of our t-corridors must be words in $F(U) \times F(W)$, and so indeed, we have the same t-corridors that gave us our area lower bound. □

5.3 A general estimate

Finally, we improve our lower bounds on the Dehn functions of mapping tori of RAAGs by applying a result of Groves and Hull to find abelian subgroups for which we can apply the estimates of Bridson–Gersten. Matt Clay gave conditions for $A(\Gamma)$ to split over $\mathbb{Z}$, and Groves and Hull completed this to include all splittings over abelian groups. Our interest will be in splittings over non-trivial abelian subgroups.

Proposition 5.3.1 (Clay, Groves–Hull [57])**.** *The group $A(\Gamma)$ splits non-trivially over an abelian (possibly trivial) group if and only if*

1. *Γ is not connected.*

2. *Γ is a complete graph.*

3. *Γ contains a separating clique.*

Definition 5.3.2. $K \subset \Gamma$ is a *separating clique* if K is a complete graph, and $\Gamma - K$ has more than one component. Call K an *invariant separating clique* if K is of maximal size among the separating cliques of Γ, and if for all $v \in K$ and $w \notin K$, $\text{link}(v) \not\subset \text{star}(w)$.

The condition we wish to avoid, $\text{link}(v) \subset \text{star}(w)$ for some $v \in K$ and $w \notin K$, holds only if the induced graph on w and the vertices of $K - \{v\}$ is also a clique on the same number of vertices.

Example 5.3.3. The subgraph K of Γ is a separating clique. It is not an invariant separating clique since $\text{link}(v) \subset \text{star}(w)$. The subgraph K' of Γ' is an invariant separating clique.

Figure 5.2: K is a separating clique. K' is an invariant separating clique.

We can immediately apply Servatius and Laurence's generating set to give a lower bound on the Dehn functions of mapping tori for graphs Γ with invariant separating cliques.

Proposition 5.3.4. *Suppose that Γ has an invariant separating clique K. If $\Phi \in \text{Aut}(A(\Gamma))$, then for some power N, Φ^N is conjugate to a map Ψ such that Ψ restricts to an automorphism ψ on K.*

Proof. For any $\Phi \in \mathrm{Aut}(A(\Gamma))$, there is some power N such that $\Phi^N = \Psi' \in \mathrm{Aut}^0(A(\Gamma))$, since $\mathrm{Aut}^0(A(\Gamma))$ has finite-index in $\mathrm{Aut}(A(\Gamma))$. Thus we restrict our attention to automorphisms Ψ' in $\mathrm{Aut}^0(A(\Gamma))$. We examine what inversions, partial conjugations, and transvections do to the subgroup $A(K)$. Inversions induce an automorphism of K, which is either the identity on K or an inversion on K. Partial conjugations have the same effect on K as inner automorphisms. Indeed,

1. If $v \in K$, then $K \subset \mathrm{star}(v)$, so the partial conjugation restricts to the identity on K.
2. If $w \notin K$, then $K - \mathrm{star}(w)$ will never have more than one component, and thus either the partial conjugation restricts to the identity on K, or to conjugation by w.

By assumption, all transvections restrict to automorphisms of K. Thus Ψ' sends K to a conjugate $g^{-1}Kg$. The automorphism $\Psi = \iota_{g^{-1}} \circ \Psi'$ will restrict to an automorphism on K. $\square$

Corollary 5.3.5. *If K is an invariant separating clique, and Φ and Ψ are as above, the Dehn function of $A(\Gamma) \rtimes_\Phi \mathbb{Z}$ is bounded below by the Dehn function of $A(K) \rtimes_\psi \mathbb{Z}$:*

$$\delta_{A(K)\rtimes_\psi \mathbb{Z}}(k) \preceq \delta_{A(\Gamma)\rtimes_\Phi \mathbb{Z}}(k).$$

Proof. Since the subgroup $A(K)$ is q.i. embedded in $A(\Gamma)$, Lemma 2.5.3 implies that the Dehn function of $A(K) \rtimes_\psi \mathbb{Z}$ provides a lower bound for the Dehn function of M_Ψ. We get the result since M_Φ and M_Ψ have equivalent Dehn functions. $\square$

5.4 Further questions

Although we have computed a few examples of the Dehn functions of mapping tori, we do not have a general strategy for these computations. Questions for the future include:

1. What automorphisms of $A(\Gamma)$ suffice to classify the Dehn functions of all mapping tori of $A(\Gamma)$?
2. How far do the methods of Section 4 extend? Can they be used inductively to find upper bounds for Dehn functions, as in Corollary 5.1.4?
3. What are the Dehn functions of mapping tori with base group $F_n \times \mathbb{Z}$?
4. What are the Dehn functions of mapping tori for Γ when Γ is connected and triangle-free?

Part II

Residual Properties of Groups

CHAPTER 1

PRELIMINARIES

In this first chapter we introduce the basic theory of residually $\mathcal{P}$ groups. Some of this theory is necessary background for Chapters 2 and 3. Other material is included to provide context for these projects.

1.1 Residually $\mathcal{P}$-groups

The following three definitions are closely related and are often studied jointly:

Definition 1.1.1. Let $\mathcal{P}$ be a property of finitely presented groups. For example, finiteness, freeness, nilpotence, solvability, or amenability. A group G is *residually* $\mathcal{P}$ if for every $g \in G - \{1\}$, there is a quotient $\phi : G \to Q$ such that $\phi(g) \neq 1$, where Q is a group with property $\mathcal{P}$.

Definition 1.1.2. A group G is *fully residually* $\mathcal{P}$ if for any finite subset $F \subset G$, there is a quotient $\psi : G \to Q$ such that Q has property $\mathcal{P}$ and $\psi \restriction F$ is one-to-one.

In particular, G being fully residually $\mathcal{P}$ implies that arbitrarily large balls in the Cayley graph G can be embedded into a family of quotients with property $\mathcal{P}$. Such a G is sometimes referred to as being approximable by $\mathcal{P}$-groups.

Definition 1.1.3. A subgroup $H \leq G$ is $\mathcal{P}$-*separable* if for every $g \in G - H$, a group with property $\mathcal{P}$ witnesses that g is not an element of H. That is, there is $\phi : G \to Q$ where Q has property $\mathcal{P}$ and $\phi(g) \notin \phi(H)$.

$\mathcal{P}$ separability is an important necessary condition for constructing residually $\mathcal{P}$ groups by HNN extensions and amalgated free products. In Chapter 2 we will show that an amalgamated product is not residually finite by showing that the amalgamating subgroup is not finitely separable.

The following observations and lemmas show some of the relationships between these three properties. Fully residually $\mathcal{P}$ groups are always residually $\mathcal{P}$. Residually finite groups are always fully residually finite, ditto for residually solvable groups.

Lemma 1.1.4. *If $\mathcal{P}$ is a property that is preserved under taking finitely many direct sums, then being being residually $\mathcal{P}$ is equivalent to being fully residually $\mathcal{P}$.*

Proof. For $F = \{g_1, \ldots, g_n\} \subset G$, consider the set $F' = \{g_i g_j^{-1} \mid i < j\}$. Note that F' does not include the identity element, and so for every $f \in F'$, there is a map to a quotient with property $\mathcal{P}$ distinguishing f from the identity: $\psi_f : G \to Q_f$. Define $\Psi : G \to \hat{Q} := \oplus_{f \in F'} Q_f$, via $g \mapsto (\psi_f(g))_{f \in F'}$. By assumption, $\hat{Q}$ has property $\mathcal{P}$. For each element $f \in F'$, $\Psi(f)$ is non-trivial in at least one coordinate, so $\Psi(f) \neq 1$. Therefore Ψ is injective on F, as desired, and so G is fully residually $\mathcal{P}$. $\square$

In contrast, freeness is not preserved by direct sums. Groups which are residually free are not in general fully residually free. For example $F_2 \times F_2$ is residually free but not fully residually free.

Now we come to the relationship between residual $\mathcal{P}$-ness and $\mathcal{P}$-separability. The most basic observation is that residual $\mathcal{P}$-ness of a group G is equivalent to $\mathcal{P}$-separability of the subgroup $\{1\}$ in G. This is a special case of the next lemma:

Lemma 1.1.5. *If $H \trianglelefteq G$ and G is residually $\mathcal{P}$, then H is $\mathcal{P}$-separable in G if G/H is residually $\mathcal{P}$. Moreover, if all quotients of a $\mathcal{P}$-group have property $\mathcal{P}$ (e.g. $\mathcal{P}$ is finiteness or solvability), the converse is true as well.*

Proof. Let $H \trianglelefteq G$, and let p be the natural map $p : G \to G/H$. When $g \notin H$, then $p(g) \neq 1$. If G/H is residually $\mathcal{P}$, there is a map $q : G/H \to Q$ such that Q has property $\mathcal{P}$ and $q(p(g)) \neq 1$. Composing, the map $q \circ p : G \to Q$ is a map from G to a $\mathcal{P}$ group such that the image of g is not in the image of H. Therefore H is $\mathcal{P}$-separable in G.

Consider the case when $\mathcal{P}$ is a property inherited by quotients and suppose that $H \triangleleft G$ is $\mathcal{P}$-separable. We show that G/H is residually $\mathcal{P}$. Let $\pi : G \to G/H$. For every $\bar{g} \notin H$, there is some $g \in G$ such that $\pi(g) = \bar{g}$. By the $\mathcal{P}$ separability of H in G, there is $N \triangleleft G$ and $p : G \to G/N$ such that $p(g) \notin p(H)$ and G/N has property $\mathcal{P}$. We can assume that p is the canonical quotient map. Since $p(g) \notin p(H)$, g is not in $p^{-1}(p(H)) = HN$. Since $HN \triangleleft G$, we can take $r : G \to G/HN$, and $r(g) \neq 1$. The map r factors as $f \circ \pi$ where $\pi : G \to G/H$ and $f : G/H \to G/HN$, and thus $f(\pi(g)) = f(\bar{g}) \neq 1$. Since G/HN is also a quotient of the $\mathcal{P}$ group G/N, G/HN has property $\mathcal{P}$. Therefore G/H is residually $\mathcal{P}$. $\square$

It is easy to see that if $H \leq G$ and G is residually $\mathcal{P}$, that H is residually $\mathcal{P}$.

Lemma 1.1.6. *If A and B are residually $\mathcal{P}$, then $A \times B$ is residually $\mathcal{P}$.*

Proof. Every element $g \neq 1$ in $A \times B - \{1\}$ can be expressed as $g = a_g b_g$ where one of these factors is non-trivial, say a_g. By the residual $\mathcal{P}$-ness of A, there is a map $q : A \to Q$ where $q(a_g) \neq 1$ and Q is $\mathcal{P}$. Composing q with the natural map from $A \times B \to A$, we get a map from $A \times B$ to a $\mathcal{P}$ group such that g has non-trivial image in Q. $\square$

1.2 Residual finiteness

Residual finiteness and residual separability have several other equivalent formulations. In the original formulation of residual finiteness, the statement is about finite quotients. This can be restated as:

Definition 1.2.1. A group G is *residually finite* if

$$\bigcap_{\substack{[G:K]<\infty \\ K \trianglelefteq G}} K = \{1\}.$$

In the next lemma, we will see that for every $H < G$ of finite index, there is a $K < H < G$ such that $K \trianglelefteq G$ and $[G : K] < \infty$.

Lemma 1.2.2. *If $H \leq G$ is finite index, then the subgroup*

$$K = \mathrm{core}(H) = \bigcap_{g \in G} g^{-1} H g,$$

called the core of H, is normal and finite index in G.

Proof. Let $X = \{gH\}$ be the space of left cosets of H. The group G acts on X on the left by $g_1 \cdot g_2 H = (g_1 g_2)H$. An element aH of X is fixed by g exactly when $g \in aHa^{-1}$: $g \cdot aH = aH$ implies that $gah_1 = ah_2$ for some $h_1, h_2 \in H$, so $g \in aHa^{-1}$. Thus K, the kernel of this action, is the set of elements that are in gHg^{-1} for all $g \in G$, that is, $K = \mathrm{core}(H)$. Since G acts by permutations on the $[G : H]$ element set X, this implies that $G/K < \mathrm{Sym}([G : H])$, and so $[G : K] < [G : H]! < \infty$. $\square$

Lemma 1.2.2 implies that if h is in every finite index normal subgroup, it is also in every finite index subgroup; if H is a finite index subgroup, it contains the finite index normal subgroup core(H). Since h is in core(H), h is in H. Therefore

$$\bigcap_{\substack{[G:K]<\infty \\ K \trianglelefteq G}} K = \bigcap_{[G:H]<\infty} H.$$

This gives us the following convenient equivalent definitions:

Definition 1.2.3. A group G is *residually finite* if

$$\bigcap_{[G:H]<\infty} H = \{1\}.$$

Definition 1.2.4. A subgroup $H \leq G$ is *finitely separable* (separable), if the intersection of all finite index subgroups of G containing H is exactly H.

The next lemma is straightforward but important:

Lemma 1.2.5. *Residual finiteness extends from finite index subgroups up to their supergroups; if $G \leq \Gamma$ and $[\Gamma : G] < \infty$, then G being residually finite implies that Γ is residually finite.*

Proof. Suppose that G is a residually finite, finite index subgroup of Γ. Let $\mathcal{A}$ and $\mathcal{B}$ be the set of finite index subgroups of G and Γ, respectively. Moreover, $\mathcal{A} \subset \mathcal{B}$, since $[\Gamma : H] = [\Gamma : G][G : H]$ when $H \leq G < \Gamma$. Therefore

$$\bigcap_{\widehat{H}\in\mathcal{B}} \widehat{H} = G \cap \left(\bigcap_{H\in\mathcal{B}} \widehat{H} \right) = \left(\bigcap_{H\in\mathcal{B}} \widehat{H} \cap G \right) \subset \bigcap_{H\in\mathcal{A}} H = \{1\},$$

and so Γ is residually finite. $\square$

One of the pay-offs of being residually finite is having a solution to the Word Problem. This also provides an obstruction to residual finiteness. It tells us for example, that there are solvable groups which are not residually finite.

Theorem 7 (McKinsey [77]). *If $G = \langle X \mid R \rangle$ is finitely presented and residually finite, then G has a solvable Word Problem.*

Proof. Given a finite presentation $\langle X \mid R \rangle$ for G, a residually finite group, and a word $w \in F(X)$, we can decide whether or not w represents the identity in G. Run in parallel two processes:

1. List all of the words representing the identity in G. If w represents the identity in G, it will eventually appear in this list. Halt if this happens.
2. List all multiplication tables for finite groups, in order of the cardinality of these groups. List all homomorphisms ϕ from G to these groups. Rewrite $w(X)$ as $w(\phi(X))$ and check if $w(\phi(X))$ represents the identity in these finite groups. Since G is residually finite, if the word w does not represent the identity there will eventually be a finite group and a homomorphism in which w has non-trivial image. Halt if this happens.

□

1.2.1 Separability and subgroup separability

A relatively rare but powerful property of groups is subgroup separability, also called locally extended residual finiteness (LERF). This amounts to all finitely generated subgroups being finitely separable.

Definition 1.2.6. A group G is LERF if every finitely generated subgroup is a separable subgroup.

Just as McKinsey's Algorithm implies that groups which are residually finite have a solution to the word problem, being LERF implies that for all finitely generated $H \leq G$, there is a solution to the membership problem in H. The group $F_2 \times F_2$ is thus not LERF: it has finitely generated subgroups for which there is no solution to the membership problem. Indeed, if $Q = \langle X | R \rangle$ is a finitely presented group with unsolvable word problem, there is a natural map $\rho : F(X) \times F(X) \to Q \times Q$. Mihailova showed that $D = \rho^{-1}(\Delta)$ does not have solvable Membership Problem, since the Membership Problem for the diagonal subgroup is equivalent to the Word Problem [79].

The group $G_{BKS} = \langle \alpha, \beta, y \mid \alpha^y = \alpha\beta, \beta^y = \beta \rangle$, studied by Burns, Karass, and Solitar, is a 3-manifold group that is not LERF. Burns, Karrass, and Solitar showed that it is not LERF by proving that the subgroup $H_{BKS} = \langle y\alpha^{-1}y^{-2}\alpha,\ \alpha^{-1} \rangle$ is not separable [32]. Their proof is essentially reproduced in Section 2.4.

If G is a residually finite group, then every $H < G$ is also residually finite. The next lemma gives the equivalent statement for separable subgroups:

Lemma 1.2.7. *The intersection of a subgroup $H < G$ with a separable subgroup $S < G$ is separable in H. That is, $H \cap S$ is separable in H if S is separable in G.*

Proof. As S is separable, $S = \bigcap_{\substack{S \leq K < G \\ [G:K] < \infty}} K$. So

$$S \cap H = \left(\bigcap_{\substack{S \leq K < G \\ [G:K] < \infty}} K \right) \cap H = \bigcap_{\substack{S \leq K < G \\ [G:K] < \infty}} (K \cap H) .$$

This expresses $S \cap H$ as the intersection of a family of finite-index subgroups in H, since $K \cap H$ is finite index in H. $\square$

Lemma 1.2.8. *Finitely generated subgroups of LERF groups are LERF.*

Finding finitely generated subgroups that are not LERF is a common strategy for showing non-LERFness. The group G_{BKS} has been an important tool for verifying other examples of non-LERFness of groups. Niblo and Wise showed that G_{BKS} virtually embeds in the fundamental group of the complement of the link of 4 circles, the RAAG $L = \langle a, b, c, d \mid [a, b] = [b, c] = [c, d] = 1 \rangle$. Therefore L is not subgroup separable. Moreover they showed:

Proposition 1.2.9 (Niblo–Wise [81])**.** *The fundamental groups of compact graph manifolds have only one obstruction to subgroup separability: the existence of an embedding of L (and hence a virtual embedding of G_{BKS}).*

The classification of LERFness for RAAGs also hinges on not containing a few non-LERF subgroups.

Proposition 1.2.10 (Metafsis–Raptis [78])**.** *The group $A(\Gamma)$ is LERF if and only if Γ contains no square as an induced subgraph and has diameter no more than 2.*

These conditions are equivalent to Γ not containing either a square or P_3, the path of length 3, as an induced subgraph. These subgraphs correspond to the non subgroup-separable groups $F_2 \times F_2$ and $A(P_3) = L$. It is clear that not containing these subgraphs is a necessary condition. Metafsis and Raptis use induction on the number of vertices of Γ to establish that it is sufficient.

1.2.2 Groups that are residually finite

Residually finite groups include finite groups, free abelian groups, finitely generated nilpotent groups, supersolvable groups, and surface groups. An important tool, which tells us many residually finite groups is the following:

Proposition 1.2.11 (Mal'cev [74])**.** *If G is a linear group, that is, if G is a finitely generated subgroup of $\mathrm{GL}(n, \mathbb{C})$, then G is residually finite.*

Corollary 1.2.12. *Finitely generated free groups are residually finite.*

Proof. This can be proved directly, but also follows as $F_2 \leq SL(2, \mathbb{Z})$:
The subgroup generated by the matrices

$$\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$$

is isomorphic to F_2, so F_2 is a linear group and thus is residually finite. The finitely generated free groups are subgroups of F_2, so they are also residually finite. □

Right angled Artin groups are linear, so they are also residually finite (and in fact are residually torsion-free nilpotent). Other important families of residually finite groups are the following.

Theorem (Hempel [60], Perelman)**.** *The fundamental groups of compact 3-manifolds are residually finite.*

Theorem (Wise [90])**.** *One-relator groups with torsion, that is, groups of the form $\langle x_1, \ldots, x_m \mid r^n \rangle$ for $m, n \geq 2$, are residually finite.*

Wise constructs a quasi-convex hierarchy for these hyperbolic groups, and thus shows that they are virtually special. Virtually special hyperbolic groups are residually finite.

1.2.3 Groups that are not residually finite

Definition 1.2.13. A group G is *Hopfian* if every surjective endomorphism of G is an automorphism. If a group has a surjective endomorphism with a non-trivial kernel it is called a *non-Hopfian group.*

Lemma 1.2.14 (Mal'cev [75])**.** *Finitely generated residually finite groups are Hopfian.*

Proof. Suppose for the contradiction that there is some non-Hopfian group G which is finitely generated and residually finite. Let $\Phi : G \to G$ be a surjective endomorphism with non-trivial kernel K. Let Q be a finite quotient of G that witnesses that the kernel is non-trivial, that is, there is $k_1 \in K - \{1\}$ such that $\psi : G \to Q$ and $\psi(k_1) \neq 1$. The existence of ψ is a consequence of the residual finiteness of G. For all $i > 1$, choose k_i such that $\Phi(k_i) = k_{i-1}$. The family of maps Φ^n are all surjective, and they are all distinct:

$$\psi \circ \Phi^i(k_{i+1}) = \psi(k_1) \neq 1, \quad \text{but} \quad \psi \circ \Phi^{i+1}(k_{i+1}) = \psi(\Phi(k_1)) = \psi(1) = 1.$$

This implies that there are infinitely many distinct quotient maps from a finitely generated group to a single finite group, which is impossible. □

Mal'cev's Lemma provides an important obstruction to residual finiteness:

Example 1.2.15. The group $BS(2,3) = \langle a, t \mid (a^2)^t = a^3 \rangle$ is non-Hopfian, and therefore is not residually finite. The map $\Phi : a \mapsto a^2,\ t \mapsto t$ witnesses that $BS(2,3)$ is non-Hopfian. The image of $[a,t]$ under Φ is $[a^2, t] = a^{-2}a^3 = a$, so Φ is a surjection. Φ has non-trivial kernel, because $[t, a^{-1}] \neq [a,t]$ but $\Phi([t, a^{-1}]) = \Phi([a,t])$.

1.3 Residual solvability

We begin with a brief review of solvable groups. For any group G, the *derived series* is defined inductively by $G^{(0)} = G$, $G^{(i+1)} = [G^{(i)}, G^{(i)}]$. A group is *solvable* if there is n such that $G^n = \{1\}$. The smallest such n is called the *derived length* of the group. Recall the following basic facts about solvable groups.

Lemma 1.3.1. *If G is solvable of derived length l then:*

1. *Subgroups H of G are solvable of derived length at most l.*
2. *Quotients G/N of G are solvable of derived length at most l.*

Lemma 1.3.2. *The derived subgroup $G^{(i)}$ is characteristic in G, i.e. it is preserved by every automorphism of G.*

For G an arbitrary group, if $N \triangleleft G$, and G/N is solvable of derived length i, then $G^{(i)} \subset N$. This implies that every solvable quotient of derived length i factors through $G \to G/G^{(i)}$. Indeed

$$\{1\} = G^{(i)}/G^{(i)} \triangleleft \ldots G^{(1)}/G^{(i)} \triangleleft G/G^{(i)}$$

shows that the quotient will have derived length i, and every other quotient $Q = G/H$ satisfies that $G^{(i)} < H$.

Lemma 1.3.3. *If*

$$1 \to N \to G \to G/N \to 1$$

is a short exact sequence and N and G/N are solvable with derived lengths n and l, respectively, then G is solvable, and it has derived length no more than $n + l$.

Proof. If $p : G \to G/N$, then the derived subgroup $G^{(i)}$ has image $(G/N)^{(i)}$. Since G/N has derived length l, every element of $G^{(l)}$ maps to $\{1\}$, and so $G^{(l)} \leq N$. Therefore $G^{(l+i)} \leq N^{(i)}$, and so $G^{(n+l)} = \{1\}$. $\square$

In finite groups, the derived subgroups must eventually stablize. That is, there is some n such that $G^{(n)} = G^{(n+i)}$ for all $i \in \mathbb{N}$.

Definition 1.3.4. A group G is residually solvable if

$$\{1\} = \bigcap_{i \in \mathbb{N}} G^{(i)}.$$

From this definition, it is clear that finite groups are either solvable or not residually solvable, since their derived subgroups always stablize. It is worth noting that although residual finiteness extends from finite index subgroups up to their supergroup, this is not the case for residual solvability. For example, A_5 is not residually solvable, even though it contains $\{1\}$ as a finite index solvable subgroup. All finite groups are either solvable or not residually solvable. If the derived series stablizes to the identity, then G is solvable. Otherwise G is not residually solvable. The proper analogy to Lemma 1.2.5, that residually finite-by-finite groups are residually finite, is the following:

Lemma 1.3.5 (Gruenberg)**.** *Residually-solvable-by-solvable groups are residually solvable. That is, if P has a solvable quotient with a residually solvable kernel, then P is residually solvable.*

Proof. Let

$$1 \to K \to P \to P/K$$

where P/K is solvable and K is residually solvable. If $g \notin K$ then g survives in P/K, so we need only worry about $g \in K$. Since P/K is solvable, say of derived length k, this implies $P^{(k)} \subset K$, and thus $P^{(k+i)} \subset K^{(i)}$. The residual solvability of K implies $\cap_{i \in \mathbb{N}} P^{(i)} \subset \cap_{i \in \mathbb{N}} K^{(i)} = \{1\}$. Therefore P is residually solvable. □

The natural question whether this result extends to residually solvable by residually solvable groups was answered in the negative by Kahrobaei in [62]. Kahrobaei gave a family of counterexamples that are amalagamated products of the form $A *_C A$, where A is residually solvable and A/C is perfect, for example if $A = F_2$ and $A/C = A_5$. She showed that these groups are residually solvable- by-residually solvable but not residually solvable.

We remind the reader of a classical result about subgroups of free products.

Lemma 1.3.6. *If $K < A * B$, and $K \cap A = K \cap B = \{1\}$, then K is free.*

Lemma 1.3.7. *If A and B are finitely generated residually solvable groups, then $A * B$ is residually solvable.*

Proof. First we show this for A and B solvable. Consider the map $\phi : A * B \to A \times B$. The group $A \times B$ is solvable. The kernel of this map is free, since A and B are preserved. Therefore $A \times B$ is a residually solvable by solvable group, and by Lemma 1.3.5, $A * B$ is residually solvable. Now suppose A and B are residually solvable. If $w \in A * B$ is non-trivial, then it can be written $a_1 b_1 \ldots a_n b_n$ for only a_1, b_n possibly trivial. As A and B are residually solvable groups, there is some $k \in \mathbb{N}$ such that all (non-trivial elements) of $a_1, \ldots a_n$ and $b_1 \ldots b_n$ are not in $A^{(k)}$ and $B^{(k)}$, respectively. Therefore w is not trivialized when it maps to $\bar{w}$ in the quotient $A/A^{(k)} * B/B^{(k)}$. By the residual solvability of $A/A^{(k)} * B/B^{(k)}$, there is a solvable quotient that does not trivialize $\bar{w}$. $\square$

1.3.1 Groups that are residually solvable

Solvable groups are trivially residually solvable. Free groups are residually solvable. It is easy to guess that the length of shortest elements grows with the index of the derived subgroup, and indeed, Elkasapy and Thom showed that for all $w_i \in G^{(i)}$, $3^i \leq |w_i|$ [47], so the intersection of the derived subgroups is trivial. Many groups (including free groups) satisfy stronger residual properties, such as residual nilpotence. Groups which are residually solvable because they are are residually nilpotent include surface groups [7] and right angled Artin groups [46]. This is a common strategy for establishing residual solvability.

Kropholler showed that the Baumslag–Solitar groups $BS(n, m)$ have free second derived subgroups [69]. Baumslag showed that all positive one-relator groups are residually solvable [9]. In her thesis, Kahrobaei studied conditions on amalgamated products that guarantee residual solvability. She showed for instance that doubles of residually solvable groups over solvably separable subgroups are residually solvable [62].

1.3.2 Groups that are not residually solvable

Perfect subgroups are an important obstruction to residual solvability.

Definition 1.3.8. A group H is *perfect* if $[H, H] = H$.

Lemma 1.3.9. *If* $1 < H < G$ *is a non-trivial perfect subgroup, then* G *is not residually solvable.*

Proof. If H is perfect, then $H \leq \cap_{i\in\mathbb{N}} G^{(i)}$ for all i, and so $\cap_{i\in\mathbb{N}} G^{(i)} \neq \{1\}$. □

1.4 Atlas of residual properties

In this section, we define a number of additional residual properties and build an 'atlas of residual properties' for the finitely presented groups that shows how some of these properties fit together. The outermost families relax the requirement

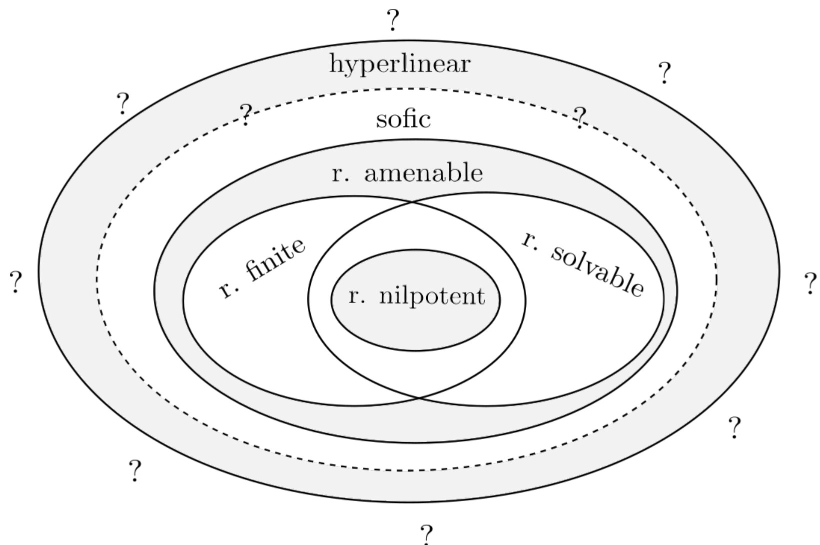

Figure 1.1: The 'residual atlas' for finitely generated groups

that maps are homomorphisms [34]. For any finite set $F \subset G$, and any $\epsilon > 0$, the outermost families require a family of functions ϕ from G into a family of metric groups, such that ϕ can be chosen to be as close to a homomorphism as we like on F (measured by ϵ), and elements of F stay uniformly away from the identity. Sofic groups map into symmetric groups with the Hamming metric:

$$d_{S_n}(\sigma_1, \sigma_2) = \frac{1}{n}\left(\# \text{ of } k \text{ s.t. } \sigma_1(k) \neq \sigma_2(k)\right).$$

Definition 1.4.1. A group G is *sofic* if for every $\epsilon > 0$ and for every finite set $F \subset G - \{1\}$, there is an $n \in \mathbb{N}$ and $\phi : G \to S_n$ such that

1. $\phi(1_G) = 1_{S_n}$

2. $\forall\, x, y \in F,\ d_{S_n}(\phi(xy), \phi(x)\phi(y)) < \epsilon$. That is, ϕ is almost a homomorphism on F.

3. There is constant $0 < r(g) < 1$ such that for every n, $\phi(g)$ acting on $\{1, \ldots, n\}$ moves at least $r(g)n$ of the elements. That is, $g \in F$ is not sent too close to the identity in S_n.

In hyperlinear groups, the family of metric groups are the unitary groups $U(n)$, the unitary matrices in $\mathrm{GL}_n(\mathbb{C})$. The metric comes from the length function $l_{U(n)}(u) = \frac{1}{2}\sqrt{\mathrm{trace}(vv^*)}$, where $v = u - I_n$.

Definition 1.4.2. A group G is *hyperlinear* if for every $\epsilon > 0$ and for every finite set $F \subset G - \{1\}$, there is an $n \in \mathbb{N}$ and $\phi : G \to U(n)$ such that $\phi(1) = 1$, $d_{U(n)}(\phi(xy), \phi(x)\phi(y)) < \epsilon$ for all $x, y \in F$, and for every $g \in F$, there is constant $r(g) > 0$ such that for every n, $l_{U(n)}(\phi(g)) \geq r(g)$.

Hyperlinearity and soficity are often established by showing that the group has stronger residual properties. Cornulier provided the first example of a sofic group that is not residually amenable [39]. Kar and Nikolov showed that an amalgamated double of $\mathrm{SL}_n(\mathbb{Z}[\frac{1}{p}])$ over a cyclic subgroup is sofic but not residually amenable [66].

Open problem 1.4.3. Are all groups hyperlinear? Are all groups sofic?

Possible candidates for groups that might not be hyperlinear are Thompson's groups. Thompson's groups T and V are simple and are not finite, solvable, or amenable, so they have neither finite, solvable, nor amenable quotients.

Since both solvable and finite groups are amenable, residually solvable and residually finite groups are automatically residually amenable. The direct sums preserve residual amenability. This allows us to construct examples of residually amenable groups that are neither residually solvable nor residually finite. For example $BS(2,3) \times A_5$ contains subgroups which are isomorphic to $BS(2,3)$ and A_5, so this group is neither residually solvable nor residually finite.

To build further examples, we can use the following:

Proposition 1.4.4. *Some extensions preserve residual properties:*

1. *(residually finite)-by-finite groups are residually finite.*
2. *(residually finite)-by-cyclic groups are residally finite.*
3. *(residually solvable)-by-solvable groups are residually solvable.*
4. *(residually solvable)-by-amenable groups are residually amenable.*
5. *(residually amenable)-by-amenable groups are sofic.*

In [10], Berlai shows that $\mathrm{Sym}_0(\mathbb{Z})$ Wr $\mathbb{Z}$ is a (residually amenable)-by-amenable group, but is not residually amenable.

Definition 1.4.5. $\mathcal{P}$ is a root property if:

1. Subgroups of $\mathcal{P}$-groups are also $\mathcal{P}$-groups.
2. Finite products of $\mathcal{P}$ groups are also $\mathcal{P}$-groups.
3. If $H \triangleleft K \triangleleft G$ and K/H and G/K are $\mathcal{P}$-groups, there is $L < H$ such that G/L is a $\mathcal{P}$-group.

Both solvability and finiteness are examples of root properties, while amenability is not a root property.

Proposition 1.4.6. *For the following $\mathcal{P}$, residually $\mathcal{P}$ groups are closed under taking free products:*

1. *$\mathcal{P}$ is finiteness.*
2. *$\mathcal{P}$ is solvability.*
3. *$\mathcal{P}$ is amenability.*

Further, sofic groups are closed under taking free products.

The second statement is Lemma 1.3.7, and the first statement can be proved in the same way. The third statement is due to Berlai and proved in [10].

Proposition 1.4.7 (Sokolov [88]). *If $\mathcal{P}$ is a root property, then wreath products of residually $\mathcal{P}$ groups are residually $\mathcal{P}$.*

This allows us to build more complicated examples of residually $\mathcal{P}$ groups.

Our favorite example, the free groups, are residually torsion-free nilpotent, so they are examples of each of these properties. Where the rest of the hyperbolic groups live in this map, and how far they extend is not known. Finite groups that are not nilpotent, like A_5, are trivial examples showing that some hyperbolic groups are not residually nilpotent or residually solvable.

Open problem 1.4.8. Are hyperbolic groups always residually finite? Are hyperbolic groups always residually amenable?

These two problems are actually equivalent; Arzhantseva showed in [3] that if there is a hyperbolic group that is not residually finite, there is also a hyperbolic group that is not residually amenable. Such a result would be important for topology, given the topological application proved by Scott [86].

Proposition 1.4.9. *Let X be a Hausdorff space and $\tilde{X}$ its universal cover. Then $\pi_1(X)$ is residually finite if and only if for any compact $C \subset \tilde{X}$, there is a finite-sheeted cover $\widehat{X}$ of X such that $\tilde{X} \to \widehat{X}$ is injective on C.*

This implies for example, that if an immersed surface or curve in X is embedded in some cover of X, that there is a finite-sheeted cover of X in which it is embedded.

In Chapter 2, we show that the hydra group doubles Γ_k are not residually finite. We do not know at this time where these groups live in the atlas.

1.5 Quantifying residual properties

The study of the quantification of residual properties began in 2012, when Khalid Bou-Rabee introduced the normal residual finiteness growth function (depth) for residually finite groups [19].

Definition 1.5.1. If $G = \langle X \mid R \rangle$ is a residually finite group and $g \in G - \{1\}$, define Q_g to be the smallest finite group such that $\psi : G \to Q_g$ and $\psi(g) \neq 1_{Q_g}$. Then $\mathrm{nfd}_X : \mathbb{N} \to \mathbb{N}$ is defined by

$$\mathrm{nfd}_X(n) := \max\{|Q_g| \mid |g|_X \leq n\}.$$

Bou-Rabee calculated nfd for nilpotent groups, and found lower bounds for nfd for $\mathrm{SL}(n, \mathbb{Z})$ and the Grigorchuk group, and set out the basic theory of this function. In Chapter 3, we will emulate this theory, so we present the basic results of nfd here. As with the Dehn function, we need an equivalence relation to handle the changes in the metric we get from changing presentation; say that $f \preceq g$ if $f(n) \leq Cg(Cn)$ and $f \asymp g$ if $f \preceq g$ and $g \preceq f$.

Proposition (Bou-Rabee)**.** *If $H \leq G$, then with respect to any finite generating set S for H and T for G, then* $\mathrm{nfd}_S \preceq \mathrm{nfd}_T$.

In particular, by taking S and T to be different finite generating sets for G, and then swapping, this proposition implies that changing presentations does not change the residual finiteness growth, up to the equivalence relation.

Proposition (Bou-Rabee)**.** *If $H \leq G$ are both finitely generated and $[G : H] < \infty$, then*

$$\mathrm{nfd}_G(n) \preceq \mathrm{nfd}_H(n)^{[G:H]}.$$

Proposition (Bou-Rabee)**.** *The group $\mathbb{Z}$ satisfies* $\mathrm{nfd}_{\mathbb{Z}}(n) \asymp \log(n)$.

The residual finiteness growth function has been particularly difficult to pin down for free groups. Currently the best known bounds are:

Theorem (Kassabov–Mattuci, Bou-Rabee)**.**

$$n^{\frac{2}{3}} \preceq \mathrm{nfd}_{F^2}(n) \preceq n^3.$$

Kassabov and Matucci build elements in F_2 requiring sufficiently large finite quotients to verify non-triviality. They use bounds on the size of identities, which are words in the free group which are trivial in all finite groups of a certain size, to establish their results [67]. Bou-Rabee's proof of the upper bound comes from his n^3 estimate of the residual finiteness growth for the group $\mathrm{SL}(2, \mathbb{Z})$.

Open question 1.5.2 (Kassabov–Matucci)**.** Is the normal residual finiteness growth of F_2 linear?

If instead of looking at the size of finite quotients that witness the nontrivial element g, we look at the minimal index $[G : H]$ such that $g \notin H$, we get the non-normal residual finiteness growth function, fd.

In her thesis [84] Patel quantified Scott's result that surface groups are LERF, and also showed that if Σ is a compact surface with negative Euler characteristic, and $\alpha \in \pi_1(\Sigma) - \{1\}$, then the minimum index of a subgroup H such that $\alpha \notin H$ is linear in the length of the geodesic representative of α. In particular, this implies a linear upper bound on fd for surface groups. Bou-Rabee, Hagen, and Patel recently found that the non-normal residual finiteness growth for F_2 is linear [21].

Theorem (Bou-Rabee-Hagen-Patel)**.** *The non-normal residual finiteness growth function is at most linear for virtually special groups. Virtually special groups include F_2 and all other RAAGs.*

Questions about how big the residual finiteness growth function could be were answered when Kharlampovich, Myasnikov, and Sapir showed that for any recursive function f, there is a residually finite group G with normal residual finiteness growth that dominates f [68]. They also constructed residually finite groups with Dehn function dominating f. Prior to this paper it was unknown whether or not residually finite groups could have super-exponential Dehn function. Their paper motivated Chapter 2 of this thesis.

In [20], Bou-Rabee examines the size of finite solvable quotients needed to distinguish elements from the identity. He shows, among other things, that if G has a finitely generated finite index solvable subgroup H, that it is possible to determine whether or not G is virtually nilpotent from his solvability growth function. His function is only defined for residually finite groups. In Chapter 3, we examine another method of quantifying residual solvability, that is also defined for groups that are not residually finite. We examine the derived length of the solvable quotient needed to show that $g \in G - \{1\}$ is not the identity.

CHAPTER 2

HYDRA GROUPS ARE NOT RESIDUALLY FINITE

2.1 Introduction

The first examples of finitely presented residually finite groups with super-exponential Dehn function were constructed in [68]:

Theorem 2.1.1 (Kharlampovich, Myasnikov, and Sapir)**.** *For any recursive function $f : \mathbb{N} \to \mathbb{N}$, there is a finitely presented residually finite solvable group G of derived length 3 for which the Dehn function $\delta_G \succcurlyeq f$.*

Their examples are sufficiently complicated that it remains interesting to find elementary examples that arise 'in nature'. One place to look is among known elementary examples of groups with large Dehn function.

Definition 2.1.2. If $H < G$ is generated by S, then the *HNN double over H* is

$$\Gamma = \langle G, p \mid [p, s], \quad s \in S \rangle$$

and the *amalgamated product double over H* is

$$\Gamma' = G *_{H=H} G = \langle G, \bar{G} \mid s = \bar{s}, \quad s \in S \rangle.$$

In [43], Dison and Riley introduced the hydra groups

$$G_k := \langle a_1, \ldots, a_k, t \mid a_1^t = a_1, \ a_i^t = a_i a_{i-1}, \ i > 1 \rangle.$$

They proved that Γ_k, the HNN double over the subgroup $H_k = \langle a_1 t, \ldots, a_k t \rangle$, and Γ'_k, the amalgamated product double over the same subgroup H_k, have Dehn functions equivalent to the Ackermann function $A_k(n)$.

The authors of [68] commented that it was unknown whether or not Γ_k is residually finite for all $k > 1$, but that they expected Γ_k would not be residually finite for $k > 1$. We confirm this.

Theorem 2.1.3. *For all $k > 1$, the groups Γ_k and Γ'_k, are not residually finite.*

In Section 2.3 we will see that separability of the subgroup H_k in G_k is necessary for the residual finiteness of Γ_k and Γ'_k. Theorem 2.1.3 is proven via:

Lemma 2.1.4. *The group H_k is not separable in G_k for any $k > 1$.*

In particular, we show that the non-separability of H_2 in G_2 implies non-separability of H_k in G_k. To see that H_2 is not a separable subgroup of G_2, we recognize (G_2, H_2) as isomorphic to the important group-subgroup pair (G_{BKS}, H_{BKS}) studied by Burns, Karrass, and Solitar in [32].

Dison and Riley produced further group-subgroup pairs that are candidates to be finitely presented groups with fast-growing Dehn functions that are residually finite. For $\mathbf{w} = (w_1, \ldots, w_k)$, where w_i is a positive word on letters in $\{a_1, \ldots, a_{i-1}\}$, define

$$G_k(\mathbf{w}) := \langle a_1, \ldots, a_k \mid a_i^t = a_i w_i,\ 1 \leq i \leq k \rangle.$$

For powers $\mathbf{r} = (r_1, \ldots, r_k)$, where $r_i \geq 0$, define the subgroup

$$H_k(\mathbf{r}) = \langle a_1 t^{r_1}, \ldots, a_k t^{r_k} \rangle.$$

In many cases, we prove that such group-subgroup pairs cannot be used to produce residually finite groups with large Dehn function. In particular:

Theorem 2.1.5. *H_k is a separable subgroup of $G_k(\mathbf{w})$ if and only if $\mathbf{w} = (1, \ldots, 1)$. Therefore the HNN double $\Gamma_k(\mathbf{w}) = \langle G_k(\mathbf{w}), p \mid h^p = h,\ h \in H_k \rangle$ is residually finite only if $G_k(\boldsymbol{w}) = F_k \times \mathbb{Z}$.*

Theorem 2.1.6. *Suppose that* $\mathbf{w} = (1, \ldots, 1, w_c, \ldots, w_k)$ *where* $w_c \neq 1$ *and* $\mathbf{r} = (r_1, \ldots, r_k)$*. Let* $[w_c]_i$ *denote the index-sum of* a_i *in* w_c*. If*

$$\sum_{i=1}^{c-1} [w_c]_i r_i \neq 0$$

then $H_k(\mathbf{r})$ *is a non-separable subgroup of* $G_k(\mathbf{w})$*.*

Remark 2.1.7. We do not have an answer to the separability of $H_k(\mathbf{r})$ in $G_k(\mathbf{w})$ when $\sum_{i=1}^{c-1}[w_c]_i r_i = 0$. For example, we do not know whether or not $\langle a_1, a_2 t, a_3\rangle$ is a separable subgroup of $\langle a_1, a_2, a_3, t \mid a_1^t = a_1, a_2^t = a_2 a_1, a_3^t = a_3 a_2\rangle$.

2.2 Dehn functions and residual finiteness of doubles

We provide here the basic ideas of Dison and Riley's proof from [43] that the HNN doubles Γ_k and the amalgamated product doubles Γ'_k have large Dehn function.

Definition 2.2.1. If $H \leq G$, are generated by S and T, respectively, the *distortion function*, $\mathrm{Dist}_H^G(n)$ measures how much longer elements of H are when written in S compared to T:

$$\mathrm{Dist}_H^G(n) := \min\{K \mid B_T(1,n) \cap H \subset B_S(1,K)\} = \max\{|g|_S \text{ s.t.} |g|_T \leq n\}.$$

Lemma 2.2.2. *The Dehn function of* Γ *satisfies* $n\mathrm{Dist}_H^G(n) \preceq \delta_\Gamma(n)$*.*

Proof. Let u_n be a witness to $\mathrm{Dist}_H^G(n)$. That is, let $u_n \in H$ and $|u_n|_G \leq n$ but $|u_n|_H = \mathrm{Dist}_H^G(n)$. Let v_n be u_n rewritten in the generators of H.

Consider $w_n = [p^n, u_n^{-1}]$, as in Figure 2.1a. In any diagram Δ for w_n, there is no choice in how p-edges are paired. Each p-edge is paired with the corresponding edge in the segment on the opposite side of the diagram. The words along the tops and bottoms of these p-corridors are words in the generators of H, and in particular, these words are always equal in H to v_n. Since there are n such p-corridors, we get $n\mathrm{Dist}_H^G(n) < \mathrm{Area}(\Delta) \leq \delta_\Gamma(4n)$. $\square$

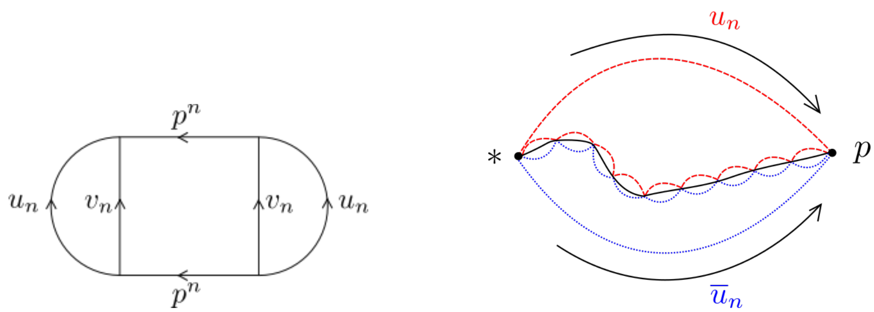

(a) v.K. diagram for the HNN double Γ

(b) v.K. diagram for the amalgamated product double Γ'

Figure 2.1: van Kampen diagrams for doubles

Lemma 2.2.3 (Bridson [26]). *The Dehn function of Γ' satisfies* $\mathrm{Dist}_H^G(n) \preceq \delta_{\Gamma'}(n)$.

Proof. Suppose that $G = \langle X|R \rangle$, and Y is the generating set of H. Let u_n be a witness to $\mathrm{Dist}_H^G(n)$, and let $w_n = u_n(\overline{u_n})^{-1}$. Let Δ be a van Kampen diagram for w_n. We will show that there is a 'path' that touches $\mathrm{Dist}_H^G(n)$ many 2-cells, so $\mathrm{Area}(\Delta)$ is at least $\mathrm{Dist}_H^G(n)$. The 2-cells in Δ may be divided into three categories: those in the generators X, those in the generators $\overline{X}$, and those of the form $\overline{y}y^{-1}$, where $y \in Y$. By drawing in a path in the interior of each 2-cell in the third category, we can build a (possibly disconnected) planar graph $\mathcal{G}$ in Δ. Since the points $*$ and p are contained in $\mathcal{G}$, we ask if they are in the same component, so that we can draw in a path as in Figure 2.1b. If they are not in the same component, they can be separated, so we can find a path in the diagram which starts at the top and ends on the bottom, that separates the components — it will not cross $\mathcal{G}$. Since the top is written in X and the bottom is written in $\overline{X}$, we must translate between X and $\overline{X}$ by crossing a 2-cell of the third kind. We can push the path from $*$ to p always to the Y side (as opposed to the $\overline{Y}$ side), so the path is labeled by a word in $Y \cup Y^{-1}$, which represents u_n. The length of this word in Y is a lower bound for the area, so in particular $\mathrm{Dist}_H^G(n) \le \delta_{\Gamma'}(n)$. □

To build groups with large Dehn function, Dison and Riley found a family of group-subgroup pairs (G_k, H_k) such that the H_k are very distorted in G_k [43]:

$$G_k := \langle a_1, \ldots, a_k, t \mid a_1^t = a_1,\ a_i^t = a_i a_{i-1},\ i > 1 \rangle$$

$$H_k = \langle a_1 t,\ \ldots,\ a_k t \rangle$$

The Ackermann function $A_k : \mathbb{N} \cup \{0\} \to \mathbb{N} \cup \{0\}$ is defined as follows:

$$A_0(n) = n + 2$$

$$A_1(n) = 2n$$

$$A_k(0) = 1,\ k \ge 2$$

$$A_k(n+1) = A_{k-1}(A_k(n)),\ k \ge 2$$

Then $A_2(n+1) = A_1(A_2(n)) = 2(A_2(n)) = 2^n$, while $A_3(n+1) = 2^{A_3(n)}$, so $A_3(n)$ is the iterated exponential $2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}$ where there are n 2's in each tower.

Lemma (Dison–Riley)**.** $\mathrm{Dist}_{H_k}^{G_k}(n) \le A_k(n)$

Proof. Consider the word $a_k^n a_2 t a_1 a_2^{-1} a_k^{-n}$. In G_k, this is a word with length no more than $2n + 2$. In Lemma 2.3.3, we will see how to recognize when words are in H_k, and how to find their length. It turns out that the length of this word, when written in the generators of H_k, has length greater than $2A_k(n)$. □

Using Lemma 2.2.2 to establish the lower bound, Dison and Riley showed that the Dehn functions of

$$\Gamma_k = \langle G_k, p \mid [a_i t, p] = 1,\ \forall\ 1 \leq i \leq k \rangle$$

$$\Gamma'_k = \langle G_k, \overline{G}_k \mid a_i t = \overline{a_i t},\ \forall\ 1 \leq i \leq k \rangle$$

are equivalent to the Ackermann function $A_k(n)$, for every k.

If Γ_k and Γ'_k were residually finite, we would see the following behavior:

Lemma 2.2.4. *If* $\Gamma = \langle G, p \mid h^p = h,\ h \in H \rangle$ *is residually finite, then* H *is separable in* G *and* G *is residually finite.*

Proof. G is residually finite since residual finiteness is inherited by subgroups. Suppose that H is not separable. We can find $g \in G - H$ such that under any homomorphism ϕ from G to a finite group, $\phi(g) \in \phi(H)$. For any map $\Phi : \Gamma \to Q$ for Q finite, consider the image of the non-trivial element $[p, g]$ under Φ. Since Φ restricts to a homomorphism on G, $\Phi(g) = \Phi(h)$ for some $h \in H$, so

$$\Phi([p, g]) = [\Phi(p), \Phi(g)] = [\Phi(p), \Phi(h)] = \Phi([p, h]) = \Phi(1) = 1.$$

Therefore, if H is not separable in G, Γ is not residually finite. $\square$

By a theorem of Baumslag and Tretkoff, this necessary condition is sufficient: if G is residually finite and H is separable, then Γ is residually finite [6].

Remark 2.2.5. For any property $\mathcal{P}$, the same proof shows that if H is not $\mathcal{P}$-separable, then Γ is not residually $\mathcal{P}$. Berlai has shown that for $\mathcal{P}$ the properties of solvability and amenability, when G is residually $\mathcal{P}$, Γ is residually $\mathcal{P}$ if and only if H is $\mathcal{P}$ separable in G [10].

Lemma 2.2.6. *If* $\Gamma' = \langle G, \overline{G} \mid h = \bar{h},\ h \in H \rangle$*, is residually finite, then* G *is residually finite and* H *is separable in* G*.*

Proof. If Γ' is residually finite, G will be too. If $g \in G - H$ is an element that cannot be separated from H in finite quotients, we will show that $g^{-1}\overline{g} \in \Gamma' - \{1\}$ is trivialized in every finite quotient. An arbitrary map Φ from Γ' to a finite group Q will factor as a pair of maps $\psi, \phi : G \to Q$ such that $\psi(h) = \phi(h)$ for all $h \in H$. To see that $\Phi(g^{-1}\overline{g}) = 1$, we show that ϕ and ψ agree on g as well. Construct $\Psi : G \to Q \times Q$, which is (ψ, ϕ). This new target group is still finite, and the image of H is contained in the diagonal. As $\Psi(g) \in \Psi(H) \subset \Delta$, this implies that $\psi(g) = \phi(g)$. Therefore $\Phi(g^{-1}\overline{g}) = 1$. □

Remark 2.2.7. If taking the direct product of groups preserves property $\mathcal{P}$ (eg. solvability, amenability), then if H is not $\mathcal{P}$-separable in G, the same proof as above implies that $\Gamma' = \langle G, \overline{G} \mid H = \overline{H} \rangle$ is not residually-$\mathcal{P}$. Kahrobaei and Berlai have shown that for $\mathcal{P}$ the properties of solvability and amenability respectively, so long as G is residually $\mathcal{P}$, Γ' is residually $\mathcal{P}$ if and only if H is $\mathcal{P}$ separable in G [63, 10].

2.3 H_k is not a separable subgroup of G_k

Lemmas 2.2.4 and 2.2.6 imply that we can show that the HNN doubles Γ_k and the amalgamated product doubles Γ'_k are not residually finite by recognizing that the subgroup H_k is not separable in G_k. In Lemma 1.1.5, we saw that if H is normal in G, then H is separable if and only if G/H is residually finite. When H is not normal, the situation is more difficult. To show that these subgroups are not separable, we use Lemma 1.2.7, and isomorphisms to (H_{BKS}, G_{BKS}). We begin by showing that H_2 is not separable in G_2.

Lemma 2.3.1. *H_2 is not a separable subgroup of G_2.*

Proof. To show that G_2 is not H_2-separable, we use the result of Burns, Karrass, and Solitar [32] that

$$H_{BKS} = \langle y\alpha^{-1}y^{-2}\alpha,\ \alpha^{-1}\rangle \quad \leq \quad G_{BKS} = \langle \alpha, \beta, y \mid \alpha^y = \alpha\beta, \beta^y = \beta\rangle$$

is not separable. We demonstrate an automorphism ϕ of G_2 from which the isomorphism carrying $(G_2, \phi(H_2))$ to (G_{BKS}, H_{BKS}) is clear. Recall that:

$$G_2 = \langle a_1, a_2, t \mid a_1^t = a_1,\ a_2^t = a_2a_1\rangle = \langle a_1, a_2, t \mid [a_1, t] = 1,\ [a_2, t] = a_1\rangle$$

$$H_2 = \langle a_1t,\ a_2t\rangle \leq G_2$$

Consider the map ϕ defined by:

$$\begin{aligned} a_2 &\mapsto a_2^{-2}ta_2 \\ t &\mapsto a_2^{-1}t^{-1}a_2 = a_1t^{-1} \\ a_1 &\mapsto a_1. \end{aligned}$$

We verify that ϕ is an endomorphism:

$$\begin{aligned} [\phi(a_1), \phi(t)] &= [a_1, a_1t^{-1}] = a_1^{-1}ta_1^{-1}a_1a_1t^{-1} = 1 = \phi([a_1, t]), \\ [\phi(a_2), \phi(t)] &= [a_2^{-2}ta_2, a_2^{-1}t^{-1}a_2] = [a_2^{-1}t,\ t^{-1}]^{a_2} = a_2^{-1}(t^{-1}a_2ta_2^{-1}tt^{-1})a_2 \\ &= [a_2, t] = a_1 = \phi(a_1). \end{aligned}$$

ϕ is also surjective:

$$\begin{aligned} a_1 &= \phi(a_1) \\ a_2 &= \phi((a_2t)^{-1}) \\ t &= \phi(t^{-1}a_1) \end{aligned}$$

In [8], G. Baumslag proved that finitely generated free-by-cyclic groups are residually finite, hence Hopfian. Because G_2 is Hopfian and ϕ is a surjective endomorphism, ϕ is an automorphism. Further, $\phi(H_2) = \langle a_2^{-1}, ta_2^{-1}t^{-2}a_2\rangle$:

$$a_1 t \mapsto a_1 a_2^{-1}t^{-1}a_2 = [a_2, t]a_2^{-1}t^{-1}a_2 = (a_2^{-1}t^{-1}a_2)t(a_2^{-1}t^{-1}a_2) = ta_2^{-1}t^{-2}a_2$$
$$a_2 t \mapsto a_2^{-2}ta_2a_2^{-1}t^{-1}a_2 = a_2^{-1}$$

where we use that t and $a_2^{-1}t^{-1}a_2$ commute. The isomorphism between $(G_2, \phi(H_2))$ and (G_{BKS}, H_{BKS}) is given by $a_2 \mapsto \alpha,\ a_1 \mapsto \beta$, and $t \mapsto y$. Burns, Karrass, and Solitar showed that H_{BKS} is not separable in G_{BKS}, so H_2 is not separable in G_2 [32]. □

Next we show that the non-separability of H_k in G_k follows from the non-separability of H_2 in G_2. There is a natural inclusion $G_2 \hookrightarrow G_k$. In the following we will abuse notation and write G_2 and H_2 for the image in G_k of G_2 and H_2 under this inclusion. Dison and Riley develop a description of elements of $G_2 \cap H_k$ in [43], which we include here.

Definition 2.3.2. Let *priority* be the ordering on generators where $a_{i+1} > a_i$. The *piece decomposition* of a word $u \in F_k = \langle a_1, \ldots, a_k\rangle$, is a grouping $u \cong \pi_1 \cdots \pi_l$ where π_i are maximal words of the form $w_i, a_k w_i, w_i a_k^{-1}$ or $a_k w_i a_k^{-1}$, and w_i is a word in $\{a_1, \ldots, a_{k-1}\}$.

Lemma 2.3.3 (Dison–Riley)**.** *A word $w = t^r u$ represents an element of H_k if and only if u has piece decomposition $u = \pi_1 \cdots \pi_l$ and these pieces satisfy that for $p_0 = r$, there is a p_i such that $t^{p_i}\pi_{i+1} \in H_k t^{p_{i+1}}$ and $p_l = 0$. When p_{i+1} exists satisfying $t^{p_i}\pi_{i+1} \in H_k t^{p_{i+1}}$, it is unique.*

For $H_2 < G_2$, we have the following simple description with which to work out whether or not a sequence of (p_i) exists. Then we illustrate Lemma 2.3.3 with an example.

1. If $\pi_i = a_1^{k_i}$, then $t^{p_{i-1}}\pi_i = (a_1t)^{k_i}t^{p_{i-1}-k_i}$ so $p_i = p_{i-1} - k_i$.
2. If $\pi_i = a_2a_1^{k_i}$, then $t^{p_{i-1}}\pi_i = a_2a_1^{k_i-p_{i-1}}t^{p_{i-1}} = a_2t(a_1t)^{k_i-p_{i-1}}t^{2p_{i-1}-k_i-1}$ so $p_i = 2p_{i-1} - k_i - 1$.
3. If $\pi_i = a_1^{k_i}a_2^{-1}$, then p_i exists when $p_{i-1} - k_i$ is odd. Let $p_{i-1} - k_i = 2\alpha_i - 1$. Then $t^{p_{i-1}}a_1^{k_i}a_2^{-1} = (a_1t)^{k_i}t^{p_{i-1}-k_i}a_2^{-1} = (a_1t)^{k_i+\alpha_i}(a_2t)^{-1}t^{\alpha_i}$, so $p_i = \alpha_i$.
4. If $\pi_i = a_2a_1^{k_i}a_2^{-1}$, then p_i exists when k_i is even. Let $k_i = 2\alpha_i$. Then $t^{p_{i-1}}a_2a_1^{2\alpha_i}a_2^{-1} = (a_2t)(a_1t)^{2\alpha_i}t^{2p_{i-1}-2\alpha_i-1} = (a_2t)(a_1t)^{p_{i-1}}(a_2t)^{-1}t^{p_{i-1}-\alpha_i}$, so $p_i = p_{i-1} - \alpha_i$.

Example 2.3.4. Let $u = a_1a_2^{-1}a_1^2a_2a_1^{-1}a_2a_1^2a_2^{-2}$. Then u has piece decomposition $(a_1a_2^{-1})(a_1^2)(a_2a_1^{-1})(a_2a_1^2a_2^{-1})(a_2^{-1})$, where the parentheses indicate the pieces. Further $t^2u \in H_2t^{-1}$: $p_0 = 2, p_1 = 1, p_2 = -1, p_3 = -2, p_4 = -3, p_5 = -1$. In contrast, tu cannot be written as H_2t^k for any k: $p_0 = 1$, but p_1 is not defined.

Lemma 2.3.5 (Dison–Riley)**.** *$G_2 \cap H_k = H_2$.*

Proof. From the definition, it is clear that $H_2 \subset G_2 \cap H_k$. Suppose $w \in G_2 \cap H_k$. Using the free-by-cyclic normal form for G_k, rewrite w as t^ru where $u \in \langle a_1, \ldots, a_k\rangle$. Since $w \in G_2$, $u \in \langle a_1, a_2\rangle$. Thus the maximum priority letter occurring in words in $G_2 \cap H_k$ is a_2. By Lemma 2.3.3, since $t^ru \in H_k$, u has a piece decomposition $u = \pi_1 \cdots \pi_l$ such that $t^{p_i}\pi_{i+1} \in H_kt^{p_{i+1}}$. As $t^{p_i}\pi_{i+1}$ contains no letters of priority greater than 2, $t^{p_i}\pi_{i+1} \in H_2t^{p_{i+1}}$. From Lemma 2.3.3, $w \in H_2$, so $G_2 \cap H_k = H_2$. □

Proof of Lemma 2.1.4. Since $H_2 = G_2 \cap H_k$ is not separable in G_2, Lemma 1.2.7 implies that H_k is not separable in G_k when $k \neq 1$. □

Lemmas 2.2.4 and 2.2.6 showed that separability of H_k in G_k is necessary for the residual finiteness of Γ_k and Γ'_k. Thus Γ_k and Γ'_k are not residually finite.

2.4 Generalizations of the Hydra groups

In the last section we saw that for every k, H_k is not a separable subgroup of G_k. We were interested in these groups because Dison and Riley showed that H_k is distorted like the Ackermann function A_k in G_k, which forces the Dehn function of the doubles to be large. In this section we consider other pairs for which the machinery of Dison and Riley show that the analogous HNN doubles have exponential or superexponential Dehn function. We will show that that many of these groups are not residually finite.

The following proposition is an extension of the example of Burns, Karrass, and Solitar in [32]. It is the key to proving Theorem 2.1.5: the subgroup H_k is separable in the generalized hydra group $G_k(\mathbf{w})$ only when $G_k(\mathbf{w}) = F_k \times \mathbb{Z}$.

Proposition 2.4.1. *If $r > 0$ and $s \geq 0$, $H_2(r, s) = \langle a_1 t^r, a_2 t^s \rangle$ is not separable in G_2. The subgroup $H_2(0, s)$ is separable for every s.*

Remark 2.4.2. We restricted $r \geq 0$ so that the techniques of Dison and Riley still hold. In particular, we wish to have the analogue of Lemma 2.3.3, which told us whether or not an element was in H_k.

Lemma 2.4.3. *$H_2(r, s)$ is separable if and only if $H_2(r, 0)$ is separable.*

Proof. For every $s \in \mathbb{Z}$ there is an automorphism, η_s of G_2 carrying $H_2(r, 0)$ to $H_2(r, s)$.

$$\eta_s : \quad a_2 \mapsto a_2 t^{s-1}, \quad t \mapsto t, \quad a_1 \mapsto a_1$$

This is an endomorphism because η_s preserves the relators $[a_1, t]$ and $a_1^{-1}[a_2, t]$:

$$\eta_s(a_1^{-1}[a_2, t]) = a_1^{-1}[a_2 t^s, t] = a_1^{-1}(a_2 a_1^s)^{-1}(a_2 a_1^{s+1}) = a_1^{-1}a_1 = 1.$$

Because η_s is surjective, it is an automorphism. The image of $H_2(r, 0)$ is $H_2(r, s)$. ☐

The subgroup $H_2(0, 0) = \langle a_1, a_2 \rangle$ is separable in G_2, since $G_2/H_2(0, 0) = \langle t \rangle$ is residually finite. Below we prove $H_2(r, 0)$ is not separable for $r \neq 0$.

Remark 2.4.4. There are no automorphisms ϕ of G_2 such that $\phi(H_2) = H_2(r, 0)$. Therefore the proof of 2.4.1 requires more than an application of Lemma 2.1.4.

Lemma 2.4.5. *The subgoup $H_2(r, 0)$ is free.*

Proof. Suppose that $w = (a_1 t^r)^{\alpha_1} a_2^{\beta_1} \cdots (a_1 t^r)^{\alpha_n} a_2^{\beta_n}$ represents the identity in $H_2(r, 0)$. We show that there is no van Kampen diagram for w is G_2. Rewriting G_2 as $\langle a_1, a_2, t \mid a_1^t = a_1,\ t^{a_2} = t a_1^{-1} \rangle$, van Kampen diagrams over G_2 for w have a_2 corridors. We consider an innermost a_2-corridor, as in Figure 2.2a.

(a) The word $(a_1 t^r)^{\alpha_i} t^{-k} = 1$ only when $\alpha_i = k = 0$

(b) a_2-corridors start and end on the same side of the diagram.

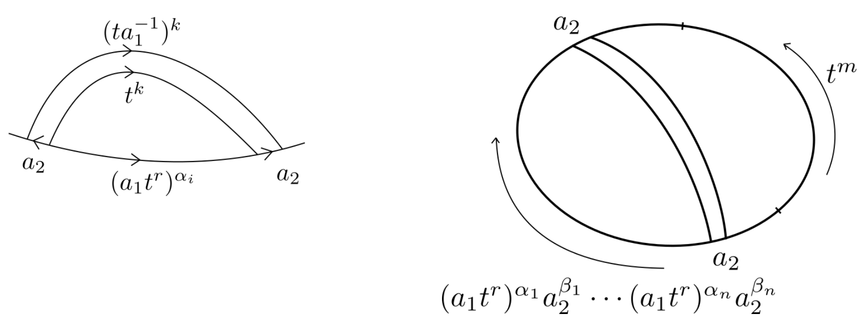

Figure 2.2: Impossible a_2-corridors in the hydra group G_2

The word along the side of the a_2 corridor will either be a power $(t a_1^{-1})^k$ or t^k, call it γ. The word along the boundary, δ, is a power of $(a_1 t^r)$. Then the word

$\delta\gamma^{-1} \in F(a_1, t)$ represents the identity $\langle a_1, t \mid [a_1, t]\rangle$. Such an equality is not possible: either the index sum of a_1 or t in $\delta\gamma^{-1}$ is non-zero. Thus words in $F(a_1 t^r, a_2)$ that contain a_2 do not represent the identity in G_2. Since $(a_1 t^r)$ has infinite order in G_2, there are no relations of the form $(a_1 t^r)^k = 1$. Thus $H_2(r, 0)$ is free. ☐

Lemma 2.4.6. $\langle t\rangle \cap H_2(r, 0) = \{1\}$

Proof. If $\langle t\rangle \cap H_2(r, 0) \neq \{1\}$, for some $\alpha_i, \beta_i \in \mathbb{Z}$ and $m \neq 0$, $w = t^{-m}(a_1 t^r)^{\alpha_1} a_2^{\beta_1} \cdots (a_1 t^r)^{\alpha_n} a_2^{\beta_n}$ represents the identity in G_2. A similar argument as above shows that w can contain no a_2-corridors. Indeed, a_2-corridors start and end on the same side of the diagram, so if there is any a_2 in w, we get an impossible innermost a_2-corridor, as in Figure 2.2a. Therefore, if $t^m \in H_2(r, 0)$, $t^m = (a_1 t^r)^k$ for some k, but clearly this is not the case. ☐

Lemma 2.4.7. *The word* $[t^{-1}, a_2^{-2} t^{-1} a_2^2] \notin H_2(r, 0)$ *when* $r \geq 1$.

Proof. The analogue of Lemma 2.3.3 holds for $H_2(r, 0)$. That is, we can decide whether or not a word $t^s u$ for $u \in \langle a_1, a_2\rangle$ is in $H_2(r, 0)$ by considering whether the rewriting of $t^{p_{i-1}}\pi_i$ as a word in $\langle a_1 t^r, a_2\rangle t^{p_i}$ can be carried out on each piece in the piece decomposition. The word $[t^{-1}, a_2^{-2} t^{-1} a_2^2]$ has piece decomposition:

$$t a_2^{-2} t a_2^2 t^{-1} a_2^{-2} t^{-1} a_2^2 \;=\; (a_2 a_1^{-1})^{-2} (a_2 a_1^{-2})^2 (a_2 a_1^{-1})^{-2} a_2^2 \;=\; a_1 a_2^{-1} | a_1^{-1} | a_2 a_1^{-1} a_2^{-1} | a_1 | a_2$$

where brackets separate the pieces. There is no p such that $a_1 a_2^{-1} \in H_2(r, 0) t^{-p}$. Suppose that there is. Then there is $h \in H_2(r, 0)$ such that $h t^{-p} = a_1 a_2^{-1}$, so

$$h t^{-p} a_2 = a_1 \quad \Rightarrow \quad h a_2 a_1^p t^{-p} = a_1 \quad \Rightarrow \quad h a_2 = a_1^{1-p} t^p \quad \Rightarrow \quad a_1^{1-p} t^p \in H_2(r, 0).$$

We can rewrite $a_1^{1-p} t^p$ as $(a_1 t^r)^{1-p} t^{r(p-1)+p} \in H_2(r, 0)$. Since $a_1 t^r \in H_2(r, 0)$, it follows that $t^{r(p-1)+p} \in H_2(r, 0)$. Lemma 2.4.6 implies that $r(p-1) + p = 0$, so

$p(r+1) = r$. Since $r > 0$, we get $p = \dfrac{r}{1+r}$, which is not an integer. Thus $[t^{-1},\ a_2^{-2}t^{-1}a_2^2] \notin H_2(r, 0)$. $\square$

Proof of Proposition 2.4.1. For $r > 0$, the arguments of Burns–Karrass–Solitar that $H_2(1, 0)$ is not separable in G_2 can be translated directly to show that $H_2(r, 0)$ is not separable in G_2 [32]. For the convenience of the reader, we repeat their argument (almost) verbatim. We drop all decoration and use H to refer to $H_2(r, 0)$ throughout this proof.

Let $\mathcal{T}$ be the group $\langle t_k \mid [t_k, t_{k+1}] = 1,\ k \in \mathbb{Z}\rangle$. To make calculations easier, Burns, Karrass, and Solitar rewrite G_2 as the HNN extension

$$G_2 = \langle \mathcal{T},\ a_2 \mid k \in \mathbb{Z},\ t_k^{a_2} = t_{k+1}\rangle.$$

The generator t_k corresponds to $t^{a_2^k}$ in our original presentation. This implies that $a_1 = [a_2, t] = a_2^{-1}t^{-1}a_2t = t_1^{-1}t_0$, and $[t^{-1}, a_2^{-2}t^{-1}a_2^2] = [t_0^{-1}, t_2^{-1}] \in \mathcal{T}$. In this generating set, $H = \langle t_1^{-1}t_0^{r+1}, a_2\rangle$.

Given an arbitrary finite-index subgroup L satisfying $H < L$, Burns, Karrass, and Solitar find a subgroup $N < L \cap \mathcal{T}$ such that

$$H \cap \mathcal{T} < N \triangleleft \mathcal{T}.$$

Analysis of the quotient $\mathcal{T}/N$ will imply that $[t_0^{-1}, t_2^{-1}]$ is contained in N, and thus in L. According to Lemma 2.4.7, $[t_0^{-1}, t_2^{-1}]$ is not an element of H. Therefore

$$[t_0^{-1}, t_2^{-1}] \ \in \left(\bigcap_{\substack{H \le L < G \\ [G:L]<\infty}} L \right) - H$$

Thus H is not separable.

If L is a finite-index subgroup $L < G$, $\mathrm{core}(L) = \cap_{g\in G} L^g$, is a finite-index normal subgroup. Moreover, $\mathrm{core}(L) \cap \mathcal{T}$ is still normal and finite index in $\mathcal{T}$. The group

N above is given by $(H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$. The majority of the work of this proof is in showing that N is normal in $\mathcal{T}$.

Lemma 2.4.8. $H \cap \mathcal{T} = \langle t_i^{-1} t_{i-1}^{r+1} \mid i \in \mathbb{Z} \rangle$.

Proof. Notice that all elements of $\mathcal{T}$ have trivial a_2 index sum, since every element in the generating set has zero a_2 index sum: $t_i = t^{a_2^i}$. The elements of H with trivial a_2 index sum are all generated by a_2 conjugates of $a_1 t^r$, so

$$H \cap \mathcal{T} \leq \langle (a_1 t^r)^{a_2^{i-1}} \mid i \in \mathbb{Z} \rangle = \langle (t_1^{-1} t_0^{r+1})^{a_2^{i-1}} \mid i \in \mathbb{Z} \rangle = \langle t_i^{-1} t_{i-1}^{r+1} \mid i \in \mathbb{Z} \rangle.$$

For the other direction, $t_i^{-1} t_{i-1}^{r+1} \in \mathcal{T}$, and $t_i^{-1} t_{i-1}^{r+1} = (t_1^{-1} t_0^{r+1})^{a_2^{i-1}} \in H$. □

Claim 2.4.9. $N = (H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$ *is normal in* $\mathcal{T}$.

Since $H \cap \mathcal{T}$ and $\mathrm{core}(L) \cap \mathcal{T}$ are invariant under conjugation by a_2, $N^{a_2} = N$. We next establish that $(H \cap \mathcal{T})^{t_0^{\pm 1}} \subset N$ by considering what happens to the generators $t_i^{-1} t_{i-1}^{r+1}$ under conjugation by $t_0^{\pm 1}$.

Lemma 2.4.10. *If* $i < 0$, *then* $(t_i^{-1} t_{i-1}^{r+1})^{t_0^{\pm 1}} \in H \cap \mathcal{T}$.

Proof. For $i < 0$, both $t_0^{-1} t_i^{(r+1)^i} \in H \cap \mathcal{T}$ and $t_i^{(r+1)^i} t_0^{-1} \in H \cap \mathcal{T}$, as

$$\begin{aligned} t_0^{-1} t_i^{(r+1)^i} &= t_0^{-1} t_{-1}^{(r+1)} (t_{-1}^{-(r+1)} t_{-2}^{(r+1)^2}) \cdots (t_{i+1}^{-(r+1)^{i-1}} t_i^{(r+1)^i}) \\ &= t_0^{-1} t_{-1}^{(r+1)} (t_{-1}^{-1} t_{-2}^{(r+1)})^{(r+1)} \cdots (t_{i+1}^{-1} t_i^{r+1})^{(r+1)^{i-1}}, \end{aligned}$$

where the second equality holds since $[t_k, t_{k+1}] = 1$ for all k. The same kind of rewriting shows $t_i^{(r+1)^i} t_0^{-1} \in H \cap \mathcal{T}$.

Next $(t_i^{-1} t_{i-1}^{r+1})^{t_0}$ can be rewritten using words of the form $t_0^{-1} t_i^{(r+1)^i}$:

$$t_0^{-1} t_i^{-1} t_{i-1}^{r+1} t_0 = t_0^{-1} t_i^{(r+1)^i} (t_i^{-(r+1)^i} t_i^{-1} t_{i-1}^{r+1} t_i^{(r+1)^i}) t_i^{-(r+1)^i} t_0 = (t_0^{-1} t_i^{(r+1)^i})(t_i^{-1} t_{i-1}^{r+1})(t_0^{-1} t_i^{(r+1)^i})^{-1}.$$

Since $t_i^{-1}t_{i-1}^{r+1}$ and $t_0^{-1}t_i^{(r+1)^i}$ are in $H \cap \mathcal{T}$, so is $(t_i^{-1}t_{i-1}^{r+1})^{t_0} \in H \cap \mathcal{T}$. Similarly,

$$t_0 t_i^{-1} t_{i-1}^{r+1} t_0^{-1} = (t_i^{(r+1)^i} t_0^{-1})^{-1}(t_i^{-1}t_{i-1}^{r+1})(t_i^{(r+1)^i}t_0^{-1})^{-1} \in H \cap \mathcal{T}.$$

$\square$

For the other half of the generators, we can only show the following weaker lemma:

Lemma 2.4.11. *If $l > 0$, then $(t_l^{-1}t_{l-1}^{r+1})^{t_0^{\pm 1}} \in (H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$.*

Proof. Because $\mathrm{core}(L) \cap \mathcal{T}$ is finite index in $\mathcal{T}$, there exists $t_i t_j^{-1} \in \mathrm{core}(L) \cap \mathcal{T}$ with $i - j < 0$. Indeed there are infinitely many generators and only finitely many cosets of $\mathrm{core}(L) \cap \mathcal{T}$. Since $\mathrm{core}(L) \cap \mathcal{T}$ is normal, $t_{i-j}t_0^{-1} = (t_i t_j^{-1})^{a_2^{-j}} \in \mathrm{core}(L) \cap \mathcal{T}$. By conjugating $t_{i-j}t_0^{-1}$ by $a_2^{(i-j)k}$ we get $t_{(i-j)(k+1)}t_{(i-j)k}^{-1} \in \mathrm{core}(L) \cap \mathcal{T}$ for all $k \in \mathbb{Z}$. Stringing these elements together, we get that $t_{(i-j)k}t_0^{-1} \in \mathrm{core}(L) \cap \mathcal{T}$ for all $k \in \mathbb{Z}$.

Given $l > 0$, choose $n = (i - j)k$ such that $n > l$. Then

$$t_0 t_l^{-1} t_{l-1}^{r+1} t_0^{-1} = (t_0 t_n^{-1})(t_n t_l^{-1} t_{l-1}^{r+1} t_n^{-1})(t_n t_0^{-1}) = (t_0 t_n^{-1})(t_0 t_{l-n}^{-1} t_{l-1-n}^{r+1} t_0^{-1})^{a_2^n}(t_n t_0^{-1}).$$

Lemma 2.4.10 implies that the middle term is an element of $H \cap \mathcal{T}$, as $l - n < 0$, and $H \cap \mathcal{T}$ is invariant under conjugation by a_2. The conjugating terms $t_n t_0^{-1} \in \mathrm{core}(L) \cap \mathcal{T}$ and so $t_0 t_l^{-1} t_{l-1}^{r+1} t_0^{-1} \in (H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$. $\square$

From Lemmas 2.4.10 and 2.4.11, we have that each of the generators of $H \cap \mathcal{T}$ is conjugated by t_0 and t_0^{-1} into $N = (H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$. From the normality of $\mathrm{core}(L) \cap \mathcal{T}$, we get $((H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T}))^{t_0^{\pm 1}} = N$, and we can conjugate $N^{t_0^{\pm 1}} \subset N$ by a_2^k for $k \in \mathbb{Z}$ to get $N^{t_k^{\pm 1}} \subset N$. Therefore $N \trianglelefteq \mathcal{T}$.

That $[t_0^{-1}, t_2^{-1}]$ is in N follows easily from N being normal in $\mathcal{T}$. Indeed, since $t_0^{-1}t_i^{(r+1)^i} \in H \cap \mathcal{T}$ for $i < 0$, it follows that $t_i^{(r+1)^i}N = t_0N$. In the quotient $\mathcal{T}/N$, the images of t_i and t_0 commute when $i < 0$. When $i > 0$, we can rewrite $[t_0, t_i] = [t_{-i}, t_0]^{a_2^i} = 1^{a_2^i} = 1$, so they too commute in the quotient. Therefore $\mathcal{T}/N$ is abelian and $[t_0^{-1}, t_2^{-1}] \in N$. Since H and $\mathrm{core}(L)$ are subgroups of L and $N = (H \cap \mathcal{T})(\mathrm{core}(L) \cap \mathcal{T})$, it follows that $N \leq L$. Therefore $[t_0^{-1}, t_2^{-1}]$ is an element of L but not of $H_2(r, 0)$, and so $H_2(r, 0)$ is not separable in G_2. $\square$

Consider the group $G_k(\mathbf{w}) = \langle a_1, \ldots, a_k, t \mid a_1^t = a_1w_1, \ldots, a_k^t = a_kw_k \rangle$, where $\mathbf{w} = (w_1, \ldots, w_k)$, with each w_i a positive word on the generators $\{a_1 \ldots a_{i-1}\}$. We now show Theorem 2.1.5: The subgroup H_k is separable in $G_k(\mathbf{w})$ if and only if $\mathbf{w} = (1, \ldots, 1)$.

Proof of Theorem 2.1.5. If $\mathbf{w} = (1, \ldots, 1)$, then $G_k = F_k \times \langle t \rangle$. G_k is subgroup separable, so in particular, H_k is separable. If $\mathbf{w} \neq (1, \ldots, 1)$, then $G_k(\mathbf{w})$ has $\mathbf{w}$ with initial segment of the form $(1, \ldots, 1, w_c, \ldots, w_k)$ where w_c is the first non-trivial word. The subgroup $\langle w_c, a_c, t \rangle$ is isomorphic to G_2 and the subgroup $\langle w_ct^{|w_c|}, a_ct \rangle$ is isomorphic to the subgroup $H_2(|w_c|, 1)$, where $|w_c|$ is the length of w_c in $\langle a_1, \ldots, a_{c-1} \rangle$. Proposition 2.4.1 implies that this subgroup is not separable in $\langle w_c, a_c, t \rangle$. Since $\langle w_ct^{|w_c|}, a_ct \rangle$ is the intersection $H_k(\mathbf{w}) \cap \langle w_c, a_c, t \rangle$, Lemma 1.2.7 implies that H_k is not separable in $G_k(\mathbf{w})$. $\square$

The most general form for which the methods of Dison and Riley apply are $H_k(\mathbf{r}) \leq G_k(\mathbf{w})$. We get only a partial characterization of separability in this case, which is a generalization of Theorem 2.1.5 and its proof.

Recall the statement of Theorem 2.1.6: Suppose $\mathbf{w} = (1, \ldots, 1, w_c, \ldots, w_k)$ and $\mathbf{r} = (r_1, \ldots, r_k)$, with conditions on $\mathbf{w}$, $\mathbf{r}$ as above. Let $[w_c]_i$ denote the index-sum of a_i in w_c. If $\sum_{i=1}^{c-1}[w_c]_i r_i \neq 0$, then $H_k(\mathbf{r})$ is not a separable subgroup of $G_k(\mathbf{w})$.

Proof of Theorem 2.1.6. We examine the subgroup

$$K = \langle a_c, w_c, t \mid w_c^t = w_c,\ a_c^t = a_c w \rangle,$$

which is isomorphic to G_2. Let $\alpha = \sum_{i=1}^{c-1}[w_c]_i r_i \neq 0$. The subgroup given by $J = \langle w_c t^{\alpha},\ a_c t \rangle$ is isomorphic to $H_2(\alpha, 1)$, so by Lemma 2.4.1, it is not separable in K. Since $J \cap H_k(\mathbf{r}) = K$, non-separability of J in K implies non-separability of $H_k(\mathbf{r})$ in $G_k(\mathbf{w})$. $\square$

2.4.1 Further questions

Separability or non-separability of $H_k(\mathbf{r})$ in $G_k(\mathbf{w})$ has not established been for $\mathbf{r}$ and $\mathbf{w}$ such that $\sum_{i=1}^{c}[w_c]_i r_i = 0$. The simplest example that we do not know is $G_3 = \langle a_1, a_2, t \mid a_1^t = a_1, a_2^t = a_2 a_1, a_3^t = a_3 a_2 \rangle$ with subgroup $H_3(\mathbf{r}) = \langle a_1, a_2 t, a_3 \rangle$. It remains to be seen if doubles of these groups are residually finite.

We also wonder what residual properties the hydra group doubles Γ_k and Γ'_k possess. Are they residually solvable, residually amenable, sofic? In the future we hope to explore these questions further.

CHAPTER 3

DERIVED DEPTH FOR RESIDUALLY SOLVABLE GROUPS

In this chapter we examine a depth function which measures the worst case derived length of solvable quotients witnessing non-triviality.

This work grows out of the work of Bou-Rabee and others (see Section 1.5) on quantifying residual properties. We are interested in measuring properties of residually solvable groups that would not require our groups to also be residually finite.

3.1 Basic theory of derived depth

In this section we define the derived depth function and minimal length function. We explain some basic properties of these functions (behavior under subgroups, direct products, wreath products with finite groups) and raise some basic questions that we cannot answer at this time.

Definition 3.1.1. For a finitely generated residually solvable group G given by the presentation $\mathcal{P} = \langle X|R\rangle$, define $dd_{\mathcal{P}} : \mathbb{N} \cup \{0\} \to \mathbb{N} \cup \{0\}$ by

$$dd_{\mathcal{P}}(n) := \min\{i \ |B(1,n) \cap G^{(i)} = \{1\}\},$$

where $B(1,n)$ is the *closed ball* of radius n.

In other words, $dd_{\mathcal{P}}(n)$ is the first derived subgroup such that $B(1,n) \cap G^{(dd_{\mathcal{P}}(n))} = \{1\}$. Every element of length less than or equal to n can be distinguished from the identity in a solvable quotient of derived length $dd_{\mathcal{P}}(n)$.

We record some basic properties of the function $dd_{\mathcal{P}}(n)$.

Proposition 3.1.2. *Basic properties of $dd_{\mathcal{P}}(n)$:*

1. *G is a non-decreasing function.*
2. *If G is solvable, with derived length k, the image of $dd_{\mathcal{P}}$ is a subset of $[0, k]$.*
3. *If G is not solvable, the image of $dd_{\mathcal{P}}$ is unbounded.*

Proof. If $B(1, n) \cap G^{(i-1)} \neq \{1\}$, then $B(1, n+1) \cap G^{(i-1)} \neq \{1\}$. Therefore $dd_{\mathcal{P}}(n) \leq dd_{\mathcal{P}}(n+1)$. The other claims follow from the definition of solvability. □

We impose the relation that $f \preceq g$ if there is $C > 0$ such that $f(n) \leq Cg(Cn)$ for all n. All non-trivial solvable groups have equivalent derived depth under this relation, since non-zero bounded functions are equivalent under this relation.

In order to calculate $dd_{\mathcal{P}}(n)$, it is useful to have the following function, which measures the length of the shortest non-trivial element in $G^{(i)}$:

Definition 3.1.3. For a finitely generated group G given by the presentation $\mathcal{P} = \langle X \mid R \rangle$, the minimum length function, $\Phi_{\mathcal{P}} : \mathbb{N} \to \mathbb{N} \cup \{0\}$, is defined by $\Phi_{\mathcal{P}}(i) := 0$ when $G^{(i)} = \{1\}$. Otherwise, $G^{(i)} \neq \{1\}$ and

$$\Phi_{\mathcal{P}}(i) := \min\{n \text{ such that } B(1, n) \cap G^{(i)} \neq \{1\}\}.$$

We put the same equivalence relation on $\Phi_{\mathcal{P}}$ as we put on $dd_{\mathcal{P}}$.

Proposition 3.1.4. *If G is residually solvable but not solvable, the function $\Phi_{\mathcal{P}}$ is non-decreasing.*

Proof. If $B(1,n) \cap G^{(i)} \neq 1$, then since $G^{(i)} \leq G^{(i-1)}$, $B(1,n) \cap G^{(i-1)} \neq 1$. Call a nontrivial element of $G^{(i)}$ which has length $\Phi_{\mathcal{P}}(i)$ a *minimum length* or shortest element of $G^{(i)}$. □

Example 3.1.5. Suppose that $G = F_2$. We list a few values of the functions dd_{F_2} and Φ_{F_2} to illustrate the definitions.

$$dd_{F_2}(0) = 0, \quad dd_{F_2}(1) = 1, \quad dd_{F_2}(2) = 1, \quad dd_{F_2}(3) = 1$$

$$dd_{F_2}(4) = 2, \ldots, \quad dd_{F_2}(13) = 2, \quad dd_{F_2}(14) = 3, \ldots$$

$$\Phi(0) = 1, \quad \Phi(1) = 4, \quad \Phi(2) = 14, \ldots$$

Lemma 3.1.6. *If f, g are increasing functions and $f \preceq g$ then $g^{-1} \preceq f^{-1}$.*

Proof. This is an easy check: $f(n) \preceq g(n)$ means that $f(n) \leq h(n) = Cg(Cn)$ for some constant $C > 0$. Therefore $h^{-1}(n) \leq f^{-1}(n)$, and $h^{-1}(n) = (1/C)g^{-1}(n/C)$. Thus

$$(1/C)g^{-1}(n/C) \leq f^{-1}(n) \quad \Leftrightarrow \quad g^{-1} \preceq f^{-1}.$$

□

Lemma 3.1.7. *If f_1 and f_2 are increasing functions and $f_1 \preceq \Phi_{\mathcal{P}} \prec f_2$, then $f_2^{-1} \prec dd_{\mathcal{P}} \preceq f_1^{-1}$.*

Proof. Suppose $f_1(k) \leq \Phi_{\mathcal{P}}(k)$ for all k. If $\Phi_{\mathcal{P}}(k) \geq f_1(k)$ then the shortest element of $G^{(k)}$ has length greater or equal to $f_1(k)$. Therefore $B(1, f_1(k) - 1) \cap G^{(k)} = \{1\}$. By definition, $dd_{\mathcal{P}}(f_1(k) - 1) \leq k$, so pre-composing with f^{-1} yields

$$dd_{\mathcal{P}}(n-1) \leq f_1^{-1}(n).$$

If $\Phi_{\mathcal{P}}(k) < f_2(k)$, the shortest element of $G^{(k)}$ has length less than $f_2(k)$. Therefore $B(1, f_2(k)) \cap G^{(k)} \neq \{1\}$. By definition, $dd_{\mathcal{P}}(f_2(k)) > k$, so

$$dd_{\mathcal{P}}(n) > f_2^{-1}(n).$$

□

Therefore, to find upper bounds for the derived depth function $dd_{\mathcal{P}}$, we want lower bounds for $\Phi_{\mathcal{P}}$ and vice versa.

Proposition 3.1.8. *If G has finite generating set S and the subgroup H has finite generating set T, then $dd_T \preceq dd_S$.*

Proof. Assume that G is not solvable. Every generator in T can be rewritten in terms of the generators of S. Since T is finite, we can define $C > 0$ to be the maximum of $|t|_S$ where $t \in T$. Let $w \in H^{(i)}$ be a shortest element: $|w|_T = \Phi_T(i)$. Then $|w|_S \leq C|w|_T$. Since $H^{(i)} \subset G^{(i)}$, $|w|_S$ is an upper bound on the length of the shortest element of $G^{(i)}$. Thus $\Phi_S(n) \leq C\Phi_T(n)$ for all n, so $\Phi_S \preceq \Phi_T$ and $dd_T(n) \preceq dd_S(n)$. □

This implies that changing finite generating sets does not change (up to equivalence), the derived depth of the group. Thus we may relax our notation to dd_G and Φ_G. A partial converse is the following:

Proposition 3.1.9. *Suppose G is residually solvable and $H \leq G$. If $G/\mathrm{core}(H)$ is solvable, then $dd_G(n) \preceq dd_H(\mathrm{Dist}_H^G(n))$.*

Proof. Let k be the derived length of the solvable quotient $G/\mathrm{core}(H)$. Since $G/\mathrm{core}(H)$ is solvable, $G^{(k)} \leq H$, and so for all i, $G^{(k+i)} \leq H^{(i)}$. Let $g \in G^{(i+k)}$ be shortest, so $|g|_G = \Phi_G(i+k)$. Since $G^{(i+k)} \leq H^{(i)}$, there is $h \in H^{(i)}$ such that

$|h|_G \leq |g|_G$. Therefore, $|h|_H \leq \mathrm{Dist}_H^G(|g|_G)$ and so $\Phi_H(i) \leq \mathrm{Dist}_H^G(\Phi_G(i+k))$ for all i. Therefore $dd_G(n) \leq dd_H(\mathrm{Dist}_H^G(n)) + k$. □

3.1.1 Derived depth and group-building operations

Lemma 3.1.10. *If A and B are residually solvable groups with derived depth functions of dd_A and dd_B, then $A \times B$ is also residually solvable, and has $dd_{A\times B}(n) \asymp \max\{dd_A(n), dd_B(n)\}$*

Proof. We only need to show that $dd_{A\times B} \preceq \max\{dd_A(n), dd_B(n)\}$.

$$(A \times B)^{(i)} \cong A^{(i)} \times B^{(i)}$$

If $B_A(1,n) \cap A^{(i)} = \{1\}$ and $B_B(1,n) \cap B^{(i)} = \{1\}$, then $B_{A\times B}(1,n) \cap \left(A^{(i)} \times B^{(i)}\right) = \{1\}$. Therefore $dd_{A\times B}(n) \leq \max\{dd_A(n), dd_B(n)\}$. □

Therefore, if G is residually solvable, $B = \oplus_{i=1}^n G$ and G have equivalent derived depth functions: $dd_G(n) \asymp dd_B(n)$.

Definition 3.1.11. The *restricted wreath product* is

$$G \text{ wr } H := \left(\bigoplus_{h\in H} G_h\right) \rtimes H,$$

where the action of H on $B = \oplus_{h\in H} G_h$ is defined as follows:

$$((a_h)_{h\in H},\ h_1) \cdot ((b_h)_{h\in H},\ h_2) = ((a_h b_{h_1^{-1}h}),\ h_1 h_2).$$

Proposition 3.1.12. *If G wr H is residually solvable and finitely presentable, then $dd_G(n) \asymp dd_{G\mathrm{wr}H}(n)$.*

Proof. If $G \neq \{1\}$, the group G wr H is finitely presentable if and only if H is finite. Since $G \hookrightarrow B \leq G$ wr H, if G wr H is residually solvable, then G is residually solvable. The subgroup H also embeds in G wr H: $\{((1)_{h\in H},\ h)\}$ is isomorphic to H, so H is residually solvable. Residual solvability and solvability are equivalent for finite groups, so H is solvable. From the short exact sequence

$$1 \to B \hookrightarrow G \text{ wr } H \to H \to 1$$

we see that B is a finite index subgroup such that $(G \text{ wr } H)/B \cong H$ is solvable. By Corollary 3.1.15, $dd_B(n)$ and $dd_{G\text{wr}H}(n)$ are equivalent. Since B is a direct sum of finitely many copies of G, $dd_G(n)$, $dd_B(n)$, and thus $dd_{G\text{wr}H}(n)$ are equivalent. □

An unfulfilled goal of this project was to understand how derived depth functions behave under free products. In Section 3.2, we calculate the derived depth of the free product of two abelian groups.

3.1.2 Derived depth and finite index subgroups

We know that if $H \leq G$ is residually solvable and finite index in G that G need not be residually solvable, but if G is also residually solvable, we wonder how the derived depth functions of G and H compare.

Open question 3.1.13. If G is residually solvable and finitely generated and H is a non-trivial finite index subgroup, when is $dd_H \asymp dd_G$?

We present a few partial answers to this question.

Lemma 3.1.14 (Lubotzky–Mann [71])**.** *If G is residually solvable and virtually solvable, then G is solvable. That is, if G is residually solvable and H is a finite index subgroup with dd_H equivalent to a constant function, then $dd_G \asymp dd_H$.*

Proof. If G is virtually solvable, there is $H \triangleleft G$ that is finite index and solvable. Let $l = [G : H]$ and m be the derived length of H. Consider $p : G \to Q$, for Q solvable. The image of H in Q, $p(H)$, is normal. The quotient $Q/p(H)$ is a solvable group with at most l elements. $Q/p(H)$ has derived length at most m, since each derived subgroup is strictly contained in its predecessor. By Lemma 1.3.3, Q has derived length at most $l + m$. Since every solvable quotient of G has derived length no more than $l + m$ and G is residually solvable, $G^{(l+m)} = \{1\}$, and so G is solvable. □

Corollary 3.1.15 (of Proposition 3.1.9)**.** *Let G be a residually solvable finitely generated group. If $G/\mathrm{core}(H)$ is a finite, solvable quotient, then H and G have equivalent derived depth functions.*

Proof. Proposition 3.1.9 implies that $dd_G(n) \preceq dd_H(\mathrm{Dist}_H^G(n)) \leq dd_H(Cn)$, so G and H have equivalent derived depth functions. □

3.2 Derived depth for some examples

We present Elkasapy and Thom's proof of bounds on Φ_{F_2}, which implies that free groups have logarithmic derived depth function. We use this to show that free-by-cyclic groups always have a linear upper bound on their derived depth functions. Applying their result, we show that the free product of two abelian groups has derived depth function equivalent to the derived depth function of F_2.

The normal residual finiteness growth of the free groups is not yet known, although many have worked on it [22, 47, 67]. Results of Elkasapy and Thom in this search settle the question of the derived depth of the free groups [47].

Theorem 3.2.1 (Elkasapy–Thom)**.** *The free group F_n satisfies that* $3^i \leq \Phi(i) \leq 4^i$.

Proof. The upper bound for $\Phi(i)$ is easy to see by induction. Take two of the shortest words in $F_n^{(i)}$. One will be w, a witness to the value of $\Phi(i)$, and the other will be v, a cyclic permutation of w. The word $[w, v]$ is a non-trivial element of $F_n^{(i+1)}$. The length of $[w, v]$ is less than $4|w| = 4\Phi(i)$. Therefore $\Phi^{(i+1)} \leq 4^{(i+1)}$.

To establish the lower bound, Elkasapy and Thom show that $3\Phi_{F_n}(i) \leq \Phi_{F_n}(i + 1)$. Their argument exploits the fact that subgroups of F_n always have Nielsen reduced bases. That is, for a subgroup H, there is a generating set S for H such that for all $s_1, s_2, s_3 \in S^{\pm 1}$, their lengths (in F_2) satisfy

$$|s_1 s_2| \geq \max(|s_1|, |s_2|) \text{ and } |s_1 s_2 s_3| > |s_1| + |s_3| - |s_2|.$$

Any freely reduced word in S when written as a word in F_n will probably not be freely reduced, but none of the generators in S (written as words in F_2) are completely canceled away during free reduction in F_n. Therefore if $w = s_{i_1}^{\epsilon_1} \dots s_{i_m}^{\epsilon_m}$ for $s_{i_k} \in S$ and $\epsilon_k \in \{\pm 1\}$, is not freely reducible in $F(S)$, then $|w| \geq m$. Further, $|w| \geq \max\{|s_{i_k}| \text{ for } 1 \leq k \leq m\}$ [72].

Let S be a Nielsen reduced basis for $F_n^{(i)}$. Let $w \in F_n^{(i+1)} = (F_n^{(i)})^{(1)}$, with $w = s_{i_1}^{\epsilon_1} \dots s_{i_m}^{\epsilon_m}$ for $s_{i_k} \in S$ and $\epsilon_k \in \{\pm 1\}$. Since w is in the derived subgroup, all generators of S in w have index sum zero, so in particular, it is possible to choose a pair of distinct generators from S to break up the word. Choose the distinguished generators t and s so that $t^{-1}s$ has the largest cancellation in F_2 of the pairs of s_{i_j} appearing in w. Here we present Elkasapy and Thom's proof for the case that $w = sw_1tw_2s^{-1}w_3t^{-1}$ (which is freely reduced in the generators S). Call the canceling piece a, so that $t = at_1$ and $s = as_1$, and $|t_1^{-1}s_1| = |t_1| + |s_1|$. Each of the elements $sw_1, tw_2s^{-1}, w_3t^{-1}$ are elements of $F_2^{(i)}$. To estimate the length of $|w|$, we want to sum $|sw_1|, |tw_2s^{-1}|$ and $|w_3t^{-1}|$, but there may be some cancellation where these words come together. The cancellation in F_n between sw_1 and tw_2s^{-1} happens on the prefix b of a: $a = b\beta$. Similarly, cancellation in F_n between tw_2s^{-1} and w_3t^{-1} happens on a suffix of s^{-1}, in particular, on c^{-1}, a suffix of a^{-1}, so $a^{-1} = \alpha c^{-1}$. Then $|w| \geq |sw_1| + |tw_2s^{-1}| + |w_3t^{-1}| - 2|b| - 2|c|$. Unless b is trivial, sw_1 is not cyclically reduced: it has prefix b and suffix b^{-1}. Since $F_n^{(i)}$ is normal, all conjugates of sw_1 are also elements of $F_2^{(i)}$. Therefore the cyclically reduced conjugate of sw_1 in $F_n^{(i)}$ has length greater than or equal to $|sw_1| - 2|b|$, so $\Phi_{F_2}(i) \leq |sw_1| - 2|b|$. Similarly, the cyclically reduced conjugates of w_3t^{-1} and tw_2s^{-1} in $F_2^{(i)}$ have length greater than or equal to $|w_3t^{-1}| - 2|c|$, and $|tw_2s^{-1}| - 2a$. Altogether then

$$3\Phi_{F_n}(i) \leq |sw_1| - 2|b| + |tw_2s^{-1}| - 2|a| + |w_3t^{-1}| - 2|c| \leq |w|.$$

A similar calculation shows that for w arbitrary, this inequality holds, and so $3\Phi_{F_n}(i) \leq \Phi_{F_n}(i+1)$. □

Corollary 3.2.2. *The free group on infinitely many generators has a logarithmic derived depth function.*

Proof. Suppose w is a minimum length word of $F_\infty^{(i)}$. It is a word on only finitely many generators, so we can consider it as a word on just say, n generators. This implies that the minimum word length for the ith derived subgroup stabilizes at rank n. Therefore the bounds from Proposition 3.2.1 imply that $|w| \geq 3^i$ and so we get an exponential lower bound for the minimum length.

Therefore, the group F_∞ satisfies the inequalities of Proposition 3.2.1. Taking inverses yields

$$\log_4(n) \leq dd_{\mathcal{P}}(n) \leq \log_3(n).$$

The upper and lower bounds are equivalent under the equivalence relation $f \preceq g$ if $f(n) \leq Cg(Cn)$ for all n. $\square$

Corollary 3.2.3. *If $G = F_n \rtimes_\phi \mathbb{Z}$, then $dd_G(n) \preceq n$. If $\mathrm{Dist}_{F_n}^G(k) \asymp P(k)$ for P a polynomial, then $dd_G(n) \asymp \log(n)$.*

Proof. By Proposition 3.1.9, $dd_G(k) \preceq dd_N(\mathrm{Dist}_N^G(k))$. If $N \triangleleft G$ are finitely generated groups, then $\mathrm{Dist}_N^G(k) \preceq 2^k$ [26]. Therefore

$$dd_G(k) \preceq dd_{F_n}(\mathrm{Dist}_{F_n}^G(k)) \preceq dd_{F_n}(2^k) = \log(2^k) = k.$$

$\square$

We now work up to showing that if A and B are abelian groups that $A * B$ has derived depth function equivalent to that of F_2 (as one would probably guess).

Lemma 3.2.4. *Some common commutator identities that we will find useful are the following:*

1. $[x, y]^{-1} = [y, x]$
2. $x^y = y^{-1}xy = x[x, y]$
3. $[x, yz] = x^{-1}z^{-1}y^{-1}xyz = x^{-1}z^{-1}xzz^{-1}x^{-1}y^{-1}xyz = [x, z][x, y]^z$

Let $K = (A * B)^{(1)}$. If A and B are abelian, abelianizing $A * B$ yields:

$$1 \to K \to A * B \to A \times B \to 1.$$

Observe that $K \cap A = K \cap B = \{1\}$, so by Lemma 1.3.6, K is free.

Definition 3.2.5. An element g of $A * B$ can be written as a product of elements of A and B: $a_1 b_1 \ldots a_n b_n$ with a_1, b_n possibly trivial. The *free product length* of g is the number of nontrivial A and B terms in this product.

Lemma 3.2.6. *For A, B abelian, $X = \{[a, b] \mid a \in A - \{1\}, b \in B - \{1\}\}$ is a generating set for $(A * B)^{(1)}$. Moreover, this is a free generating set for $(A * B)^{(1)}$.*

Proof. It is clear that $X \subset (A * B)^{(1)}$. To show that X is a generating set, we show that all elements of the the generating set $Y = \{[u, v] \mid u, v \in A * B\}$ can be written as words in X. Words in $A * B$ have the form $u = a_1 b_1 \cdots a_n b_n$ where $a_i \in A$ and $b_i \in B$, where only a_1 or b_n can be trivial. Letting $v = \alpha_1 \beta_1 \cdots \alpha_m \beta_m$, we induct on $n + m$, starting at $n + m = 2$:

$$[u, v] = b_1^{-1} a_1^{-1} \beta_1^{-1} \alpha_1^{-1} a_1 b_1 \alpha_1 \beta_1 = [b_1, a_1][a_1, \beta_1 b_1][\beta_1 b_1, \alpha_1][\alpha_1, \beta_1],$$

which is checked by expanding — recall that A and B are abelian.

Suppose $n + m = k$, and all elements of Y with $n + m < k$ can be written as products of elements in $X \cup X^{-1}$. If $u = a_1 b_1 \cdots a_n b_n$, and $v = \alpha_1 \beta_1 \cdots \alpha_m \beta_m$, let $v = v' \alpha_n \beta_n$. Then

$$[u, v] = [u, v' \alpha_n \beta_n] = [u, \alpha_n \beta_n][u, v']^{\alpha_n \beta_n}$$

By induction, both $[u, \alpha_n \beta_n]$ and $[u, v']$ can be written as products of elements of $X \cup X^{-1}$. All that remains to show is that for all $a \in A$ and $b \in B$, $[a, b]^{\alpha_n \beta_n}$ can be written as such a product. Indeed,

$$[a, b]^{\alpha_n \beta_n} = \beta_n^{-1} \alpha_n^{-1} a^{-1} b^{-1} a b \alpha_n \beta_n = [\beta_n, a\alpha_n][a\alpha_n, b\beta_n][b\beta_n, \alpha_n][\alpha_n, \beta_n].$$

Therefore X is a generating set for $(A * B)^{(1)}$.

To show that X is a free generating set, we show that a freely reduced word $w \in F(X)$ with length n in $F(X)$ represents an element with free-product length at least n in $A * B$. Moreover, for any product of generators in $X \cup X^{-1}$ of the form $[a_1, b_1][b_2, a_2] \cdots [a_n, b_n]$, the element it represents has prefix $a_1^{-1}b_1^{-1}$ and suffix a_nb_n. This is clear for the case $n = 1$. Inducting on n, let $w = [a_1, b_1] \cdots [a_{n-1}, b_{n-1}][b_n, a_n] = w'[b_n, a_n]$. Then w' represents a word in $(A * B)^{(1)}$ that has free product length at least $m \geq n - 1$, and w' represents $a_1^{-1}b_1^{-1}A_3B_3 \ldots B_{m-2}a_{n-1}b_{n-1}$. Therefore in $A * B$, w is equal to

$$a_1^{-1}b_1^{-1}A_3B_3 \ldots B_{m-2}a_{n-1}b_{n-1}b_n^{-1}a_n^{-1}b_na_n$$

If $b_{n-1}b_n^{-1} \neq 1$, then we add 3 to the free product length of the word, and the word begins and ends as promised. If there is cancellation and $b_{n-1}b_n^{-1} = 1$, then $a_{n-1}a_n^{-1} \neq 1$. Otherwise w ended in $[a_{n-1}, b_{n-1}][b_{n-1}, a_{n-1}]$ and was not reduced. Letting $A_{m-1} = a_{n-1}a_n^{-1}$, we see that the free product length of the word is at least $m + 1$ and so is greater than or equal to n. Moreover, the word begins with $a_1^{-1}b_1^{-1}$ and ends with b_na_n, as promised. This is a free generating set, as we have shown that no freely reduced elements of $F(X)$ represent the identity in $A * B$. $\square$

Lemma 3.2.6 gave lower bounds on the free-product lengths of products of elements of X. In Lemma 3.2.7, we rewrite elements of $(A * B)^{(1)}$ written in the standard form of $A * B$, into words in $X \cup X^{-1}$.

Lemma 3.2.7. *If $w = a_1b_1 \cdots a_nb_n \in (A * B)^{(1)}$, then*

$$w = [a_1^{-1},\ b_1^{-1}] \prod_{i=2}^{n-1} [(b_1 \cdots b_{i-1})^{-1},\ (a_1 \cdots a_i)^{-1}][(a_1 \cdots a_i)^{-1},\ (b_1 \cdots b_i)^{-1}].$$

Some of these terms may be trivial (if $a_1 \cdots a_k = 1$ or $b_1 \cdots b_k = 1$ for some k), but there are never more than 3 terms in a row that are equal to 1, and after removing these terms, the word is freely reduced in $F(X)$.

Proof. We induct on n. If $n = 2$, then $w = a_1 b_1 a_2 b_2 \in (A * B)^{(1)}$ if and only if $w = [a_1^{-1},\ b_1^{-1}]$. For the induction, suppose $w = a_1 b_1 \cdots a_n b_n$ and let w' be the suffix $a_3 b_3 \cdots a_n b_n$. Then

$$w = a_1 b_1 a_2 b_2 w' = [a_1^{-1}, b_1^{-1}][b_1^{-1}, (a_1 a_2^{-1})] a_1 a_2 b_1 b_2 w'.$$

Since $[a_1^{-1}, b_1^{-1}][b_1^{-1}, (a_1 a_2^{-1})] \in (A * B)^{(1)}$, then so is $a_1 a_2 b_1 b_2 w'$. By hypothesis, the word $a_1 a_2 b_1 b_2 w'$ has the form

$$[(a_1 a_2)^{-1}, (b_1 b_2)^{-1}] \prod_{i=3}^{n-1} [(b_1 b_2 \cdots b_{i-1})^{-1},\ (a_1 a_2 \cdots a_i)^{-1}][(a_1 a_2 \cdots a_i)^{-1},\ (b_1 b_2 \cdots b_i)^{-1}].$$

Putting these together we get the desired form.

Let us abbreviate $a_1 \cdots a_i = A_i^{-1}$, and $b_1 \cdots b_i = B_i^{-1}$. Our word has the form $[A_1, B_1][B_1, A_2] \cdots [A_{n-1}, B_{n-1}]$. If $A_i = 1$, then the two terms $[B_{i-1}, A_i]$ and $[A_i, B_i]$ are trivial. Since $a_j \neq 1$ for any j, and $a_i^{-1} A_{i-1} = A_i$ and $a_{i+1}^{-1} A_i = A_{i+1}$, if $A_i = 0$, this implies that $A_{i-1} \neq 1$ and $A_{i+1} \neq 1$. This is illustrated in Figure 3.1. We read the word off by following the arrows. An arrow starting at v_s and ending at v_e represents $[v_s, v_e]$. The red arrows and vertices represent being trivial. Blue represents forced non-triviality. Suppose $B_{i-1}, B_i \neq 1$, as in Figure 3.1a.

(a) A cartoon of $A_i = 1$, but $B_{i-1}, B_i \neq 1$.

(b) A cartoon of $A_i = 1$ and $B_i = 1$.

Figure 3.1: Collapse in rewriting $w \in (A * B)^{(1)}$ as a word in $F(X)$.

Simplifying, we get $[A_{i-1}, B_{i-1}][B_{i-1}, A_i][A_i, B_i][B_i, A_{i+1}] = [A_{i-1}, B_{i-1}][B_i, A_{i+1}]$.

Since $b_i^{-1}B_{i-1} = B_i$, and in particular $B_{i-1} \neq B_i$, there is no further collapse. Therefore, to have more than two terms in a row that are equal to 1, either $B_{i-1} = 1$ or $B_i = 1$. We consider the second case, illustrated in Figure 3.1b. The term $[A_i, B_i] = 1$, and $[B_i, A_{i+1}] = 1$, but $[A_{i+1}, B_{i+1}] \neq 1$, since neither $A_{i+1}, B_{i+1} \neq 1$. Then

$$[A_{i-1}, B_{i-1}][B_{i-1}, A_i][A_i, B_i][B_i, A_{i+1}][A_{i+1}, B_{i+1}] = [A_{i-1}, B_{i-1}][A_{i+1}, B_{i+1}],$$

and there can be no further collapse here. □

Corollary 3.2.8. *If $w \in (A * B)^{(1)}$ has free product length $2n$ and $v \in F(X)$ represents w, then*

$$\frac{n}{2} - 1 \leq |v|_X \leq 2n.$$

If $w \in (A * B)^{(1)}$, and $v \in F(X)$ represents w, if $w = a_1 b_1 \cdots a_n b_n$, then the free product length of w satisfies:

$$|v|_X \leq 2n \leq 4|v|_X + 4.$$

Proof. Lemma 3.2.7 implies $w = a_1 b_1 \cdots a_n b_n$ is a product of $2n - 3$ elements, with some possibly trivial: $[A_1, B_1][B_1, A_2][A_2, B_2] \cdots [A_{n-1}, B_{n-1}]$. The second part of the lemma showed that there are never more than 3 trivial commutators in a row. Therefore $|v|_X \geq (2n - 3)/4$. The second set of inequalities follows immediately. □

The group $(\mathbb{Z}/2\mathbb{Z}) * (\mathbb{Z}/2\mathbb{Z})$ is solvable, so it has a constant derived depth function. Otherwise, we get logarithmic derived depth functions.

Corollary 3.2.9. *If A and B are abelian, and not both isomorphic to $\mathbb{Z}/2\mathbb{Z}$ then $A * B$ has logarithmic derived depth function*

Proof. Let $Y = \{a_1, \ldots, a_l\} \cup \{b_1 \ldots, b_m\}$ be a generating set for $A * B$, coming from generating sets for A and B. Choose $g \in (A * B)^{(m)}$ such that g realizes $\Phi_Y(m)$. Let K be the free-product length of g and let $v \in F(X)$ represent g. Since $g \in (A * B)^{(m)}$, $v \in ((A * B)^{(1)})^{(m-1)}$, and by Corollary 3.2.2, $3^{(m-1)} \leq |v|_X$. By Corollary 3.2.8, $|v|_X \leq K \leq |g|_Y$. Therefore $3^{m-1} \leq |g|_Y = \Phi_Y(m)$ and $dd_Y(n) \preceq \log(n)$. If $A * B \ncong \mathbb{Z}_2 * \mathbb{Z}_2$, then $A * B$ has a free subgroup, producing a logarithmic lower bound on $dd_{A*B}(n)$. Therefore $dd_{A*B}(n) \asymp \log(n)$. $\square$

These techniques may extend to establish the derived depth for free products of solvable groups, but this calculation has not yet been done. Although we anticipate a result of the form $A * B$ has derived depth equivalent to $M(n) = \max\{\log(n), dd_A(n), dd_B(n)\}$, it is not clear if the technique developed here can be extended.

3.3 Further questions

Many questions remain unanswered:

1. If G is a residually solvable group with finite index normal subgroup H such that G/H is non-solvable, how do $dd_G(n)$ and $dd_H(n)$ compare?
2. Is the derived depth function a quasi-isometry invariant in the class of residually solvable groups? That is, if two residually solvable groups are q.i. do they have equivalent derived depth functions?
3. What kinds of derived depth functions occur? We presented here several classes of groups with logarithmic derived depth. We have not yet produced an example with a non-equivalent derived depth function.

4. How do the derived depth functions of groups of the form $A * B$ and $A *_C B$ behave? How do they depend on A, B, and C?
5. What are the derived depth functions for the Baumslag-Solitar groups? None of the strategies we used above are obviously applicable. What are the derived depth functions for wreath products that are not finitely presentable?
6. In many of the examples of G we have considered, for i large enough, $G^{(i)}$ is a free group. It would be helpful to have a more diverse family of examples to draw upon.

BIBLIOGRAPHY

[1] I. Agol. The Virtual Haken Conjecture. *Documenta mathematica*, 18, 2013.

[2] J. Alonso. Ingalits isoprimtriques et quasi-isomtries. *Comptes Rendus de l'Acadmie des Sciences*, 311(12):761–764, 1990.

[3] G. Arzhantseva. Asymptotic approximations of finitely generated groups. In J. González-Meneses, M. Lustig, and E. Ventura, editors, *Extended Abstracts Fall 2012: Automorphisms of Free Groups*, pages 7–15. Springer International Publishing, Cham, 2014.

[4] A. Ash. Small-dimensional classifying spaces for arithmetic subgroups of general linear groups. *Duke Mathematical Journal*, 51(2), 1984.

[5] A. Baudisch. Subgroups of semifree groups. *Acta Mathematica Academiae Scientiarum Hungaricae*, pages 19–28, 1981.

[6] B. Baumslag and M. Tretkoff. Residually finite HNN extensions. *Communications in Algebra*, 6(2):179–194, 1978.

[7] G. Baumslag. On generalised free products. *Mathematische Zeitschrift*, pages 423–438, 1962.

[8] G. Baumslag. Finitely generated cyclic extensions of free groups are residually finite. *Bulletin of the Australian Mathematical Society*, 5:87–94, 1971.

[9] G. Baumslag. Positive one-relator groups. *Transactions of the American Mathematical Society*, 156, 1971.

[10] F. Berlai. Residual properties of free products. *Communications in Algebra*, 44:2959–2980, 2016.

[11] M. Bestvina and N. Brady. Morse theory and finiteness properties of groups. *Inventiones mathematicae*, 129(3):445–470, 1997.

[12] M. Bestvina and M. Feighn. A combination theorem for negatively curved groups. *J. Differential Geom.*, 35(1):85–101, 1992.

[13] M. Bestvina and M. Feighn. Addendum and correction to: “A combination theorem for negatively curved groups”. *J. Differential Geom.*, 43(4):783–788, 1996.

[14] M. Bestvina, M. Feighn, and M. Handel. The Tits Alternative for Out(F_n) I: Dynamics of exponentially-growing automorphisms. *Annals of Mathematics*, 151(2):517–623, 2000.

[15] M. Bestvina, M. Feighn, and M. Handel. The Tits Alternative for Out(F_n) II: A Kolchin type theorem. *Annals of Mathematics*, 161(1):1–59, 2005.

[16] O. Bogopolski, A. Martino, and E. Ventura. The automorphism group of a free-by-cyclic group in rank 2. *Communications in Algebra*, 35(5):1675–1690, 2007.

[17] A. Borisov and M. Sapir. Polynomial maps over finite fields and residual finiteness of mapping tori of group endomorphisms. *Inventiones mathematicae*, 160(2):341–356, 2005.

[18] A. Borisov and M. Sapir. Polynomial maps over p-adics and residual properties of mapping tori of group endomorphisms. *International Mathematics Research Notices*, 2009(16):3002–3015, 2009.

[19] K. Bou-Rabee. Quantifying residual finiteness. *Journal of Algebra*, 323(3):729 – 737, 2010.

[20] K. Bou-Rabee. Approximating a group by its solvable quotients. *New York Journal of Mathematics*, 17:699–712, 2011.

[21] K. Bou-Rabee, M.F. Hagen, and P. Patel. Residual finiteness growths of virtually special groups. *Mathematische Zeitschrift*, 279(1):297–310, 2015.

[22] K. Bou-Rabee and D. B. McReynolds. Asymptotic growth and least common multiples in groups. *Bulletin of the London Mathematical Society*, 43(6):1059–1068, 2011.

[23] B. H. Bowditch. Relatively hyperbolic groups. *International Journal of Algebra and Computation*, 22(03):1250016, 2012.

[24] N. Brady, M. R Bridson, M. Forester, and K. Shankar. Snowflake groups, Perron-Frobenius eigenvalues and isoperimetric spectra. *Geometry and Topology*, 13:141 – 187, 2009.

[25] M. Bridson. On the geometry of normal forms in discrete groups. *Proceedings of the London Mathematical Society*, 67:596–616, 1993.

[26] M. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*. Number 319 in Grundlehren der mathematischen Wissenschaften. Springer-Verlag Berlin Heidelberg, 2 edition, 2009.

[27] M. R. Bridson. Fractional isoperimetric inequalities and subgroup distortion. *Journal of the American Mathematical Society*, 12(4):1103–1118, 1999.

[28] M.R. Bridson and S.M. Gersten. The optimal isoperimetric inequality for torus bundles over the circle. *The Quarterly Journal of Mathematics. Oxford*, 47(185):1–23, 1996.

[29] M.R. Bridson and D. Groves. *The quadratic isoperimetric inequality for mapping tori of free group automorphisms*, volume 203 of *Memoirs of the American Mathematical Society*. American Mathematical Society, 2010.

[30] M.R. Bridson and C. Pittet. Isoperimetric inequalities for the fundamental groups of torus bundles over the circle. *Geometriae Dedicata*, 49(2):203–219, 1994.

[31] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geometric and Functional Analysis*, 10:1071–1089, 2000.

[32] R. G. Burns, A. Karrass, and D. Solitar. A note on groups with separable finitely generated subgroups. *Bulletin of the Australian Mathematical Society*, 36:153–160, 8 1987.

[33] J.O. Button and R.P. Kropholler. Non-hyperbolic free-by-cylic and one-relator groups. arXiv: 1503.01989v1, March 2015.

[34] V. Capraro and M. Lupini. Introduction to sofic and hyperlinear groups and Connes' embedding conjecture. *Lecture Notes in Mathematics*, 2136, 2015.

[35] R. Charney, J. Crisp, and K. Vogtmann. Automorphisms of 2-dimensional right angled Artin groups. *Geometry and Topology*, 2007.

[36] R. Charney, N. Stambaugh, and K. Vogtmann. Outer space for untwisted automorphisms of right-angled Artin groups. *Geometry and Topology*, 2016.

[37] R. Charney and K. Vogtmann. Finiteness properties of automorphism groups of right-angled Artin groups. *Bulletin of the London Mathematical Society*, 41(1):94–102, 2009.

[38] R. Charney and K. Vogtmann. Subgroups and quotients of automorphism groups of RAAGs. *Low-dimensional and Symplectic Topology*, 82, 2011.

[39] Y. Cornulier. A sofic group away from amenable groups. *Mathematische Annalen*, 350(2):269–275, 2011.

[40] M. Culler and K. Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones mathematicae*, 84(1):91–119, 1986.

[41] M. B. Day. Peak reduction and finite presentations for automorphism groups of right-angled Artin groups. *Geometry & Topology*, pages 817–855, 2009.

[42] R. Diestel. *Graph theory*. Number 173 in Graduate Texts in Mathematics. Springer-Verlag, Berlin, 3 edition, 2005.

[43] W. Dison and T. R. Riley. Hydra groups. *Commentarii Mathematici Helvetici. A Journal of the Swiss Mathematical Society*, 88(3):507–540, 2013.

[44] C. Droms. Isomorphisms of graph groups. *Proc. Amer. Math. Soc.*, 100(3):407–408, 1987.

[45] C. Druţu and M. Sapir. Non-linear residually finite groups. *Journal of Algebra*, 284:174–178, 2005.

[46] G. Duchamp and D. Krob. The lower central series of the free partially commutative group. *Semigroup Forum*, 45(1):385–394, 1992.

[47] A. Elkasapy and A. Thom. On the length of the shortest non-trivial element in the derived and the lower central series. *Journal of Group Theory*, 18:793–804, 2013.

[48] D. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson, and W. Thurston. *Word Processing in Groups*. Jones and Bartlett Publishers, 1992.

[49] B. Farb. Relatively hyperbolic groups. *Geometric and Functional Analysis*, 8(5):810 –840, 1998.

[50] M. Feighn and M. Handel. Mapping tori of free group automorphisms are coherent. *Annals of Mathematics*, 149(3):1061–1077, 1999.

[51] F. Gautero and M. Lustig. The mapping torus group of a free group automorphism is hyperbolic relative to the canonical subgroups of polynomial growth. arXiv: 0707.0822v2, June 2013.

[52] R. Geoghegan, M. L. Mihalik, M. Sapir, and D. T. Wise. Ascending HNN extensions of finitely generated free groups are hopfian. *Bulletin of the London Mathematical Society*, 33(3):292–298, 2001.

[53] S. M. Gersten. Dehn functions and $l1$-norms of finite presentations. In G. Baumslag and C. F. Miller, editors, *Algorithms and Classification in Combinatorial Group Theory*, pages 195–224. Springer New York, New York, NY, 1992.

[54] S.M. Gersten. Cohomological lower bounds for isoperimetric functions on groups. *Topology*, 37(5):1031 – 1072, 1998.

[55] S.M. Gersten and T.R. Riley. Some duality conjectures for finite graphs and their group theoretic consequences. *Proceedings of the Edinburgh Mathematical Society*, 2:389–421, 2005.

[56] M. Gromov. Hyperbolic groups. In S. Gersten, editor, *Essays in Group Theory*, volume 8 of *MSRI Publications*, pages 75–265. Springer, 1987.

[57] D. Groves and M. Hull. Abelian splittings of right-angled Artin groups. arXiv:1502.00129[math.GR], 2015.

[58] M. Handel and W.P. Thurston. New proofs of some results of nielsen. *Advances in Mathematics*, 56(2):173–191, 1985.

[59] J. L. Harer. The virtual cohomological dimension of the mapping class group of an orientable surface. *Inventiones mathematicae*, 84(1):157–176, 1986.

[60] J. Hempel. Residual finiteness for 3-manifolds. *Combinatorial Group Theory and Topology*, pages 379–396, 1984.

[61] L. Ji. Well-rounded equivariant deformation retracts of Teichmüller spaces. *L'Enseignement Mathematique*, 2014.

[62] D. Kahrobaei. *Doubles of residually solvable groups*, volume 1 of *Algebra and Discrete Mathematics*, pages 192–200. World Scientific, 2008.

[63] D. Kahrobaei. Doubles of residually solvable groups. In *Aspects of infinite groups*, volume 1 of *Algebra and Discrete Mathematics*, pages 192–200. World Sci. Publ., Hackensack, NJ, 2008.

[64] E. R. Van Kampen. On some lemmas in the theory of groups. *American Journal of Mathematics*, 55(1):268–273, 1933.

[65] I. Kapovich. Mapping tori of endomorphisms of free groups. *Communications in Algebra*, 28(6):2895–2917, 2000.

[66] A. Kar and N. Nikolov. A non-LEA sofic group. arXiv:1405.1620 [math.GR], 2015.

[67] M. Kassabov and F. Matucci. Bounding the residual finiteness of free groups. *Proc. Amer. Math. Soc.*, 139:2281–2286, 2011.

[68] O. Kharlampovich, A. Myasnikov, and M. Sapir. Algorithmically complex residually finite groups. arXiv:1204.6506, March 2013.

[69] P. H. Kropholler. Baumslag-Solitar groups and some other groups of cohomological dimension two. *Commentarii Mathematici Helvetici*, 65(1):547–558, 1990.

[70] M. Laurence. A generating set for the automorphism group of a graph group. *Journal of the London Mathematical Society*, 52:318–334, 1995.

[71] A. Lubotzky and A. Mann. On groups of polynomial subgroup growth. *Inventiones mathematicae*, 104:521–533, 1991.

[72] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory*. Springer, 1977.

[73] N. Macura. Quadratic isoperimetric inequality for mapping tori of polynomially growing automorphisms of free groups. *Geometric & Functional Analysis GAFA*, 10(4):874–901, 2000.

[74] A.I. Malcev. On isomorphic matrix representations of infinite groups of matrices. *Matematiceskij sbornik*, 8:405–422, 1940. (English transl. in: Amer. Math. Soc. Transl. (2) 45 (1965) 118).

[75] A.I. Malcev. On homomorphisms onto finite groups. *Uchen. Zap. Ivanov Gos. Ped. Inst.*, 18:49–60, 1956. (English transl. in: Amer. Math. Soc. Transl. (2) 119 (1983) 6779).

[76] J. McCarthy. A Tits-Alternative for subgroups of surface mapping class groups. *Transactions of the American Mathematical Society*, 291(2):583–612, 1985.

[77] J.C.C. McKinsey. The decision problem for some classes of sentences without quantifiers. *Journal of Symbolic Logic*, 8:61–76, 1943.

[78] V. Metafsis and E. Raptis. On the profinite topology of right-angled artin groups. *Journal of Algebra*, 320:1174–1181, 2008.

[79] K.A. Mihajlov. Problma vhodni dl prmyh proizvdnij grupp:the occurrence problem for direct products of groups. *Matematiceskij sbornik*, 70:241–251, 1966.

[80] M. Newman. *Integral matrices*. Academic Press, New York, 1972.

[81] G. A. Niblo and D. T. Wise. Subgroup separability, knot groups and graph manifolds. *Proceedings of the American Mathematical Society*, 129(3):685–693, 2001.

[82] D.V. Osin. *Relatively hyperbolic groups: Instrinsic geometry, algebraic properties, and algorithmic problems*, volume 179 of *Memoirs of the American Mathematical Society*. American Mathematical Society, 2006.

[83] P. Papasoglu. On the asymptotic cone of groups satisfying a quadratic isoperimetric inequality. *Journal of Differential Geometry*, 44:789–806, 1996.

[84] P. Patel. On a theorem of Peter Scott. *Proc. Amer. Math. Soc.*, 142:2891–2906, 2014.

[85] M. Sapir, J-C. Birget, and E. Rips. Isoperimetric and isodiametric functions of groups. *Annals of Mathematics*, 156:345–466, 2002.

[86] P. Scott. Subgroups of surface groups are almost geometric. *Journal of the London Mathematical Society*, s2-17(3):555–565, 1978.

[87] H. Servatius. Automorphisms of graph groups. *Journal of Algebra*, 126:34–60, 1989.

[88] E. V. Sokolov. A characterization of root classes of groups. *Communications in Algebra*, 43(2):856–860, 2015.

[89] J. Tits. Free subgroups in linear groups. *Journal of Algebra*, 20(2):250 – 270, 1972.

[90] D.T. Wise. The structure of groups with a quasiconvex hierarchy. Preprint. Available from http://www.math.mcgill.ca/wise/papers.html, 2011.

[91] Leonard Van Wyk. Graph groups are biautomatic. *Journal of Pure and Applied Algebra*, 94(3):341 – 352, 1994.