CERTAIN FREE GROUP FUNCTIONS AND UNTANGLING CLOSED CURVES ON SURFACES

BY

NEHA GUPTA

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2016

Urbana, Illinois

Doctoral Committee:

Professor Emeritus Paul Schupp, Chair
Professor Ilya Kapovich, Director of Research
Associate Professor Chris Leininger
Professor Igor Mineyev

Abstract

This thesis primarily addresses the problem of untangling closed geodesics in finite covers of hyperbolic surfaces. Our motivation comes from results of Scott and Patel. Scott's result tells us that one can always untangle a closed geodesic on a hyperbolic surface in a finite degree cover. Our goal is to quantify the degree of this cover in which the geodesic untangles in terms of the length of the geodesic. Our approach is to introduce and study the notions of primitivity, simplicity and non-filling index functions for finitely generated free groups. In joint work with Ilya Kapovich we obtain lower bounds for these functions and relate these free group results back to the setting of hyperbolic surfaces. Chapters 1-6 in parts comprise of a joint paper with Kapovich that is under review. Chapter 7 discusses the problem of Nielsen equivalence in a particular class of generic groups.

To my Biji, for never leaving my side, not when I crossed the Caspian Sea, not when I crossed the Atlantic Ocean, and not even when she crossed the seven seas as I know them.

Acknowledgments

This thesis is mathematically and philosophically, a sum of it's parts. No words can justify my gratitude to Ilya Kapovich, for being an exemplary adviser, for his infinite patience with me, and his ready and honest advice. I learnt mathematics from him, but more importantly I learnt what it means to be a morally correct mathematician. For this, and for much else, I thank him wholeheartedly. I would also like to thank my thesis committee members: Paul Schupp, Chris Leininger, and Igor Mineyev. Thank you for the spicy fish, and the spicier conversations.

Thanks to Yuliy Baryshnikov for one of my favourite proofs in this thesis. Thanks also to Priyam Patel for her mathematics and her friendship, both of which have been equally motivating. Thanks to Kasra Rafi for showing me that the best mathematical conversations happen while swimming in the Adriatic Sea. I also thank Doron Puder and Khalid Bou-Rabee for their useful conversations. In the same breath, I thank Geetha Venkataraman for being the role model that one day I aspire to emulate.

My friends, near and far, have been undeniably wonderful. Thank you Sepideh for being my first friend, and my endless source of comfort. Those pistachios are to die for. Thank you Elyse for your laughter, and your support. Thank you Sarah and Brian, for the ice cream, the tequila, and for flying halfway across the world for mangoes and me. Thanks also to Pat and Derek for always being a text or a short walk away. Suvani, Pranav, Ritwik, and Kanika, thank you so much for never letting our physical distances gnaw, and for being my constants throughout. Manav, thank you for being my pick-me-up and my laughter conspirator. Aditi and Venu, thank you for the text support, for our two wonderful trips, and for the prospect of many more. John, thanks for the tough love.

My mother, I thank for her overflowing love. You are my biggest inspiration, and if it weren't for you, I simply would not have been here. My father, I thank for his faith in me, and his endless encouragement. Papa, if I had your strength, I could probably write another thesis or two! Cherry, you are my pride and my joy. Thank you for playing pretend-big-sister when I needed you to, and for being my little sister when I wanted you to. These past years would have been infinitely more difficult were it not for Don Juan's singing, warmth, and love; and were it not for Margarita's support for the choices we made, and the freedom she gives to make the decisions we will make. Thank you Carmen for always having my back, and for the trips past and future. Thank you Jaime for Apollo, and for your newly learnt best hugs ever. Finally, and most importantly, thank you Juan for being my biggest strength and my favourite weakness. There is nothing that I cannot do with you beside me, and there is nothing that I would want to do otherwise.

Table of Contents

List of Symbols

RF	Residually finite.
LERF	Locally extended residually finite.
$\pi_1(X)$	Fundamental group of X.
∂S	Boundary of a surface S.
F_N	Free Group of rank N.
F_n	Free Group of rank n.
$F(A)$	Free group with free basis A.
R_N	The N-rose i.e. a wedge of N loop edges, each labeled by a generator.
CV_N	Projectivized Culler-Vogtmann Outer Space of rank N.
cv_N	Unprojectivized Culler-Vogtmann Outer Space of rank N.
$\overline{\mathrm{CV}}_N$	Compactification of projectivized Culler-Vogtmann Outer Space of rank N.
$\lvert w\rvert_A$	Freely reduced length of a word $w \in F(A)$.
$\lvert\lvert w\rvert\rvert_A$	Cyclically reduced length of a word $w \in F(A)$.

Chapter 1

Introduction

Chapters 1-6 address the problem of untangling closed curves on hyperbolic surfaces by looking at certain free group functions instead. A group G is *residually finite* (RF) if for every nontrivial element $g \in G$, there exists a finite index subgroup H of G such that $g \notin H$. A group G is called *locally extended residually finite* (LERF) if for any finitely generated subgroup G' of G and any element $g \in G$ with $g \notin G'$, there exists a finite index subgroup H of G which contains G' but not g. This property of a group is often called "subgroup separability".

Let Σ be a compact connected surface with a hyperbolic metric ρ and with (possibly empty) geodesic boundary. In [59, 60] Scott proved that $\pi_1(\Sigma)$ is *subgroup separable* or *LERF*. In this context, that would mean that for every finitely generated subgroup $K \leq \pi_1(\Sigma)$ and every $g \in \pi_1(\Sigma)$ such that $g \notin K$ there exists a subgroup $H \leq \pi_1(\Sigma)$ of finite index in $\pi_1(\Sigma)$ such that $K \leq H$ but $g \notin H$. (Scott's result dealt with the case of a closed surface S since in the case $\partial S \neq \emptyset$, the group $\pi_1(S)$ is free and hence known to be subgroup separable by a much older result of Hall [30]).

In the same work [59] Scott showed that if γ is a closed geodesic on (Σ, ρ) then there exists a finite cover $\hat{\Sigma} \to \Sigma$ such that γ lifts to a simple closed geodesic in $\hat{\Sigma}$, where $\hat{\Sigma}$ is given the hyperbolic structure obtained by the pull-back of ρ. As customary in the context of hyperbolic surfaces, the term "closed geodesic" here assumes that the curve in question is not a proper power in the fundamental group of the surface. Figure 1 depicts Scott's result in a picture:

Scott's result raises lots of interesting questions. We are interested in quantifying the degree of this cover in which a closed geodesic untangles in terms of the length of the geodesic, or in terms of its self-intersection number.

Patel [51] obtained quantitative versions of Scott's subgroup separability result and of his result about lifting a closed geodesic to a simple one in a finite cover. She proved that for every Σ as above there exists a hyperbolic metric ρ_0 on Σ such that every closed geodesic of length L on (Σ, ρ_0) lifts to a simple closed geodesic in some finite cover of Σ of degree $\leq 16.2L$. Since the length functions on $\pi_1(\Sigma)$ coming from any two hyperbolic structures on Σ are bi-Lipschitz equivalent, it follows that for any hyperbolic structure ρ on

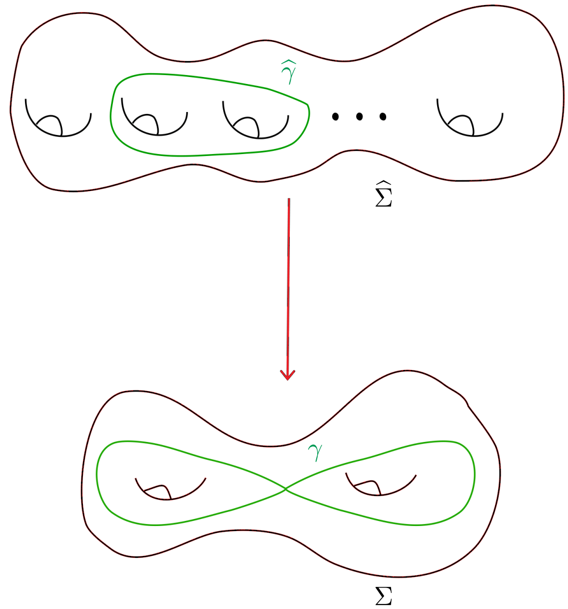

Figure 1.1: A pictorial representation of Scott's result

Σ there is some constant $c > 0$ such that every closed geodesic of length L on (Σ, ρ) lifts to a simple closed geodesic in some finite cover of Σ of degree $\leq cL$.

In order to quantify Scott's result in terms of the length of the geodesic, we need some definitions. If ρ is a hyperbolic structure on Σ, for every closed geodesic γ on (Σ, ρ) we denote by $\deg_{\Sigma,\rho}(\gamma)$ the smallest degree of a finite cover of Σ such that γ lifts to a simple closed geodesic in that cover. For $L \geq sys(\rho)$ (where $sys(\rho)$ is the shortest length of a closed geodesic on (Σ, ρ)) we define $f_{\Sigma,\rho}(L)$ to be the maximum of $\deg_{\Sigma,\rho}(\gamma)$ taken over all closed geodesics γ on (Σ, ρ) of length $\leq L$. Patel's result mentioned above implies that for every hyperbolic structure ρ on Σ there is $c > 0$ such that $f_{\Sigma,\rho}(L) \leq cL$ for all $L \geq sys(\rho)$.

Kasra Rafi observed that a simple closed geodesic on a hyperbolic surface is a particular example of a non-filling curve. Thus for a hyperbolic surface (Σ, ρ) as above and for a closed geodesic γ on Σ we can also define $\deg^{nfill}_{\Sigma,\rho}(\gamma)$ to be the smallest degree of a finite cover of Σ such that γ lifts to a non-filling closed geodesic in that cover. Again, we define $f^{nfill}_{\Sigma,\rho}(L)$ to be the maximum of $\deg^{nfill}_{\Sigma,\rho}(\gamma)$ taken over all closed geodesics γ on (Σ, ρ) of length $\leq L$. Thus, in view of Patel's result, we have $f^{nfill}_{\Sigma,\rho}(L) \leq f_{\Sigma,\rho}(L) \leq cL$ for all $L \geq sys(\rho)$.

However, up to now, nothing has been known about lower bounds for $f_{\Sigma,\rho}(L)$ or $f^{nfill}_{\Sigma,\rho}(L)$. (Note that the first place where the question about quantitative properties of $f_{\Sigma,\rho}(L)$ was raised, although somewhat indirectly, appears to have been the paper of Rivin [56]). In general, obtaining lower bounds for quantitative results related to residual finiteness is quite difficult, and is often harder than obtaining upper bounds. Recently there has been a significant amount of research regarding quantitative aspects of residual finiteness; see, for example [8, 6, 11, 12, 13, 14, 10, 17, 24, 41, 42, 51, 56, 19, 9, 16, 18]. We will discuss some of these results in more detail below.

Even though the question we are addressing here is a geometric one, our approach to answering it involves translating this question to the algebraic setting of free groups. We define some "analogous functions" for free groups, and find bounds for these free group functions (See Chapters 3, 4, 5). We then translate these free group bounds to our original surface setting (See Chapter 6). All of this is joint work with Ilya Kapovich.

Let $N \geq 2$ be an integer and let F_N be the free group of rank N. If A is a free basis of F_N, for an element $g \in F_N$ we denote by $|g|_A$ the freely reduced length of g over A and we denote by $||g||_A$ the cyclically reduced length of g over A. A classic result of Marshall Hall [30] (see also [35] for a modern proof using Stallings subgroup graphs), proves that finitely generated free groups are subgroup separable. More precisely, Hall proved that if $K \leq F_N$ is a finitely generated subgroup and $g \in F_N - K$ then there exists a subgroup $H \leq F_N$ of finite index such that $g \notin H$, $K \leq H$, and, moreover, K is a free factor of H. It is not hard to adapt the proof of this result to show that for every $g \in F_N, g \neq 1$ there exists a subgroup $H \leq F_N$ of finite index such that $g \in H$ and that g is a *primitive* element of H, that is, that g belongs to some free basis of H. In fact, a simple argument using Stallings subgroup graphs (Proposition 3.1.5) shows that if A is a free basis of F_N and w is a nontrivial cyclically reduced word in $F(A)$ of length n then there exists a subgroup $H \leq F_N$ with $[F_N : H] = n$ such that $w \in H$ is a primitive element of H. For a nontrivial element $g \in F_N$ we define the *primitivity index* $d_{prim}(g) = d_{prim}(g; F_N)$ as the minimum of $[F_N : H]$ where H varies over all subgroups of finite index in F_N containing g as a primitive element. Given a free basis A of F_N, for $n \geq 1$ we then define $f_{prim}(n) = f_{prim}(n; F_N)$ as the maximum of $d_{prim}(g)$ where g varies over all nontrivial freely reduced words of length $\leq n$ in $F_N = F(A)$ which are not proper powers in F_N. It is not hard to see that $f_{prim}(n)$ does not depend on the choice of a free basis A of F_N; we call $f_{prim}(n)$ the *primitivity index function* for F_N. Thus $f_{prim}(n)$ is the smallest monotone non-decreasing function such that for every nontrivial root-free $g \in F_N$ we have $d_{prim}(g) \leq f_{prim}(|g|_A)$.

A nontrivial element $g \in F_N$ is called *simple* in F_N if g belongs to some proper free factor of F_N. A nontrivial element $g \in F_N$ is called *filling* in F_N if g does not belong to a vertex group of a nontrivial splitting of F_N over the trivial or maximal infinite cyclic subgroup. See Chapter 2 for more precise definitions and a discussion of these notions. Note that for $1 \neq g \in F_N$, if g is primitive then g is simple, and if g is simple then g is non-filling. For a nontrivial element $g \in F_N$ let $d_{simp}(g) = d_{simp}(g; F_N)$ be the smallest index $[F_N : H]$

where H varies over all subgroups of finite index in F_N such that $g \in H$ and that g is simple in H. Finally, let $d_{nfill}(g) = d_{nfill}(g; F_N)$ be the smallest index $[F_N : H]$ where H varies over all subgroups of finite index in F_N such that $g \in H$ and that g is non-filling in H. Thus by definition, we have $d_{nfill}(g) \leq d_{simp}(g) \leq d_{prim}(g)$. For $n \geq 1$ we then define the *simplicity index function* $f_{simp}(n) = f_{simp}(n; F_N)$ as the maximum of $d_{simp}(g)$ where g varies over all nontrivial freely reduced words of length $\leq n$ in $F_N = F(A)$ that are not proper powers in F_N. Also, for $n \geq 1$ we then define the *non-filling index function* $f_{nfill}(n) = f_{nfill}(n; F_N)$ as the maximum of $d_{nfill}(g)$ where g varies over all nontrivial freely reduced words of length $\leq n$ in $F_N = F(A)$ that are not proper powers in F_N.

In view of Proposition 3.1.5 mentioned above, for every nontrivial $g \in F_N$ we have $d_{simp}(g) \leq d_{prim}(g) \leq ||g||_A \leq |g|_A$, and hence $f_{nfill}(n) \leq f_{simp}(n) \leq f_{prim}(n) \leq n$ (see Lemma 3.1.6 for details).

Note that in defining these functions, our analogies are clear. We have replaced the idea of a geodesic in the surface setting with that of a group element becoming "nice" (non-filling/simple/primitive) in a finite index subgroup of the group. Thus, in the free group setting, we want to quantify the index of a subgroup in which a group element becomes "nice" (non-filling/simple/primitive) in terms of the word length of the group element. In the surface setting, we want to quantify the degree of a cover in which a geodesic becomes simple in terms of the length of the geodesic.

In general, we are interested in the following types of questions:

- Understand the actual asymptotics of the "worst-case" index functions $f_{nfill}(n), f_{simp}(n), f_{prim}(n)$ for free groups and of their geometric counterparts $f_{\Sigma,\rho}(L)$ or $f^{nfill}_{\Sigma,\rho}(L)$.
- Find specific sequences of elements in free groups or curves on surfaces realizing this "worst-case" behavior or at least exhibiting reasonably fast growth of the corresponding index and degree functions.
- Understand the asymptotics of the indexes $d_{prim}(g_n), d_{simp}(g_n), d_{nfill}(g_n)$ and of $\deg_{\Sigma,\rho}(\gamma_n)$, $\deg^{nfill}_{\Sigma,\rho}(\gamma_n)$ for various "natural" sequences of group elements $g_n \in F_N$ or closed geodesics γ_n on (Σ, ρ).
- Understand the relationship between the index functions for free groups and the degree functions for surfaces, and relate both to other functions measuring quantitative aspects of residual properties of free and surface groups.

Our first main result provides a lower bound for $f_{nfill}(n; F_N)$:

Theorem A. Let $N \geq 2$ and let $F_N = F(A)$ where $A = a_1, \ldots, a_N$. Then there exists a constant $c > 0$ and an integer $M \geq 1$ such that for all $n \geq M$ we have

$$f_{prim}(n) \geq f_{simp}(n) \geq f_{nfill}(n) \geq c\frac{\log n}{\log \log n}.$$

For a finitely generated group G equipped with a finite generating set A, the *residual finiteness growth function* $\mathrm{RF}_G(n)$ is defined as the smallest number d such that for every nontrivial element $g \in G$ of word-length $\leq n$ with respect to A there exists a subgroup of index at most d in G that does not contain g.

For a free group F_N with a free basis A, Khalid Bou-Rabee relates $f_{prim}(n, F_N)$ to the residual finiteness growth function $\mathrm{RF}_{F_N}(n)$. Namely, he shows (Theorem 4.3.1) that for $n \geq 1$ one has $f_{prim}(4n+4, F_N) \geq \mathrm{RF}_{F_N}(n)$. Using a recent result of Kozma and Thom [42] about lower bounds for $\mathrm{RF}_{F_N}(n)$, Bou-Rabee then shows (Corollary 4.3.2) that for all sufficiently large n one has

$$f_{prim}(4n+4) \geq \exp\left(\left(\frac{\log(n)}{C \log\log(n)}\right)^{1/4}\right).$$

Note that this lower bound behaves almost like $n^{1/4}$. Moreover, if we assume Babai's Conjecture on the diameter of Cayley graphs of permutation groups, then for all sufficiently large n we have an almost linear lower bound:

$$f_{prim}(4n+4) \geq n^{\frac{1}{C \log\log(n)}}.$$

Bou-Rabee's homological trick used in Theorem 4.3.1 does not work for the index functions $f_{simp}(n)$ and $f_{nfill}(n)$. Thus for these functions the lower bound given by Theorem A remains the best known bound.

We also obtain a lower bound on $d_{simp}(w_n)$ and $d_{nfill}(w_n)$ where w_n is a "random" freely reduced word in $F(A)$ of length $n >> 1$:

Theorem B. Let $N \geq 2$ and let $F_N = F(A)$ where $A = \{a_1, \ldots, a_N\}$.

Then there exist constants $c > 0$, $D_1 > 1$, $1 > D_2 > 0$ such that for $n \geq 1$ and for a freely reduced word $w_n \in F(A)$ of length n chosen uniformly at random from the sphere $S(n)$ of radius n in $F(A)$ we have

$$P_{\mu_n}\left(d_{simp}(w_n) \geq c \log^{1/3} n\right) \geq_{n\to\infty} 1 - O\left((D_1)^{-n^{D_2}}\right)$$

and

$$P_{\mu_n}\left(d_{nfill}(w_n) \geq c \log^{1/5} n\right) \geq_{n\to\infty} 1 - O\left((D_1)^{-n^{D_2}}\right)$$

so that

$$\lim_{n\to\infty} P_{\mu_n}\left(d_{simp}(w_n) \geq c \log^{1/3} n\right) = 1$$

and

$$\lim_{n\to\infty} P_{\mu_n}\left(d_{nfill}(w_n) \geq c \log^{1/5} n\right) = 1$$

Here μ_n is the uniform probability distribution on the n-sphere $S(n) \subseteq F_N = F(A)$.

It remains an interesting question to understand the actual behavior of $d_{simp}(w_n)$ and $d_{nfill}(w_n)$ on "random" elements $w_n \in F_N$ and, in particular, to see if $d_{simp}(w_n)$ and $d_{nfill}(w_n)$ admit sublinear upper bounds.

Finally, in Chapter 6 we relate the above results for free groups to the original motivating questions about the degree functions for hyperbolic surfaces. Thus, applying Theorem B suitably, we obtain:

Theorem C. Let (Σ, ρ) be a compact connected hyperbolic surface with $b \geq 1$ geodesic boundary components. Then there exists $C' > 0$ such that for all sufficiently large L we have

$$f_{\Sigma,\rho}(L) \geq f_{\Sigma,\rho}^{nfill}(L) \geq C' \frac{\log L}{\log \log L}.$$

Similarly, using Theorem A, we obtain :

Theorem D. Let Σ be a compact connected surface with a hyperbolic structure ρ and with (possibly empty) geodesic boundary. Let $\Sigma_1 \subseteq \Sigma$ be a compact connected subsurface with ≥ 3 boundary components, each of which is a geodesic in (Σ, ρ). Let $x \in \Sigma_1$ and let A be a free basis of $\pi_1(\Sigma_1, x)$.

Let $w_n \in F(A) = \pi_1(\Sigma_1, x)$ be a freely reduced word of length n over $A^{\pm 1}$ generated by a simple non-backtracking random walk on $F(A) = \pi_1(\Sigma_1, x)$. Let γ_n be the closed geodesic on (Σ, ρ) in the free homotopy class of w_n.

Then there exist constants $c > 0, K' \geq 1$ such that

$$\lim_{n\to\infty} Pr(\deg_{\Sigma,\rho}(\gamma_n) \geq c \log^{1/3} n) = 1$$

and such that with probability tending to 1 as $n \to \infty$ we have that $w_n \in \pi_1(\Sigma, x)$ is not a proper power and that $n/K' \leq \ell_\rho(\gamma_n) \leq K'n$.

Chapters 2-6 are in a joint paper with Kapovich on the ArXiv:1411.5523. In the original November 2014 version of the ArXiv paper we used Theorem D to obtain, for all sufficiently large L, a lower bound

$$f_{\Sigma,\rho}(L) \geq c \log^{1/3} L,$$

where (Σ, ρ) is a closed hyperbolic surface. At the time this was the only known lower bound for $f_{\Sigma,\rho}(L)$. Motivated by our work, Jonah Gaster [23] subsequently obtained a linear lower bound $f_{\Sigma,\rho}(L) \geq cL$ and exhibited a specific sequence of curves γ_n in Σ, living in a pair-of-pants subsurface of Σ, realizing this lower bound. Thus, Gaster shows that for his sequence of curves, one needs a cover of degree at least cL to untangle. He does not address what one needs for a random sequence of curves, which is what we work with. Further, since Gaster's curves are already non-filling in Σ and have $\deg^{nfill}_{\Sigma,\rho}(\gamma_n) = 1$, his proof does not provide any lower bounds for $f^{nfill}_{\Sigma,\rho}(L)$. Thus, for the moment the lower bound for $f^{nfill}_{\Sigma,\rho}(L)$ given by Theorem C remains the best bound known.

In Section 6.3 we also relate our results to the versions of $f_{\Sigma,\rho}(L)$ and $f^{nfill}_{\Sigma,\rho}(L)$ that do not involve a hyperbolic metric and use the geometric intersection number $i([\gamma],[\gamma])$ instead of the hyperbolic length of γ in their definitions. Also, in Section 3.2 we prove algorithmic computability of the indexes $d_{prim}(g, F_N)$ $d_{simp}(g, F_N)$, $d_{nfill}(g, F_N)$ and of the corresponding index functions $f_{prim}(n), f_{simp}(n), f_{nfill}(n)$; see Theorem 3.2.14 and Theorem 3.2.18.

In Chapter 7 we discuss the question of establishing Nielsen equivalence classes of generating tuples in a certain class of random groups. The precise meaning of what we mean by "random" here is addressed in Section 7.2.2 of Chapter 7. Jakob Nielsen defined the notion of Nielsen Equivalence in the 1920s [49, 50]. If G is a group, $n \geq 1$, and $\tau = (g_1, \dots, g_n)$ is an ordered n-tuple of elements in G, an *elementary Nielsen transformation* on τ is one of the following three types of moves:

1. For some $i \in \{1, \dots, n\}$ replace g_i in τ by g_i^{-1}
2. For some $i \neq j$, $i, j \in \{1, \dots, n\}$ interchange g_i and g_j in τ
3. For some $i \neq j$, $i, j \in \{1, \dots, n\}$ replace g_i in τ by $g_i g_j^{\pm 1}$

Two n-tuples $\tau = (g_1, \dots, g_n)$ and $\tau' = (g'_1, \dots, g'_n)$ are called *Nielsen equivalent*, denoted $\tau \sim_{\mathrm{NE}} \tau'$, if there exists a finite chain of elementary Nielsen transformations taking τ to τ'.

In general, it is quite hard to distinguish between Nielsen equivalence classes of n-tuples that generate the same group. Our specific goal in Chapter 7 is to extend a result of Kapovich and Weidmann [39] which essentially shows that in a certain class of groups, there exist generating n-tuples $(a_1, \dots, a_n)$, and $(b_1, \dots, b_n)$ such that the $(2n-1)$-tuples
$(a_1, \dots, a_n, \underbrace{1, \dots, 1}_{n-1 \text{ times}})$ and $(b_1, \dots, b_n, \underbrace{1, \dots, 1}_{n-1 \text{ times}})$ are not Nielsen-equivalent. In particular, this group admits at least 2 Nielsen equivalence classes of generating n-tuples. We show that in fact for any $k \geq 2$, there exists a class groups that admits precisely k Nielsen equivalence classes of generating n-tuples. More precisely, the main result here is:

Theorem E. Let $k \geq 2$, $n \geq 2$ be arbitrary integers. Then there exists a generic set $\mathcal{R}$ of kn-tuples

$$\tau = (u_{11}, \dots, u_{1n}, u_{21}, \dots, u_{2n}, \dots, u_{k1}, \dots, u_{kn})$$

where for each $\tau \in \mathcal{R}$, $i \in \{1, \dots, k\}$, $j \in \{1, \dots, n\}$, u_{ij} is a cyclically reduced word in $F(a_{i1}, \dots, a_{in})$.

Further, $|u_{11}| = \ldots = |u_{1n}| = \ldots = |u_{k1}| = \ldots = |u_{kn}|$ and such that the following holds for each $\tau = (u_{11}, \ldots, u_{1n}, u_{21}, \ldots, u_{2n}, \ldots, u_{k1}, \ldots, u_{kn})$:

Let G be a group given by the presentation:

$$G =< a_{11}, \ldots, a_{1n}, \ldots, a_{k1}, \ldots, a_{kn} | a_{1j} = u_{2j}(\underline{a_2}), a_{2j} = u_{3j}(\underline{a_3}), \ldots, a_{(k-1)j} = u_{kj}(\underline{a_k}),$$
$$a_{kj} = u_{1j}(\underline{a_1}), \text{ for } 1 \leq j \leq n >, \qquad (*)$$

where for $i \in \{1, \ldots, k\}$, $(\underline{a_i}) = (a_{i1}, \ldots, a_{in})$. Then G is a torsion-free word-hyperbolic one-ended group of rank n admitting precisely k Nielsen equivalence classes of generating n-tuples.

We use "genericity" conditions and small cancellation conditions to obtain this result. As an immediate corollary one gets that:

Corollary A. For all integers $k \geq 2, n \geq 2$ there exists a non-elementary torsion-free word-hyperbolic one-ended group of rank n such that G has exactly k distinct Nielsen equivalence classes of n-tuples that generate G.

It should be noted that for the case $k = 1$, the above result holds due to Kapovich and Schupp [37].

Finally, in Chapter 8 we discuss a short aside on the index of the image of a finite index subgroup of $F(a, b)$ in a quotient $F(a, b)/N$. We also discuss some interesting open questions with regards to the material presented in Chapters 1-6.

Chapter 2

Preliminaries

2.1 Graphs and Edge Paths

The exposition in this section follows that of [36].

Definition 2.1.1. A *graph* is a 1-dimensional cell-complex. The 0-cells of Γ are called *vertices* and we denote the set of vertices of Γ by $V\Gamma$. The open 1-cells of Γ are called *topological edges* of Γ and the set of topological edges are denoted by $E_{top}\Gamma$.

Every topological edge of Γ is homeomorphic to the open interval $(0, 1)$ and thus, when viewed as a 1-manifold, admits two possible orientations. An *oriented edge* of Γ is a topological edge with a choice of orientation on it. We denote by $E\Gamma$ the set of all oriented edges of Γ. If $e \in E\Gamma$ is an oriented edge, we denote by $\bar{e}$ the same underlying edge with the opposite orientation. Note that for every $e \in E\Gamma$ we have $\bar{e} \neq e$ and $\bar{\bar{e}} = e$; thus $: E\Gamma \to E\Gamma$ is an involution with no fixed points.

Since Γ is a cell-complex, every oriented edge $e \in E\Gamma$ comes equipped with the orientation-preserving attaching map $j_e : [0, 1] \to \Gamma$ such that j_e maps $(0, 1)$ homeomorphically to e and such that $j_e(0), j_e(1) \in V\Gamma$ (though not necessarily distinct). For $e \in E\Gamma$ we call $j_e(0)$ the *initial vertex* of e, denoted $o(e)$, and we call $j_e(1)$ the *terminal vertex* of e, denoted $t(e)$. Thus, by definition, $o(\bar{e}) = t(e)$ and $t(\bar{e}) = o(e)$.

For any vertex $x \in V\Gamma$, the *degree of* x in Γ denoted by $deg(x)$ is the cardinality of the set $\{e \in E\Gamma | o(e) = x\}$.

An *orientation* on a graph Γ is a partition $E\Gamma = E_+\Gamma \sqcup E_-\Gamma$ such that for an edge $e \in E\Gamma$ we have $e \in E_+\Gamma$ if and only if $\bar{e} \in E_-\Gamma$.

An *edge-path* p in Γ is a sequence of edges $e_1, e_2, \ldots, e_k$ with $e_i \in E\Gamma$ for all i and $o(e_j) = t(e_{j-1})$ for all $2 \leq j \leq k$. The length $|p|$, of the path p is the number of edges in p, that is $|p| = k$. We put $o(p) = o(e_1)$, and $t(p) = t(e_k)$. We define $p^{-1} := e_k, e_{k-1}, \ldots, e_1$. A path p in a graph Γ is *reduced* if it does not contain any sub-paths of the form e, e^{-1} where $e \in E\Gamma$ is an edge.

Definition 2.1.2. For two graphs Γ_1 and Γ_2, a morphism or a graph-map $f : \Gamma_1 \to \Gamma_2$ is a continuous map f such that $f(V\Gamma_1) \subseteq V\Gamma_2$ and such that the restriction of f to any topological edge $e \in \Gamma_1$ is a homeomorphism between e and some topological edge e' of Γ_2. Thus a morphism $f : \Gamma_1 \to \Gamma_2$ naturally defines functions $f : E\Gamma_1 \to E\Gamma_2$ and $f : V\Gamma_1 \to V\Gamma_2$ such that for any $e \in E\Gamma_1$ we have $f(\bar{e}) = \overline{f(e)} \in E\Gamma_2$, $o(f(e)) = f(o(e))$ and $t(f(e)) = f(t(e))$.

Definition 2.1.3. Let Γ be a graph and $x \in V\Gamma$. Then the core of Γ at x is defined as:

$$Core(\Gamma, x) = \cup\{p \,|\, \text{where}\, p\, \text{is a reduced path in}\, \Gamma\, \text{from}\, x\, \text{to}\, x\}$$

Note that $Core(\Gamma, x)$ is a connected subgraph of Γ containing x. If $Core(\Gamma, x) = \Gamma$ we say that Γ is a *core graph with respect to* x. The graph $Core(\Gamma, x)$ has no degree 1 vertices except possibly x itself.

Proposition-Definition 2.1.4. Let Γ be a graph, and $x \in V\Gamma$. Choose a maximal subtree $T \subseteq \Gamma$, and an orientation $E\Gamma = E_+\Gamma \sqcup E_-\Gamma$. For $e \in E\Gamma$ define $[x, o(e)]_T$ to be the unique reduced path in T from x to $o(e)$, and let $s_e := [x, o(e)]_T\, e\, [t(e), x]_T$. Let $S_T := \{s_e \,|\, e \in E_+\Gamma - T\}$. Then $\pi_1(\Gamma, x)$ is free and S_T is a free basis of $\pi_1(\Gamma, x)$.

We call S_T the free basis of $\pi_1(\Gamma, x)$ *dual to* T.

We need to explicitly say how to rewrite elements of $\pi_1(\Gamma, x)$ in terms of the basis S_T, both as freely reduced words and cyclically reduced words.

Proposition 2.1.5. *Let $\gamma \in \pi_1(\Gamma, x)$ and T be as above. Suppose $E_+\Gamma - T = \{e_1, \ldots, e_m\}$. Then $S_T = \{s_{e_i} | 1 \le i \le m\}$. Then:*

1. *Rewriting γ as a freely reduced word in S_T: Delete from γ all edges of T and replace each $e_i^{\pm 1}$ by $s_{e_i}^{\pm 1}$. The result is a freely reduced word over S_T representing $\gamma \in \pi_1(\Gamma, x)$.*

2. *Rewriting γ as a cyclically reduced word in S_T: First cyclically reduce the edge-path γ by removing the maximal initial and terminal segments of γ that cancel in the concatenation $\gamma\gamma$. The result is a subpath γ_1 of γ such that γ_1 is a closed cyclically reduced path (though γ_1 maybe based at a vertex different from x). Now apply the previous procedure to γ_1: delete all edges of T and replace each $e_i^{\pm 1}$ by $s_{e_i}^{\pm 1}$. The result is the cyclically reduced form of $\gamma \in \pi_1(\Gamma, x)$ over S_T.*

2.2 Graphs and subgroups

In a seminal paper from 1983 Stallings [63] used labeled graphs to study subgroups of free groups. We give a brief exposition of the relevant definitions and results below and refer the reader to[35] for details.

Recall that we fix for the free group $F_N = F(A) = F(a_1, \ldots, a_N)$ (where $N \geq 2$), a distinguished free basis $A = \{a_1, \ldots, a_N\}$. If w is a word in $\Upsilon = A \sqcup A^{-1}$, we will denote by $\underline{w}$ the freely reduced word in Υ obtained from w by performing all possible (if any) free reductions.

Definition 2.2.1. An A-graph Γ consists of an oriented graph where every edge $e \in E\Gamma$ is labeled by a letter $\mu(e) \in A \sqcup A^{-1}$ in such a way that $\mu(\bar{e}) = (\mu(e))^{-1}$. Multiple edges between vertices and loops at a vertex are allowed. An A-graph Γ is said to be *folded* if there does not exist a vertex x and two distinct edges e_1, e_2 with $o(e_1) = o(e_2) = x$ such that $\mu(e_1) = \mu(e_2)$. Otherwise Γ is said to be *non-folded.*

An A-graph Γ is said to be *A-regular* if for every vertex $x \in V\Gamma$ and for every a_i, there is precisely one outgoing edge at x labeled by a_i and precisely one incoming edge at x labeled by a_i (thus, in particular, an A-regular graph is folded).

If Γ is an A-graph and $p = e_1, \ldots, e_k$ is an edge-path in Γ, then p has a label which is a word in $A \sqcup A^{-1}$ and we denote this label by $\mu(p) = \mu(e_1)\mu(e_2)\ldots\mu(e_k)$. The definitions immediately imply:

Lemma 2.2.2. *An A-graph Γ is folded if and only if the label of every reduced path in Γ is a freely reduced word.*

Definition 2.2.3. Let Γ be an A-graph. Suppose e_1, e_2 are distinct edges with a common initial vertex $o(e_1) = o(e_2) = x$ and with the same label $\mu(e_1) = \mu(e_2) = a \in A \sqcup A^{-1}$. We fold the two edges e_1 and e_2 into a single edge e with $\mu(e) = a$. The resulting A-graph Γ' is said to have been obtained from Γ via a *fold.*

The following proposition is from [63] and immediately follows from definitions:

Proposition 2.2.4. *Let Γ be a connected A-graph, and let Γ' be obtained from Γ via a fold. Then $rank(\pi_1(\Gamma')) \leq rank(\pi_1(\Gamma))$.*

In this thesis, when we refer to the *rank of a graph*, we will always mean the *rank of the fundamental group of the graph.*

Definition 2.2.5. For any two A-graphs Γ_1 and Γ_2, a map $f : \Gamma_1 \to \Gamma_2$ is an A-morphism if f is a graph-map such that $\mu(e) = \mu(f(e))$.

For $F_N = F(a_1, \ldots, a_N)$ we define the *standard N-rose* R_N to be the wedge of N loop-edges each labeled by $a_1, \ldots, a_N$ respectively, at a vertex x_0. Then $F(A) = \pi_1(R_N, x_0)$.

For Γ an A-graph, $x \in V\Gamma$ and μ as before, we can define a map $\mu_\# : \pi_1(\Gamma, x) \to F(A)$ as $p \mapsto \underline{\mu(p)}$. This map is a group homomorphism.

Notation 2.2.6. For Γ an A-graph, $x \in V\Gamma$ we say that (Γ, x) represents the subgroup $H := \mu_\#(\pi_1(\Gamma, x)) \leq F(A)$.

Stallings also showed that any A-graph can be folded without changing the image of the induced homomorphism $\mu_{\#}$.

Proposition-Definition 2.2.7. [63, 35] Let $H \leq F(A)$. Then there exists a connected, folded A-graph Γ with $x_0 \in V\Gamma$ such that $\Gamma = Core(\Gamma, x_0)$ and (Γ, x_0) represents

$$H = \{\mu(p) \,|\, p \,\text{is a reduced path in}\, \Gamma \,\text{from}\, x_0 \,\text{to}\, x_0\} \leq F(A)$$

Moreover, such a (Γ, x_0) is unique. This graph (Γ, x_0) is called the *Stallings subgroup graph of* H with respect to A.

If (Γ, x_0) is the Stallings subgroup graph for H, then the labeling map $\mu : \pi_1(\Gamma, x_0) \to H$ is a group isomorphism. If $T \subseteq \Gamma$ is a maximal tree and $S_T = \{s_e | e \in E_+(\Gamma - T)\}$ is the dual free basis of $\pi_1(\Gamma, x_0)$, then $\mu(S_T) = \{\mu(s_e) | e \in E_+(\Gamma - T)\}$ is a free basis of H.

Given a tuple $\tau = (g_1, \ldots, g_k)$ of elements from $F(A)$, we can construct a graph S_τ with base vertex x_0 such that $\mu_{\#}(\pi_1(S_\tau, x_0)) =< g_1, \ldots, g_k >$ as follows: for $1 \leq i \leq k$, if $g_i \neq 1$ draw a circle c_i at vertex x_0 such that the label of c_i is the reduced word representing g_i.

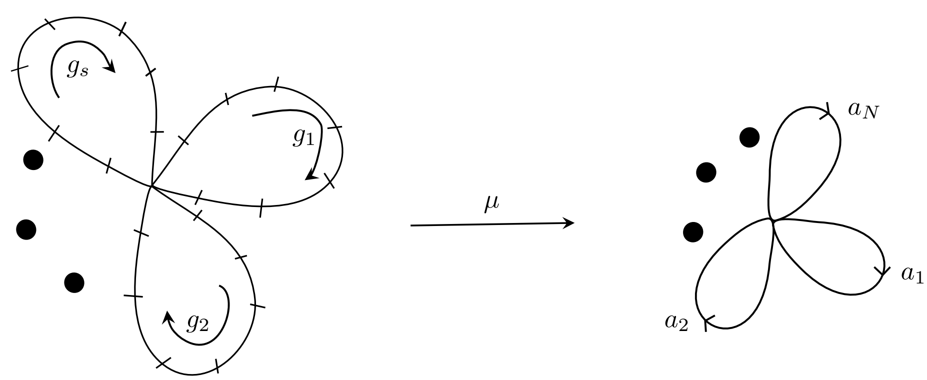

Figure 2.1: The map μ from S_τ to R_N

In particular since folding a graph does not change the image of the induced homomorphism, we have that:

Lemma 2.2.8. *Let Γ be an A-graph such that $\mu_{\#} : \pi_1(\Gamma, x) \to \pi_1(R_N, x_0)$ is surjective. Then there exists a finite sequence of A-graphs*

$$\Gamma = \Gamma_0, \Gamma_1, \ldots, \Gamma_n = R_N$$

such that Γ_i can be obtained from Γ_{i-1} by a single fold for $1 \leq i \leq n$.

Remark 2.2.9. Note that by Proposition 2.2.4, in the above sequence if $rank(\Gamma_0) = N$, then for every $1 \leq i \leq n$, we have that $rank(\Gamma_i) = N$.

The remainder of this section is relevant for Chapter 7. We will now describe some transformations of labeled graphs that were first described by Arzhantseva and Ol'shanskii [3]. Here we describe them as defined in [37]. Let us fix a quotient group $G = F(A)/N$ where N is a normal subgroup of $F(A)$ for the remaining part of this section.

An *arc* in a graph Γ is a simple path, possibly closed, where every intermediate vertex of the path has degree two in Γ.

Definition 2.2.10. (Arzhantseva-Ol'shanskii "AO move") Let $p = p_1 p' p_2$ be a reduced edge-path in an A-graph Γ such that p' is an arc of Γ and the paths p_1, p_2 do not overlap p'. Let p have initial vertex x, terminal vertex y, and label $\mu(p) = v$. Let z be a reduced word such that $v =_G z$.
We now modify Γ by adding a new arc q from x to y with label z and removing all the edges of p' from Γ. We will say that the resulting A-graph Γ' is obtained from Γ by an *AO-move*.

This move essentially "completes a relator cycle" by adding in the path q and then removing the arc p'. Further note that the first part of the move (adding q) decreases the Euler characteristic by one while the second part (removing p') increases the Euler characteristic by one. Hence, the Euler characteristic is unchanged by an AO-move.

The following two propositions are from [37].

Proposition-Definition 2.2.11. Let Γ be a connected A-graph with a base-vertex x. Then the labeling of paths gives rise to a homomorphism

$$\phi : \pi_1(\Gamma, x) \to G$$

such that for every path p from x to x we have $\phi([p]) =_G \mu(p)$ in G where $[p]$ stands for the equivalence class of p in $\pi_1(\Gamma, x)$. We say that $H = \phi(\pi_1(\Gamma, x)) \leq G$ is the subgroup represented by $\pi_1(\Gamma, x)$.
Moreover the following is true:

(1) If Γ is finite then $image(\phi)$ is finitely generated. In addition, Γ has Euler characteristic $1 - k$ if and only if the free group $\pi_1(\Gamma, x)$ has rank k and hence $\phi(\pi_1(\Gamma, x))$ can be generated by k elements.

(2) If x_0 is another vertex of Γ then the pairs (Γ, x) and (Γ, x_0) define conjugate subgroups of G.

(3) Every finitely generated subgroup of G can be represented in this fashion for some finite connected Γ. Moreover, if $H \leq G$ is k-generated, then H can be represented by a connected A-graph of Euler characteristic $\geq 1 - k$.

Proposition 2.2.12. *Let Γ be a connected A-graph with a base-vertex x. Suppose Γ' is obtained from Γ by a finite sequence of folds and AO- moves and that x' is the image of x in Γ'. Then the pairs (Γ, x) and (Γ', x') define the same subgroup of G.*

Moreover, if Γ' is obtained from Γ by removing a degree-one vertex, then for every pair of vertices y of Γ and y' of Γ', the pairs (Γ, y) and (Γ', y') define conjugate subgroups of G.

2.3 Primitive, simple, and non-filling elements

Definition 2.3.1 (Primitive elements and simple elements)**.** In the free group F_N, a non-trivial element $g \in F_N$ is called *primitive* in F_N if g belongs to some free basis of F_N.

In the free group F_N, a non-trivial element $g \in F_N$ is called *simple* in F_N if g belongs to a proper free factor of F_N. That is, there exist non-trivial subgroups H, K of F_N such that $g \in H$ and F_N is the free product $F_N = H * K$.

Definition 2.3.2 (Non-filling elements)**.** An element $g \in F_N$ is said to be *non-filling* in F_N if there exists a splitting of F_N as an amalgamated free product $F_N = K *_C L$ or as an HNN-extension $F_N = \langle K, t | t^{-1} C t = C' \rangle$, such that $C \leq F_N$ is either trivial or a maximal cyclic subgroup, such that in the $F_N = K *_C L$ case $C \neq K, C \neq L$, and such that $g \in K$.

An element $g \in F_N$ is said to be *filling* in F_N if g is not non-filling.

Remark 2.3.3. Note that if $g \in F_N$ is primitive, then it is also simple. Similarly, if $g \in F_N$ is simple, then g is non-filling.

Also, for elements of F_N the properties of being primitive, being simple and being non-filling are preserved under applying arbitrary automorphisms of F_N.

2.3.1 Culler-Vogtmann Outer Space and Closure

There are multiple descriptions of Culler-Vogtmann Outer Space that exist. We will describe one in terms of actions on trees. Please refer to [5, 22, 32, 34, 43] for details. An $\mathbb{R}$-tree is a space T with metric d where any two points $P, Q \in T$ are joined by a unique arc and this arc is isometric to the interval $[0, d(P, Q)] \subset \mathbb{R}$. We'll assume that an $\mathbb{R}$-tree T always comes with a left action of F_N on T by isometries. Any isometry w of T is either *elliptic*, in which case it fixes at least one point of T , or else it is *hyperbolic*, in which case there is an axis $Ax(w)$ in T, isometric to $\mathbb{R}$, which is w-invariant, and along which w acts as translation. A tree action is called *small* [43] if any two group elements that fix pointwise a non-trivial arc in T commute. It is called *very small* [43] if moreover:

(i) the fixed set $Fix(g) \subset T$ of any elliptic element $1 \neq g \in F_N$ is a segment or a single point (i.e., no branching), and

(ii) $Fix(g) = Fix(g^m)$ for all $g \in F_N$ and $m \geq 1$.

Let T, T' be $\mathbb{R}$-trees; a map $f : T \to T'$ is called a *homothety* if f is F_N-equivariant and bijective, and if there is some positive real number λ such that for any $x, y \in T$, we have $d_{T'}(f(x), f(y)) = \lambda d_T(x, y)$.

Let $N \geq 2$. The *unprojectivized Outer space* cv_N consists of all minimal free and discrete isometric actions of F_N on $\mathbb{R}$-trees (where two such actions are considered equal if there exists an F_N -equivariant isometry between the corresponding trees). For each $g \in F_N$ and $T \in \text{cv}_N$, we define *translation length*, denoted $||g||_T$, by $||g||_T = \inf\{d(x, gx)|x \in T\}$. Further, *projectivized Culler-Vogtmann outer space* is the subset $\text{CV}_N \subseteq \text{cv}_N$ which consists of those trees T for which the quotient metric graph T/F_N has volume 1. The *closed Outer space of rank* N, denoted $\overline{\text{CV}}_N$, is the topological space whose underlying set consists of homothety classes of very small isometric actions of F_N on $\mathbb{R}$-trees.

The following proposition relates the property of being filling in F_N to $\overline{\text{CV}}_N$:

Proposition 2.3.4. *[34, 62] Let $1 \neq g \in F_N$. Then the following conditions are equivalent:*

1. *The element g is filling in F_N.*
2. *For every minimal very small isometric action of F_N on a nontrivial simplicial tree T we have $||g||_T > 0$.*
3. *For every minimal very small isometric action of F_N on a nontrivial $\mathbb{R}$-tree T we have $||g||_T > 0$.*

Proof. The proof of this statement is implicit in [34, 62] but we sketch the argument for completeness.

Part (3) directly implies part (2). Since the simplicial splittings that appear in Definition 2.3.2 are very small, part (2) also directly implies part (1).

To see that part (1) implies part (3), suppose that $1 \neq g \in F_N$ is filling but that there exists a minimal very small isometric action of F_N on a nontrivial $\mathbb{R}$-tree T we have $||g||_T = 0$. Then a result of Guirardel [28] implies that there exists a very small minimal simplicial F_N-tree T' with $||g||_{T'} = 0$. Taking the quotient graph of groups T'/F_N and collapsing all edges except one in this graph gives us a splitting of F_N as in Definition 2.3.2 such that g is conjugate to a vertex group element for that splitting. This contradicts the assumption that g is filling in F_N. Thus (1) implies (3), as required.

□

2.4 Whitehead Graphs

We now describe the relationship between simple elements, primitive elements, and Whitehead graphs.

Definition 2.4.1. [Whitehead graph] Let $F_N = F(A)$ be as before and let $w \in F_N$ be a nontrivial cyclically reduced word. Let c be the first letter of w. The word wc is then freely reduced.

The *Whitehead graph of* w with respect to A, denoted by $Wh_A(w)$, is an undirected graph whose set of vertices $V(Wh_A(w)) = \Upsilon$. Edges are added as follows: For $a, b \in V(Wh_A(w))$, there is an undirected edge joining a^{-1} and b if ab or $b^{-1}a^{-1}$ occurs as a subword of wc.

Example 2.4.2. For $w = ab^2 \in F(a,b) = F(A)$, $wc = abba$, and then $Wh_A(w)$ is given by:

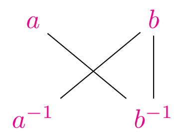

Note that if $\tilde{w}$ is a cyclic permutation of w or of w^{-1} then $Wh_A(w) = Wh_A(\tilde{w})$.

For an arbitrary $1 \neq g \in F_N$, we put $Wh_A(g) := Wh_A(w)$, where w is the cyclically reduced form of g in $F(A)$.

Recall that a *cut vertex* in a graph Δ is a vertex x such that $\Delta - \{x\}$ is disconnected. Note that if Δ has at least one edge and is disconnected, then Γ does possess a cut vertex; namely any end-vertex of an edge of Δ is a cut vertex in this case.

Example 2.4.3. Observe that the Whitehead graph in 2.4.2 has a cut vertex at the vertex b:

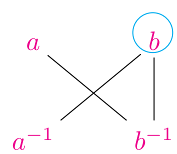

Generalizing a result of Whitehead, Stallings established the relationship between simple elements and Whitehead graphs [64]:

Proposition 2.4.4. *[64] Let $F_N = F(A)$, where $N \geq 2$ and let $g \in F(A)$ be a cyclically reduced word. If g is simple, then the Whitehead graph $Wh_A(g)$ has a cut vertex.*

Notice that Remark 2.3.3 implies that if $g \in F(A)$ is primitive, then $Wh_A(g)$ has a cut vertex.

Remark 2.4.5. Stallings' definition of Whitehead graphs differs slightly from our definition. Assume the same setting as in Definition 2.4.1. Stallings adds an edge from a^{-1} to b for *each* occurrence of a subword

ab in wc. Let us call the Whitehead graph of a cyclically reduced word w under Stallings' definition Γ, and the corresponding graph under our definition Γ_1. It is clear that $V(\Gamma) = V(\Gamma_1)$. Further it is easily checked that $x \in V(\Gamma)$ is a cut-vertex in Γ if and only if $x \in V(\Gamma_1)$ is a cut-vertex in Γ_1. Thus Proposition 2.4.4 holds for our definition of Whitehead graphs just as well.

Finally, note that if the Whitehead graph of an element has a circuit that contains all the vertices, then it can not have a cut vertex. This occurs, for instance, when the string $a_N^2 a_1^2 a_2^2 \dots a_N^2$ occurs as a subword of a cyclically reduced form of g. In this case g is not simple (and hence not primitive) as its Whitehead graph does not have a cut vertex. We state this explicitly as a corollary of Proposition 2.4.4:

Corollary 2.4.6. *Let $F_N = F(A)$, where $N \geq 2$ and $A = \{a_1, \dots, a_N\}$. If a cyclically reduced word $w \in F(A)$ contains the subword $a_N^2 a_1^2 a_2^2 \dots a_N^2$ then w is not simple (and hence not primitive) in $F(A)$.*

The Whitehead graph, as defined above, records the information about two-letter subwords in the cyclically reduced form w of a nontrivial element $g \in F_N = F(A)$. There are also generalizations of the Whitehead graph recording the information about k-letter subwords of w, where $k \geq 2$ is a fixed integer. These generalizations are commonly known as "Rauzy graphs" or "initial graphs" and naturally occur in the study of geodesic currents on free groups [31, 32, 33].

We do not formally define these "level k" versions of the Whitehead graph here because we only need the following specific statement related to the $k = 3$ case:

Proposition 2.4.7. *[21] Let $F_N = F(A)$, where $N \geq 2$ and $A = \{a_1, \dots, a_N\}$. Let w be a nontrivial cyclically reduced word in $F(A)$ such that for every freely reduced word $v \in F(A)$ with $|v| = 3$ the word v occurs a subword of a cyclic permutation of w or of w^{-1}.*

Then w is filling in F_N (and, in particular, w is non-simple and non-primitive in F_N).

Chapter 3

Primitivity, Simplicity, and Non-Filling Index Functions

In this chapter we formally introduce the free group analogues of our surface group functions $\deg_{\Sigma,\rho}(\gamma)$, $f_{\Sigma,\rho}(L)$, and $\deg_{\Sigma,\rho}^{nfill}(\gamma)$, $f_{\Sigma,\rho}^{nfill}(L)$ (see Chapter 1).

3.1 Analogous Functions in Free Groups

In 1949 Marshall Hall Jr. proved in [30] that any finitely generated subgroup of a free group F_N is a free factor of a finite index subgroup of F_N. We state the result in a more precise form, as stated in [63]:

Proposition 3.1.1. *[63] Let $\alpha_1, \dots, \alpha_k, \beta_1, \dots, \beta_l$ be elements of a free group F_N. Let S be a subgroup of F_N generated by $\{\alpha_1, \dots, \alpha_k\}$. Suppose $\beta_i \notin S$ for $i = 1, \dots, l$. Then there exists a subgroup S' of finite index in F_N, such that $S \subset S'$, $\beta_i \notin S'$ for $i = 1, \dots, l$, and there exists a free basis of S' having a subset that is a free basis of S.*

If we pick $g \neq 1 \in F_N$ and apply the above result to the infinite cyclic subgroup $S = \langle g \rangle$, we get that there must exist a finite index subgroup S' of F_N such that g is a primitive element in S' (and hence $g \in S'$ is non-simple and non-filling in S').

This fact motivates the following definition:

Definition 3.1.2. [Primitivity, simplicity and non-filling indexes]

Let $N \geq 2$ be an integer and let F_N be a free group of rank N. Let $1 \neq g \in F_N$.

Define the *primitivity index* $d_{prim}(g) = d_{prim}(g, F_N)$ of g in F_N to be the smallest possible index for a subgroup $L \leq F_N$ containing g as a primitive element.

Define the *simplicity index* $d_{simp}(g) = d_{simp}(g, F_N)$ to be the smallest possible index for a subgroup $L \leq F_N$ containing g as a simple element.

Finally, define the *non-filling index* $d_{nfill}(g) = d_{nfill}(g, F_N)$ to be the smallest possible index for a subgroup $L \leq F_N$ containing g as a non-filling element.

As noted above, Proposition 3.1.1 implies that for every nontrivial $g \in F_N$ we have $d_{nfill}(g) \leq d_{simp}(g) \leq d_{prim}(g) < \infty$.

Definition 3.1.3 (Primitivity, simplicity and non-filling index functions)**.** Let F_N be a free group of rank $N \geq 2$ and let A be a free basis of F_N. For any $n \geq 1$ define the *primitivity index function* for F_N as:

$$f_{prim}(n) = f_{prim}(n; F_N) := \max_{\substack{1 \leq |g|_A \leq n,\, g \neq 1 \\ g \text{ not a proper power in } F_N}} d_{prim}(g)$$

Similarly, for $n \geq 1$ define the *simplicity index function* for F_N as:

$$f_{simp}(n) = f_{simp}(n; F_N) := \max_{\substack{1 \leq |g|_A \leq n,\, g \neq 1 \\ g \text{ not a proper power in } F_N}} d_{simp}(g)$$

Finally, for for $n \geq 1$ define the *non-filling index function* for F_N as:

$$f_{nfill}(n) = f_{nfill}(n; F_N) := \max_{\substack{1 \leq |g|_A \leq n,\, g \neq 1 \\ g \text{ not a proper power in } F_N}} d_{nfill}(g)$$

It is easy to see that the definitions of $f_{prim}(n; F_N)$, $f_{simp}(n; F_N)$ and $f_{nfill}(n; F_N)$ do not depend on the choice of a free basis A of F_N. Note that $f_{prim}(n)$ is the smallest monotone non-decreasing function such that for every non-trivial root-free $g \in F_N$ we have $d_{prim}(g) \leq f_{prim}(|g|_A)$; similar reformulations hold for $f_{simp}(n)$ and $f_{nfill}(n)$.

Again, the purpose behind defining these free group functions is to find analogues for the corresponding surface group functions. In the surface setting we are addressing the problem of "untangling" a closed geodesic in a finite cover, and we wish to quantify the degree of that cover in terms of the length of the geodesic. In free groups our notion of "untangling" is to make a group element primitive or simple or non-filling. We then want to quantify the smallest index of a subgroup in which our group element "untangles" in terms of the word length of our group element. So now, we will quantify these free group functions.

It turns out that finding upper bounds in this setting is quite straightforward. We recall the following well-known fact, which is Lemma 8.10 in [35]:

Lemma 3.1.4. *Let Γ be a finite folded A-graph. Then there exists a finite folded A-regular graph Γ' such that Γ is a subgraph of Γ' and such that $V\Gamma = V\Gamma'$.*

Proposition 3.1.5. *For every non-trivial cyclically reduced word $w \in F(A)$ of length n, there exists a finite index subgroup $H \leq F(A)$ of index n such that $w \in H$ is primitive in H.*

Proof. Take the word w of length n and write it on a circle of simplicial length n. Pick a vertex x as the base vertex. Call this graph (Γ_w, x). By Lemma 3.1.4 we can complete this graph to a finite cover (Γ'_w, x) of the N-rose without adding any extra vertices. Thus (Γ'_w, x) has n vertices and represents a subgroup H of F_N of index precisely n. The fact that w is realized as the label of a simple closed curve in (Γ'_w, x) implies that w is a primitive element in H. It is clear that $w \in H$ by definition of H. Note that since (Γ', x) has no extra vertices, a maximal tree T of (Γ, x) consists of all but one edge of the simple closed curve representing w. Let $e \in E_+\Gamma' - T$. Then $\mu(s_e) = w$ and hence w is primitive. See Figure 3.1 for a pictorial proof.

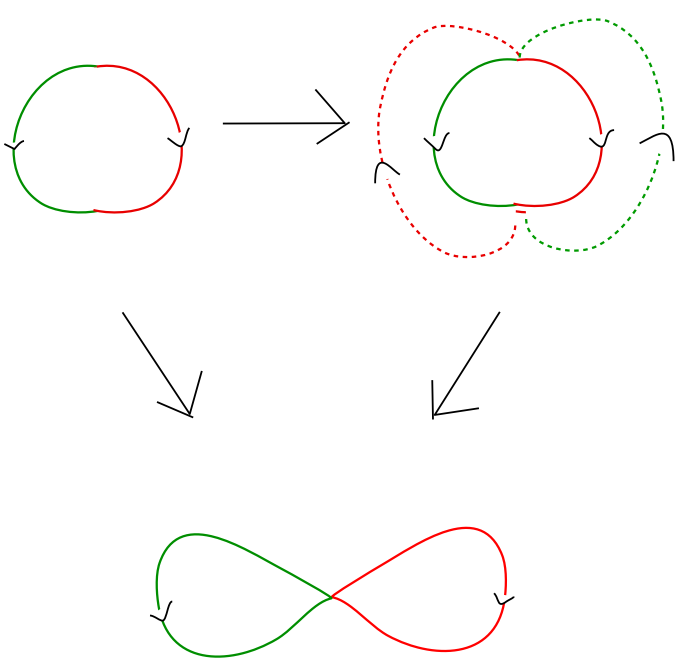

Figure 3.1: Proof by Picture for Proposition 3.1.5

□

Proposition 3.1.5, together with the definitions, directly implies:

Lemma 3.1.6. *Let $N \geq 2$ and let F_N be free of rank N. Then the following hold:*

1. *If $1 \neq g \in F_N = F(A)$ then*

$$d_{nfill}(g) \leq d_{simp}(g) \leq d_{prim}(g) \leq ||g||_A \leq |g|_A = n.$$

2. *For every $n \geq 1$ we have*

$$f_{nfill}(n) \leq f_{simp}(n) \leq f_{prim}(n) \leq n.$$

3. *Let* $1 \neq g \in F_N$ *and let* $\alpha \in Aut(F_N)$. *Then* $d_{prim}(g) = d_{prim}(\alpha(g))$, $d_{simp}(g) = d_{simp}(\alpha(g))$ *and* $d_{nfill}(g) = d_{nfill}(\alpha(g))$.

4. *If* $1 \neq g \in F_N$ *and* $k \geq 1$ *is an integer, then* $d_{simp}(g^k) \leq d_{simp}(g)$ *and* $d_{nfill}(g^k) \leq d_{nfill}(g)$.

In particular, part (3) of the above lemma shows that for g_1, g_2 conjugate non-trivial elements of F_N, we have $d_{prim}(g_1) = d_{prim}(g_2)$, $d_{simp}(g_1) = d_{simp}(g_2)$ and $d_{nfill}(g_1) = d_{nfill}(g_2)$.

As noted above, if $1 \neq g \in F_N$ and $k \geq 1$ is an integer, then $d_{simp}(g^k) \leq d_{simp}(g)$ and $d_{nfill}(g^k) \leq d_{nfill}(g)$. However, the function $d_{prim}(g)$ does not behave well under taking powers, as demonstrated by the following lemma:

Lemma 3.1.7. *For any* $a_i \in \{a_1, \dots, a_N\}$, *and any positive integer* n, $d_{prim}(a_i^n) = n$.

Proof. As noted above, for every nontrivial $g \in F_N$ we have $d_{simp}(g) \leq d_{prim}(g) \leq ||g||_A$. Thus $d_{prim}(a_i^n) \leq ||a_i^n||_A = n$. We need to show that $d_{prim}(a_i^n) \geq n$.

Let $d = d_{prim}(a_i^n)$ and let $H \leq F_N$ be a subgroup of index d such that $a_i^n \in H$ and that a_i^n is a primitive element of H. Let $(\Gamma, *)$ be the d-fold cover of R_N corresponding to H, so that for the covering map $p : \Gamma \to R_N$ have $\pi_1(\Gamma, *) \cong H$ and $p_\# = \mu : \pi_1(\Gamma, *) \to H \leq F_N = \pi_1(R_N, x_0)$ is an isomorphism.

The fact that $a_i^n \in H$ implies that there exists a reduced closed path γ from $*$ to $*$ in Γ with $\mu(\gamma) = a_i^n$. Since a_i^n is primitive in H, the element γ is primitive in $\pi_1(\Gamma, *)$.

Since a_i^n is cyclically reduced, the closed path γ is also cyclically reduced. We claim that γ is a simple closed path in Γ. Indeed, suppose not. Then $\gamma = \gamma_1^k$ where $k \geq 2$ and where γ_1 is a simple closed path at $*$ in Γ with label $a_i^{n/k}$. Therefore γ is a proper power in $\pi_1(\Gamma, *)$, which contradicts the fact that γ is primitive in $\pi_1(\Gamma, *)$. Thus indeed γ is a simple closed path in Γ with label a_i^n. This means that the full p-preimage of the i-th petal of R_N, labeled a_i, in Γ consists of $\geq n$ distinct topological edges. Therefore the degree d of the cover $p : \Gamma \to R_N$ satisfies $d \geq n$.

Thus $d = d_{prim}(a_i^n) \geq n$. Since we already know that $d_{prim}(a_i^n) \leq n$, it follows that $d_{prim}(a_i^n) = n$, as required. ☐

Avoiding the bad behavior of $d_{prim}(g)$ under taking powers of g, demonstrated by Lemma 3.1.7, is the main reason why in Definition 3.1.2 we take the maximum over all root-free nontrivial elements $g \in F_N$ with $|g|_A \leq n$ rather than over all nontrivial $g \in F_N$ with $|g|_A \leq n$.

3.2 Algorithmic computability of $d_{prim}(g)$, $d_{simp}(g)$, and $d_{nfill}(g)$

In this section we will establish algorithmic computability of $d_{prim}(g)$, $d_{simp}(g)$, and $d_{nfill}(g)$. Consequently, we will also establish the algorithmic computability of $f_{prim}(n)$, $f_{simp}(n)$, and $f_{nfill}(n)$.

We first need to recall some basic definitions and facts related to Whitehead automorphisms and Whitehead's algorithm. We only briefly cover this topic here and refer the reader for further details to [45, pp. 30-35] and to [48, 38, 33, 57] for some of the recent developments. As before, $F_N = F(A) = F(a_1, \ldots, a_N)$ is the free group of rank $N \geq 2$ with a free basis $A = \{a_1, \ldots, a_N\}$.

Definition 3.2.1 (Whitehead automorphisms)**.** A *Whitehead automorphism* τ of $F_N = F(A)$ with respect to A is an automorphism τ of $F(A)$ of one of the following types:

1. There exists a permutation t of $\Upsilon = A \sqcup A^{-1}$ such that $\tau|_\Upsilon = t$. In this case τ is called a *relabeling automorphism* or a *Whitehead automorphism of the first kind.*
2. There exists an element $a \in \Upsilon$ which we call the *multiplier* such that for any $x \in \Upsilon$, $\tau(x) \in \{x, xa, a^{-1}x, a^{-1}xa\}$. In this case τ is called a *Whitehead automorphism of the second kind.*

Note that since $\tau \in Aut(F(A))$, if τ is a Whitehead automorphism of the second kind with multiplier a, then $\tau(a) = a$. Also for any $a \in \Upsilon$, the inner automorphism corresponding to conjugation by a is a Whitehead automorphism of the second kind.

Definition 3.2.2 (Automorphically minimal and Whitehead minimal elements)**.** An element $g \in F(A) = F_N$ is *automorphically minimal* in $F(A)$ with respect to a basis A of F_N if, for every $\phi \in Aut(F(A))$ we have $||g||_A \leq ||\phi(g)||_A$.

An element $g \in F(A)$ is *Whitehead minimal* in $F(A)$ with respect to a free basis A if, for every Whitehead automorphism τ of $F(A)$ we have $||g||_A \leq ||\tau(g)||_A$.

Note that neither Whitehead automorphisms of the first kind nor inner automorphisms change the cyclically reduced length of an element.

The following proposition summarizes the key known facts regarding Whitehead's algorithm (see [68] for the original proof by Whitehead and see [45, Proposition 4.17] for a modern exposition):

Proposition 3.2.3 (Whitehead's Theorem)**.** *Let $N \geq 2$ and let $F_N = F(A)$ be free of rank N with a free basis A. Then:*

1. *An element $g \in F(A)$ is automorphically minimal in $F(A)$ with respect to a basis A if and only if g is Whitehead minimal in $F(A)$ with respect to A. (Hence $g \in F(A)$ is not automorphically minimal with respect to A if and only if there exists a Whitehead automorphism τ such that $||\tau(g)||_A < ||g||_A$).*

2. *Whenever $u, v \in F(A)$ are Whitehead minimal with respect to A such that the orbits $Aut(F(A))u = Aut(F(A))v$ (so that, in particular, $||u||_A = ||v||_A$), then there exists a sequence of Whitehead automorphisms $\tau_1, \ldots, \tau_m$ of $F(A)$ with respect to A such that $\tau_m ... \tau_1(u) = v$ and that $||\tau_i ... \tau_1(u)||_A = ||u||_A$ for $i = 1, ..., m$.*

Note that part (2) of Proposition 3.2.3 holds even if u, v are conjugate in $F(A)$ since conjugation by an element of $A^{\pm 1}$ is a Whitehead automorphism.

3.2.1 Algorithmic computability of $d_{prim}(g)$ and $d_{simp}(g)$

The following useful lemma explicitly states the relationship between primitivity, simplicity and Whitehead minimality:

Lemma 3.2.4. *Let $1 \neq w \in F(A) = F_N$.*

1. *w primitive in $F(A)$ if and only if every (equivalently, some) Whitehead minimal form $\widetilde{w}$ of w has $||\widetilde{w}||_A = 1$.*

2. *w is simple in $F(A)$ if and only if some Whitehead minimal form $\widetilde{w}$ of w misses an $a_i^{\pm 1}$.*

3. *w is simple in $F(A)$ if and only if every Whitehead minimal cyclically reduced form $\widetilde{w}$ of w misses an $a_i^{\pm 1}$.*

Proof. Part (1) of the lemma is well-known and follows directly from Proposition 3.2.3.

If some Whitehead minimal form $\widetilde{w}$ of w misses an $a_i^{\pm 1}$, then w is simple in $F(A)$ as $w \in F(B)$ where $B = A - \{a_i\}$ and $F(B)$ is a proper free factor of $F(A)$.

Conversely, suppose that w is simple in $F(A)$. Then there exists an automorphism ϕ of $F(A)$ such that the cyclically reduced form $\widehat{w}$ of $\phi(w)$ misses $a_N^{\pm 1}$.

Claim 1. We claim that some Whitehead minimal form of $\widehat{w}$ also misses $a_N^{\pm 1}$.

We prove this claim by induction on $||\widehat{w}||_A$. If $||\widehat{w}||_A = 1$, then the claim clearly holds. Suppose now that $||\widehat{w}||_A = m > 1$ and that the claim has been established for all nontrivial cyclically reduced words in $F(a_1, \ldots, a_{N-1})$ of length $\leq m - 1$.

If $\widehat{w}$ is already Whitehead minimal in $F(A)$ then we are done as the claim holds in this case.

If $\widehat{w}$ is not Whitehead minimal in $F(A)$ then there exists a Whitehead automorphism τ of $F(A)$ such that $||\tau(\widehat{w})||_A < ||\widehat{w}||_A$. Note first that since the cyclically reduced length of $\widehat{w}$ changes under τ, we must have that τ is a Whitehead automorphisms of the second kind that is not an inner automorphism.

Let $a \in \Upsilon = A \sqcup A^{-1}$ be the multiplier of τ. If $a = a_N^{\pm 1}$, since $\widehat{w}$ is a cyclically reduced word in $F(A)$ that misses the letter $a_N^{\pm 1}$, the definition of a Whitehead automorphism implies that there can be no cancellation in $\tau(\widehat{w})$ between the letters $\{a_1, \ldots, a_{N-1}\}$ when a cyclically reduced form of $\tau(\widehat{w})$ is computed. Hence $||\tau(\widehat{w})||_A \geq ||\widehat{w}||_A$, contrary to the fact that $||\tau(\widehat{w})||_A < ||\widehat{w}||_A$. Therefore $a \in \{a_1, \ldots, a_{N-1}\}^{\pm 1}$. We then define a Whitehead automorphism τ' of $F(a_1, \ldots, a_{N-1})$ with respect to $\{a_1, \ldots, a_{N-1}\}$ as $\tau' = \tau|_{\{a_1,\ldots,a_{N-1}\}}$. Hence $\tau(\widehat{w}) = \tau'(\widehat{w})$. Thus $\tau(\widehat{w})$ still misses $a_N^{\pm 1}$ and $||\tau(\widehat{w})||_A < ||\widehat{w}||_A = m$. Applying the inductive hypothesis to $\tau(\widehat{w})$, we conclude that some Whitehead minimal form $\widetilde{w}$ of $\tau(\widehat{w})$ in $F(A)$ misses $a_N^{\pm 1}$. Then $\widetilde{w}$ is also a Whitehead minimal form of $\hat{w}$, and Claim 1 is verified.

Thus we have established part (2) of the lemma.

To see that part (3) holds, note that if every Whitehead minimal cyclically reduced form $\widetilde{w}$ of w misses an $a_i^{\pm 1}$ then w is simple in $F(A)$.

Now suppose w is simple in $F(A)$. From (2) we know that there is a $\widetilde{w}$ Whitehead minimal cyclically reduced form of w that misses $a_N^{\pm 1}$. Let w' be another Whitehead minimal cyclically reduced form of w in $F(A)$. Then $Aut(F(A))w' = Aut(F(A))\widetilde{w}$, and so by part (2) of Proposition 3.2.3, there exists a sequence of Whitehead automorphisms $\tau_1, \ldots, \tau_m$ of $F(A)$ with respect to A such that $\tau_m ... \tau_1(\widetilde{w}) = w'$ and that $||\tau_i ... \tau_1(\widetilde{w})||_A = ||w'||_A$ for $i = 1, ..., m$.

For $j = 0, 1, \ldots, m$ denote $w_j = \tau_j ... \tau_1(\widetilde{w})$, where $w_0 = \widetilde{w}$.

Claim 2. We claim that for each $j = 0, \ldots, m$ the cyclically reduced form of w_j misses some $a_i^{\pm 1}$.

We will establish Claim 2 by induction on j.

If $j = 0$ then $w_0 = w$ and there is nothing to prove. Suppose now that $j \geq 1$ and that the claim has been verified for w_{j-1}.

Thus the cyclically reduced form of w_{j-1} misses some $a_i^{\pm 1}$. If τ_j is a Whitehead automorphism of the first kind, it is clear that the cyclically reduced form of $\tau_j(w_{j-1}) = w_j$ still misses some $a_k^{\pm 1}$ (this $a_k^{\pm 1}$ is not necessarily $a_i^{\pm 1}$). Suppose now that τ_j is a Whitehead automorphism of the second kind. The restriction that $||\tau_j(w_{j-1})||_A = ||w_{j-1}||$ forces the condition that either $\tau_j(w_{j-1})$ is equal to w_{j-1} after cyclic reduction, or else τ_j is a Whitehead automorphism of the second kind with multiplier $a \in B \sqcup B^{-1}$ where $B = \{x \in A \sqcup A^{-1} | x \text{ occurs in the cyclically reduced form of } w_{j-1}\}$ (in particular $a \neq a_i^{\pm 1}$). In both cases we see that the cyclically reduced form of w_j still misses $a_i^{\pm 1}$, as required. This completes the inductive step and the proof of Claim 2.

Applying Claim 2 with $j = m$ shows that the cyclically reduced form of $w' = w_m$ misses some $a_i^{\pm 1}$, and part (3) of the lemma is proved. □

Proposition 3.2.5. *Let $1 \neq g \in H \leq F(A)$, where H is a proper free factor of $F(A)$. Then the following hold:*

1. *The element g is primitive in H if and only if g is primitive in F_N.*

2. *There is an algorithm which decides, given $g \in F(A)$, whether or not $g \in F(A)$ is primitive.*

3. *There is an algorithm which given $g \in F(A)$, whether or not $g \in F(A)$ is simple.*

Proof. We first prove part (1). The "only if" direction is obvious. Thus we assume that $g \in H$ is primitive in F_N.

Let $K \leq F_N$ be such that $F_N = H * K$. Let $\mathcal{B}_H = \{h_1, \dots, h_l\}$ be a free basis for H, and $\mathcal{B}_K = \{k_1, \dots, k_m\}$ be a free basis for K. Then $\mathcal{B}_F = \{h_1, \dots, h_l, k_1, \dots, k_m\}$ is a free basis for F_N (here $l+m = n$).

Since $g \in H$, then g is a freely reduced word over $\mathcal{B}_H$, with cyclically reduced form w. We prove that g is primitive in H by induction on the length m of w.

If w has length 1, then g is primitive in H, as required. If w has length $m > 1$, then the fact that w is primitive in F_N implies that w is not Whitehead minimal in F_N with respect to the free basis $\mathcal{B}_F$ of F_N. Hence there exists a Whitehead automorphism τ of F_N with respect to $\mathcal{B}_F$ such that $||\tau(w)||_{\mathcal{B}_F} < m$. By the same argument as in the proof of Lemma 3.2.4, we see that there exists a Whitehead automorphism τ' of $H = F(\mathcal{B}_H)$ such that $\tau'(w) = \tau(w)$. Then $\tau(w) = \tau'(w) \in H$ is primitive in F_N with $||\tau(w)||_{\mathcal{B}_F} < m$. Therefore by the inductive hypothesis the element $\tau(w) = \tau'(w)$ is primitive in H. Since $\tau' \in Aut(H)$, it follows that w is also primitive in H, as required. Thus part (1) of the proposition holds.

To prove parts (2) and (3) for $g \in F(A) = F(a_1, \dots, F_N)$, we find a Whitehead minimal form $\widetilde{g}$ in $F(A)$. By part (1) of Lemma 3.2.4, $||\widetilde{g}||_A = 1$ if and only if g is primitive in $F(A)$. By part (3) of Lemma 3.2.4, $\widetilde{g}$ misses some $a_i^{\pm 1}$ if and only if w is simple in $F(A)$. □

Remark 3.2.6. The algorithm described in part (2) of Proposition 3.2.5 is due to Whitehead [68]. The first algorithms for deciding whether an element of F_N is simple in F_N were provided by Stallings [64] and Stong [65] in 1990s. Their algorithms are somewhat different from the algorithm given in part (3) of Proposition 3.2.5 above, but they are also based on using Whitehead's algorithm.

Definition 3.2.7 (Principal quotient)**.** Following the terminology of [35], for a finite connected A-graph Γ_1 and a folded A-graph Γ_2, we say that Γ_2 is a *principal quotient* of Γ_1 if there exists a surjective A-morphism $\Gamma_1 \to \Gamma_2$.

Definition 3.2.8. Let $w \in F_N = F(A)$ be a nontrivial cyclically reduced word. We denote by C_w the A-graph which is a simplicial circle subdivided into $n = ||w||_A$ topological edges, such that the label of the closed path of length n corresponding to going around this circle once from some vertex $*$ to $*$ is the word w.

By definition, the graph C_w has a distinguished base-vertex $*$. Thus a principal quotient of C_w also come equipped with a distinguished base-vertex. We say that (Γ, x) is a *principal quotient* of C_w if Γ is a finite connected folded A-graph, if $x \in V\Gamma$ and if there exists a surjective A-morphism $f : C_w \to \Gamma$ such that $f(*) = x$.

Note that if (Γ, x) is a principal quotient of C_w, then there exists a unique path $\gamma_{w,x}$ in Γ starting with x and with label w, and, moreover, this path is closed and passes through every topological edge of Γ.

The following lemma is an immediate corollary of the definitions:

Lemma 3.2.9. *The following hold:*

1. *Let Γ_1 be a finite connected A-graph and Γ_2 be a finite folded A-graph. Then Γ_2 is a principal quotient of Γ_1 if and only if Γ_2 can be obtained from Γ_1 by the following procedure: choose some partition $V\Gamma_1 = V_1 \sqcup \cdots \sqcup V_m$ (with all $V_i \neq \emptyset$), then for each $i = 1, \ldots, m$ collapse V_i to a single vertex to get an A-graph Γ_1', and then fold the graph Γ_1' to obtain Γ_2.*

2. *If $w \in F_N = F(A)$ is a nontrivial cyclically reduced word and Γ is a finite connected folded A-graph, then Γ is a principal quotient of C_w if and only if Γ' is a core graph and there exists a closed path γ_w in Γ with label w such that γ_w passes through every topological edge of Γ'.*

A priori it is unclear that the functions $f_{prim}(n)$ and $f_{simp}(n)$ are even computable for a given F_N. We now give an algorithm that calculates $d_{prim}(g)$ and $d_{simp}(g)$ for any non-trivial g. This would then show that the functions $f_{prim}(n)$ and $f_{simp}(n)$ are indeed algorithmically computable.

Definition 3.2.10. Let $1 \neq g \in F_N = F(A)$ and let $w \in F(A)$ be the cyclically reduced form of g. We denote by $\mathcal{G}_0(w)$ the set of all finite connected folded basepointed A-graphs (Γ, x) such that there exists a closed path γ from x to x labeled w with the property that γ passes through every topological edge of Γ at least once and such that either the labeling map $\Gamma \to R_N$ is not a covering (that is, there exists a vertex of Γ of degree $< 2N$), or the labeling map $\Gamma \to R_N$ is a covering and the element $\gamma \in \pi_1(\Gamma, x)$ is simple in $\pi_1(\Gamma, x)$.

We denote by $\mathcal{G}(w)$ the set of all finite connected folded basepointed A-graphs (Γ, x) such that there exists a closed path γ from x to x labeled w with the property that γ passes through every topological edge of Γ at least once and such that the element $\gamma \in \pi_1(\Gamma, x)$ is primitive in $\pi_1(\Gamma, x)$.

Let $(\Gamma, x) \in \mathcal{G}(w)$ or $(\Gamma, x) \in \mathcal{G}_0(w)$. Since w is cyclically reduced and γ passes through every topological edge of Γ at least once, every vertex of Γ has degree ≥ 2, so that Γ is a core graph.

Note further that the condition that γ is simple in $\pi_1(\Gamma, x)$ is equivalent to the condition that w is simple in the subgroup $H \leq F_N$ represented by (Γ, x). This follows from the fact that the labeling map gives an isomorphism $\mu : \pi_1(\Gamma, x) \to H$, with $\mu(\gamma) = w$.

We recall the following basic fact:

Lemma 3.2.11 ([35], p.13)**.** *Let Γ be a folded connected A-graph and let Γ' be a connected subgraph of Γ. Let $*$ be a vertex of Γ'. If $H' \leq F(A)$ is the subgroup represented by $(\Gamma', *)$ and H is the subgroup represented by $(\Gamma, *)$, then H' is a free factor of H.*

Remark 3.2.12. In the setting of Lemma 3.2.11, $\pi_1(\Gamma', *)$ is a free factor of $\pi_1(\Gamma, *)$.

Proposition 3.2.13. *Let $1 \neq g \in F_N = F(A)$ and let $w \in F(A)$ be the cyclically reduced form of g. Then the following hold:*

1. *The number $d_{prim}(g)$ equals to the minimum of $\#V\Gamma$, taken over all $(\Gamma, x) \in \mathcal{G}(w)$.*

2. *The number $d_{simp}(g)$ equals to the minimum of $\#V\Gamma$, taken over all $(\Gamma, x) \in \mathcal{G}_0(w)$.*

Proof. We give a proof of part (2). The proof of part (1) is very similar in nature. However, it additionally involves using part (1) of Proposition 3.2.5 to prove one of the inequalities. For $1 \neq g \in F_N = F(A)$ and $w \in F(A)$ the cyclically reduced form of g, let $\overline{d_{simp}}(g) = \min_{(\Gamma,x) \in \mathcal{G}_0(w)} \#V\Gamma$. First suppose that $H \leq F_N$ such that $[F_N : H] = d_{simp}(g) = d_{simp}(w)$, and that $w \in H$ is simple in H. Let (Γ, x) be the graph representing H as in Proposition-Definition 2.2.7. We have that $\#V\Gamma = d_{simp}(w)$. Since $w \in H$, there exists a path γ from x to x in Γ with label w. Also since $w \in H$ is simple in H, $\gamma \in \pi_1(\Gamma, x)$ is simple in $\pi_1(\Gamma, x)$. Let $\Gamma' \subseteq \Gamma$ be the subgraph spanned by γ. Then γ is a path from x to x in Γ' that passes through every topological edge in Γ' at least once. If $\Gamma' = \Gamma$, then the labeling map $\Gamma' \to R_N$ is a covering. Since γ is simple in $\Gamma = \Gamma'$, we have $(\Gamma', x) \in \mathcal{G}_0(w)$. Since $\#V\Gamma' = \#V\Gamma = d_{simp}(g)$, we have that $\overline{d_{simp}}(g) \leq d_{simp}(g)$. If $\Gamma' \neq \Gamma$, then $\#V\Gamma' \leq \#V\Gamma$ and $\#E\Gamma - \#E\Gamma' \geq 1$. From Remark 3.2.12, (Γ', x) is a proper free factor of (Γ, x). In this case the labeling map $\Gamma' \to R_N$ is not a covering and $(\Gamma', x) \in \mathcal{G}_0(w)$. Thus $\overline{d_{simp}}(g) \leq d_{simp}(g)$.

Conversely suppose that $(\Gamma, x) \in \mathcal{G}_0(w)$ with $\#V\Gamma = \overline{d_{simp}}(g)$. Let γ be the closed path from x to x labeled by w such that γ passes through every topological edge of Γ at least once. If the labeling map $\Gamma \to R_N$ is a covering then $\gamma \in \pi_1(\Gamma, x)$ is simple in $\pi_1(\Gamma, x)$ by definition of $\mathcal{G}_0(w)$. Let H be the subgroup represented by (Γ, x). H is then a subgroup of F_N of index $\overline{d_{simp}}(g)$ with $w \in H$ and w simple in H. Hence $d_{simp}(g) = d_{simp}(w) \leq \overline{d_{simp}}(g)$. If the labeling map $\Gamma \to R_N$ is not a covering, we use Lemma 3.1.4 to complete (Γ, x) to a finite cover $(\widehat{\Gamma}, x)$ of R_N without adding any extra vertices and by adding at least one edge. Again from Remark 3.2.12, (Γ, x) is a proper free factor of $(\widehat{\Gamma}, x)$. Hence $\gamma \in \pi_1(\widehat{\Gamma}, x)$ is simple in

$\pi_1(\widehat{\Gamma}, x)$. Let H be the subgroup represented by $(\widehat{\Gamma}, x)$. We have shown that $w \in H$ is simple in H. Since $\#V\widehat{\Gamma} = \#V\Gamma = \overline{d_{simp}}(g)$, we see that $d_{simp}(g) \leq \overline{d_{simp}}(g)$. □

We can now prove:

Theorem 3.2.14. *Let $F_N = F(A)$, where $N \geq 2$ and where $A = \{a_1, \dots, a_N\}$ is a free basis of F_N. Then:*

1. *There exists an algorithm that, given $1 \neq g \in F_N$, computes $d_{prim}(g)$ and $d_{simp}(g)$.*
2. *There exists an algorithm that, for every $n \geq 1$ computes $f_{prim}(n)$ and $f_{simp}(n)$*

Proof. Let $1 \neq g \in F_N$ and let w be the cyclically reduced form of g. Note that a finite connected folded base-pointed A-graph (Γ, x) admits a closed path γ from x to x labeled w and passing through every topological edge of Γ at least once if and only if (Γ, x) is a principal quotient of C_w with x being the image of the base-vertex $*$ of C_w.

Therefore we can algorithmically find all the graphs in $\mathcal{G}_0(w)$ as follows: List all partitions on VC_w. For each partition of VC_w as a disjoint union of nonempty subsets $V_1, \dots V_m$, collapse V_i to a single vertex for $i = 1, \dots, m$, and fold the resulting graph to obtain a principal quotient (Γ, x) of C_w, with x being the image of the base-vertex $*$ of C_w. Let γ be the path from x to x in Γ labeled w (so that, by construction, γ passes through every topological edge of Γ at least once). Then check whether the labeling map $\Gamma \to R_N$ is a covering, that is, whether it is true that every vertex of Γ has degree $2N$. If $\Gamma \to R_N$ is not a covering, the graph (Γ, x) belongs to $\mathcal{G}_0(w)$. If $\Gamma \to R_N$ is a covering, check, using the algorithm from part (3) of Proposition 3.2.5, whether or not $\gamma \in \pi_1(\Gamma, x)$ is simple in the finite rank free group $\pi_1(\Gamma, x)$. If $\gamma \in \pi_1(\Gamma, x)$ is simple in $\pi_1(\Gamma, x)$, we conclude that the graph (Γ, x) belongs to $\mathcal{G}_0(w)$, and $\gamma \in \pi_1(\Gamma, x)$ is not simple in $\pi_1(\Gamma, x)$, we conclude that he graph (Γ, x) does not belong to $\mathcal{G}_0(w)$. Performing this procedure for each partition of VC_w as a disjoint union of nonempty subsets produces the finite set $\mathcal{G}_0(w)$. Proposition 3.2.13 then implies that $d_{simp}(g) = d_{simp}(w) = \min\{\#V\Gamma : (\Gamma, x) \in \mathcal{G}_0(w)\}$.

The algorithm for computing $d_{prim}(g) = d_{prim}(w)$ is similar. We first find all the graphs in $\mathcal{G}(w)$ as follows. Enumerate all partitions of VC_w as a disjoint union of nonempty subsets. For each such partition $V_1, \dots V_m$ collapse each V_i, $i = 1, \dots, m$, to a vertex and then fold the result to get a principal quotient (Γ, x) of C_w. There is a path γ from x to x in Γ labeled w. Then check, using the algorithm from part (2) of Proposition 3.2.5,, whether or not $\gamma \in \pi_1(\Gamma, x)$ is primitive in the free group $\pi_1(\Gamma, x)$. If yes, we conclude that $(\Gamma, x) \in \mathcal{G}(w)$ and if not, we conclude that $(\Gamma, x) \notin \mathcal{G}(w)$. This procedure algorithmically computes the set $\mathcal{G}(w)$.

Proposition 3.2.13 then implies that $d_{prim}(g) = d_{prim}(w) = \min\{\#V\Gamma : (\Gamma, x) \in \mathcal{G}(w)\}$. Thus part (1) of the theorem is verified.

Part (2) now follows directly from part (1) using the definitions of $f_{prim}(n)$ and $f_{simp}(n)$.

$\square$

Remark 3.2.15. The complexity of the algorithms for computing $d_{simp}(g)$ and $d_{prim}(g)$ given in part (1) of Theorem 3.2.14 is super-exponential in $n = ||g||_A$. The reason is that enumerating all principal quotients of the graph C_w requires listing all partitions of the n-element set VC_w. The *Bell number* B_n, which is the number of all partitions of an n-element set, grows roughly as n^n.

3.2.2 Algorithmic computability of $d_{nfill}(g)$

We now want to give an algorithm for computing $d_{nfill}(g)$. Computationally this algorithm is not nearly as nice as the algorithms for computing $d_{simp}(g)$ and $d_{prim}(g)$ described above.

The following result provides a useful characterization of filling elements:

Proposition 3.2.16. *Let $1 \neq g \in F_N$. Then $g \in F_N$ is filling if and only if $Stab_{Out(F_N)}([g])$ is finite.*

Proof. Solie [61, Lemma 2.42, Lemma 2.44] proves that if $g \in F_N$ is non-filling then $Stab_{Out(F_N)}([g])$ is infinite. Thus the the "if" direction of Proposition 3.2.16 holds.

Let us now prove the "only if" direction. Suppose $Stab_{Out(F_N)}([g])$ is infinite. Choose a basepoint $[T_0] \in \mathrm{CV}_N$. Since the action of $Out(F_N)$ on CV_N is properly discontinuous and since $\overline{\mathrm{CV}}_N$ is compact, it follows that there exist an infinite sequence of distinct elements $\phi_n \in Stab_{Out(F_N)}([g])$ and a point $[T] \in \overline{\mathrm{CV}}_N - \mathrm{CV}_N$ such that $\lim_{n\to\infty} [T_0]\phi^n = [T]$. Then for some sequence of scalars $c_n \geq 0$ with $c_n \to 0$ as $n \to \infty$ we have $\lim_{n\to\infty} c_n T_0 \phi_n = T$ in $\overline{\mathrm{cv}}_N$. Since $\phi_n([g]) = [g]$, it follows that $||g||_T = \lim_{n\to\infty} c_n ||\phi_n(g)||_{T_0} = 0$. Then by Proposition 2.3.4 the element g is not filling in F_N, as required. $\square$

Proposition 3.2.17. *Let $F_N = F(A)$, where $N \geq 2$ and where $A = \{a_1, \dots, a_N\}$ is a free basis of F_N. Then there exists an algorithm that, given a nontrivial element $g \in F_N$, decides whether or not g is filling in F_N.*

Proof. Let $g \in F_N = F(A)$ be a nontrivial freely reduced word. By a result of McCool [46] the group $Stab_{Out(F_N)}([g])$ is finitely generated and, moreover, we can algorithmically compute a finite generating set $Y = \{\psi_1, \dots, \psi_k\}$ of $Stab_{Out(F_N)}([g])$.

In view of Proposition 3.2.16 we next need to determine if $H := \langle Y \rangle \leq Out(F_N)$ is finite. Wang and Zimmermann [67] prove that for $N > 2$, the maximum order of a finite subgroup of $Out(F_N)$ is $2^N N!$. Also, the word problem for $Out(F_N)$ is solvable (even solvable in polynomial time [58]). Thus we then start building the Cayley graph $Cay(H; Y)$ of H with respect to Y. Using solvability of the word problem in

$Out(F_N)$, for any finite m we can algorithmically construct the ball $B(m)$ of radius m cantered at identity in $Cay(H;Y)$. We construct the balls $B(2^N N!)$ and $B(1+2^N N!)$. By the result of Wang and Zimmermann mentioned above, the group H is finite if and only if $B(2^N N!) = B(1+2^N N!)$.

Thus we can algorithmically decide whether or not $Stab_{Out(F_N)}([g])$ is finite, and hence, by Proposition 3.2.16, whether or not g is filling in F_N.

□

Theorem 3.2.18. *Let $F_N = F(A)$, where $N \geq 2$ and where $A = \{a_1, \dots, a_N\}$ is a free basis of F_N. Then:*

1. *There exists an algorithm that, given $1 \neq g \in F_N$, computes $d_{nfill}(g)$.*

2. *There exists an algorithm that, for every $n \geq 1$ computes $f_{nfill}(n)$*

Proof. Part (2) follows directly from part (1) and from the definition of $f_{nfill}(n)$. Thus we only need to establish part (1).

Given $g \in F_N$, let w be the cyclically reduced form of g. Let C_w and its principle quotients be as in Definitions 3.2.8, 3.2.7. Enumerate all principle quotients of C_w as $\{\Gamma_1, \dots, \Gamma_k\}$. For each Γ_i with $1 \leq i \leq k$, two possibilities arise:

Case (i) (Γ_i *is not a finite cover of* R_N): In this case, we call Γ_i a "success". In this case we can complete Γ_i to a finite cover Γ'_i of R_N and now $\pi_1(\Gamma_i)$ is a free factor of $\pi_1(\Gamma'_i)$. Hence w is simple in in the subgroup represented by Γ'_i i.e. w is not filling in the subgroup represented by Γ'_i.

Case (ii) (Γ_i *is a finite cover of* R_N): In this case there is a closed loop γ_i in Γ_i with label w. We then use the algorithm from Proposition 3.2.17 to check whether γ_i is filling in $\pi_1(\Gamma_i)$. If γ_i is not filling in $\pi_1(\Gamma_i)$, and we call Γ_i a success.

Finally, observe that $d_{nfill}(g) = \min\{V\Gamma_i \,|\, \Gamma_i \text{ is a "success"}\}$ where this equality is established in a manner similar to that in Proposition 3.2.13. Thus part (1) of the theorem is proved. □

Chapter 4

A Lower Bound for the Free Group Index Functions

In this Chapter we will prove Theorem A.

4.1 Special words and finite covers

The main goal of this section is to find a suitable sufficient condition implying that a given freely reduced word is filling in a given finite index subgroup of F_N represented by a finite cover of the rose R_N. Similarly we find a suitable sufficient condition implying that a given freely reduced word is not simple in a subgroup of F_N represented by a given finite cover of the rose R_N.

These goals are accomplished by constructing "simplicity blocking" and "filling forcing" words in F_N of controlled length, provided by Proposition 4.1.12 and Proposition 4.1.7 below. Since the proofs of these Propositions are somewhat technical, we first illustrate the idea of their proof by obtaining a related simpler statement, given in Lemma 4.1.1 below. The proof of Lemma 4.1.1 is due to Yuliy Baryshnikov. We then adapt the idea of this proof to obtain Proposition 4.1.7 and Proposition 4.1.12.

Lemma 4.1.1. *Let $N \geq 2$. Then there exists a constant $c_0 = c_0(N) > 0$ with the following property. Let $(\Gamma, *)$ be a connected d-fold cover of the N-rose R_N, where $d \geq 1$. Then there exists a freely reduced word $v = v(\Gamma)$ with $|v| \leq c_0 d^2$ such that for every vertex $x \in V\Gamma$ the path $p(x, v)$ from x labeled by v in Γ passes through every topological edge of Γ at least once.*

Proof. The graph Γ is a connected $2N$-regular graph with d vertices and Nd topological edges. We can view Γ as a directed graph where the directed edges are labeled by elements of A (and without using A^{-1}). Then Γ is a connected directed graph where the in-degree of every vertex is equal to N, which is also equal to the out-degree of every vertex. Hence there exists an Euler circuit in Γ beginning and ending at $*$ consisting of edges labeled by elements of A that transverses each topological edge exactly once. Let v_1 be the label of this Euler circuit. Then v_1 is freely reduced and no a_i^{-1} occurs in v_1 for $i = 1, \dots, N$. Enumerate the vertices as $V\Gamma = \{x_1, x_2, \dots, x_d\}$ with $* = x_1$. Starting at the vertex x_2 follow a path p_1 with label v_1. Denote the terminal vertex of p_1 by z_1. Let p_1' be an Euler circuit in Γ starting and ending at z_1 and consisting only of edges labeled by elements of A. Let v_2 be the label of this path p'_1. Note that since we only consider positively labeled edges, the path $p_2 = p_1 p'_1$ is reduced and its label $v_1 v_2$ is a positive (and hence freely reduced) word over A. We now inductively define a positive word v_i over A given that the positive words $v_1, \ldots, v_{i-1}$ where $i \in \{1, \ldots, d\}$ have already been defined. Starting at vertex x_i we follow a path p_{i-1} with label $v_1 \ldots v_{i-1}$. Denote the terminal vertex of the path p_{i-1} by z_{i-1}. Let p'_{i-1} be an Euler circuit at z_{i-1} that transverses every positively labeled edge exactly once. Let v_i be the label of this path p'_{i-1}. This process is illustrated below:

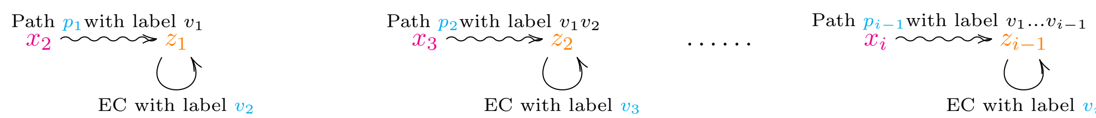

We define our word $v := v_1 v_2 \ldots v_d$. Since following a path with label $v_1 \ldots v_i$ at any vertex v_i already passes through every topological edge of Γ at least once, so does following a path with label v. Since each $|v_i| = Nd$ for $1 = 1, \ldots, d$, we have that $|v| = Nd^2$. □

4.1.1 Simplicity blocking words and finite covers

In the above proof the concatenation argument always produces reduced edge-paths because we only deal with edges and paths labeled by positive words over A. By contrast, in proving Proposition 4.1.7 simple concatenation does not always work as it may result in paths that are not reduced. Also, instead of paths labeled by v passing through every edge of Γ, we need to ensure a more complicated condition which implies that all paths labeled by v in Γ pass through a certain "simplicity-blocking" path $\alpha(\Gamma, T)$, which is defined below.

Definition 4.1.2. Let Γ be a finite connected folded A-graph, let $T \subseteq \Gamma$ be a maximal tree in Γ and let S_T be the corresponding basis of $\pi_1(\Gamma, *)$. Let $u = y_1 \ldots y_n$ be a nontrivial freely reduced word over $S_T^{\pm 1}$. Thus each y_i corresponds to an edge $e_i \in E(\Gamma - T)$. We define a reduced path $\delta(u)$ in Γ as

$$\delta(u) := [*, o(e_1)]_T e_1 [t(e_1), o(e_2)]_T e_2 \ldots \ldots e_n [t(e_n), *]_T.$$

Note that if $d = \# V\Gamma$ then T has $\leq d-1$ topological edges and hence $|\delta(u)| \leq n + (n+1)(d-1) = nd + d - 1 = d(n+1) - 1$.

Definition 4.1.3. Let $(\Gamma, *)$ be a finite folded core graph with a base-vertex $*$. Let $T \subseteq \Gamma$ be a maximal subtree in Γ. Let $S_T = \{b_1, \ldots, b_r\}$ be the basis of $\pi_1(\Gamma, *)$ dual to T.

Define a reduced edge-path $\alpha(\Gamma, T)$ from $*$ to $*$ in Γ as

$$\alpha(\Gamma, T) := \delta(b_r^2 b_1^2 \dots b_r^2).$$

Remark 4.1.4. Note that the path $\alpha(\Gamma, T)$ is reduced and represents the element $b_r^2 b_1^2 \dots b_r^2$ in $\pi_1(\Gamma, *)$. The following proposition demonstrates the "simplicity-blocking" property of $\alpha(\Gamma, T)$. The word $b_r^2 b_1^2 \dots b_r^2$ has length $2r + 2$ and hence $|\alpha(\Gamma, T)| \leq d(2r + 3) - 1$ where $d = \#V\Gamma$. In particular, if Γ is a d-fold cover of the rose R_N, then $r = d(N - 1) + 1$ and

$$|\alpha(\Gamma, T)| \leq d(2d(N - 1) + 3) - 1 \leq 2d^2(N - 1) + 4d.$$

Proposition 4.1.5. *Let Γ be as in Definition 4.1.3 with T a maximal tree in Γ. Let S_T and $\alpha(\Gamma, T)$ be as before. Let $\gamma \in \pi_1(\Gamma, *)$ be such that γ is represented by a cyclically reduced circuit in Γ containing $\alpha(\Gamma, T)$ as a subpath. Then γ is not simple in $\pi_1(\Gamma, *)$.*

Proof. We first use Proposition 2.1.5 to rewrite γ as a cyclically reduced word w in $S_T = \{b_1, \dots, b_r\}$. Then the occurrence of $\alpha(\Gamma, T)$ in γ produces an occurrence of the reduced word $b_r^2 b_1^2 \dots b_r^2$ in w. Hence, by Corollary 2.4.6, in this case γ is not simple in $F(b_1, \dots, b_r) = \pi_1(\Gamma, *)$.

□

Lemma 4.1.6. *Let Γ be a finite connected core graph with d vertices. Suppose that $\pi_1(\Gamma)$ has rank ≥ 2. Then for any any two edges $e_1, e_2 \in E(\Gamma)$, there exists a reduced path $p(e_1, e_2)$ starting at e_1, ending at e_2, and with $|p(e_1, e_2)| \leq 3d$.*

Proof. Pick a graph $\Gamma' \subseteq \Gamma$ such that Γ' is a finite, connected, core graph with $\pi_1(\Gamma')$ of rank 2 and $e_1, e_2 \in E\Gamma'$. Then there are precisely three possibilities for Γ'. It can be the wedge of two circles, or a theta-graph (a circle with a line segment joining two points on the circle), or a barbell graph (two circles attached to two ends of a line segment). We will show that the result holds for the graph Γ', and hence holds for our graph Γ. Our proof is essentially going to be a proof by picture for each of these three cases. In Figure 4.1, green edges (or arrows) indicate e_1 and blue edges (or arrows) indicate e_2. We indicate the path $p(e_1, e_2)$ in red with the $\bullet$ representing the starting point of $p(e_1, e_2)$ and the $\rightarrow$ representing the direction. The path $p(e_1, e_2)$ starts at $o(e_1)$ and ends at $t(v_1)$. We call a "cusp" any vertex that is at the intersection of maximal arcs. The idea behind finding this path $p(e_1, e_2)$ is always to travel along e_1 to the nearest cusp. Then if one is required to go back on the same path one has already been on to get to e_2, one instead travels along a disjoint loop at the cusp. Now one can go back to e_2 and the path $p(e_1, e_2)$ will be reduced. If after traveling from e_1 to the cusp one can get to e_2 without compromising the fact that the path $p(e_1, e_2)$ is reduced, then one simply goes to e_2 and the path $p(e_1, e_2)$ so obtained is reduced. From Figure 4.1 we see that the result holds.

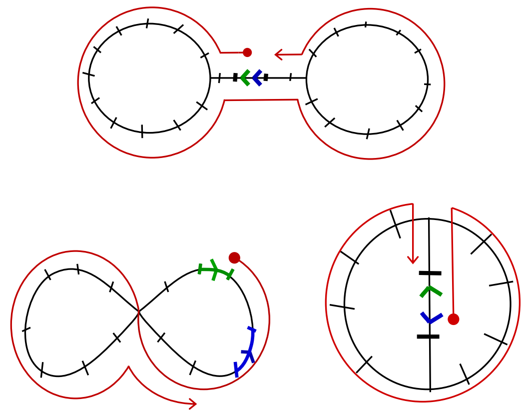

Figure 4.1: Proof by picture for Lemma 4.1.6

□

The following fact plays a key role in the proof of Theorem B:

Proposition 4.1.7. *Let $N \geq 2$. Then there exists a constant $c_0 = c_0(N) > 0$ with the following property: Let $(\Gamma, *)$ be a connected d-fold cover of the N-rose R_N, where $d \geq 1$ and let $T \subseteq \Gamma$ be a maximal subtree of Γ. Then there exists a freely reduced word $v = v(\Gamma, T)$ with $|v| \leq c_0 d^3$ such that for every vertex $x \in V\Gamma$ the path $p(x, v)$ from x labeled by v in Γ contains $\alpha(\Gamma, T)$ as a subpath.*

Proof. Let us begin by enumerating the vertices of $V\Gamma = x_1, x_2, \ldots, x_d$. Let $H \leq F_N$ be the subgroup of index d that is represented by $(\Gamma, *)$. For a maximal tree T in $(\Gamma, *)$ and let $S_T = \{b_1, b_2, \ldots, b_r\}$ be the corresponding basis of $\pi_1(\Gamma, *)$. By Remark 4.1.4, we have $|\alpha(\Gamma, T)| \leq 2d^2(N-1) + 4d$.

Let e be the first edge of the path $\alpha(\Gamma, T)$. Starting at the vertex $x_1 \in V\Gamma$, there exists a unique path $[x_1, *]_T$ of length $\leq d - 1$ with terminal edge e_1 (say). Lemma 4.1.6 then gives us a reduced path $p(e_1, e) = e_1 p' e$ of length $\leq 3d$. Let the word v_1 be the label of the path $p_1 = [x_1, *]_T p' \alpha(\Gamma, T)$. Note that $|v_1| = |p_1| \leq 2d^2(N-1) + 8d - 3$.

Starting at the vertex x_2 we follow a path p_1' that has label v_1. Let e_2 be the terminal edge of the path p_1'. Then from Lemma 4.1.6, the path $p(e_2, e) = e_2 p_1'' e$ is reduced with $|p(e_2, e)| \leq 3d$, and hence $|p_1''| \leq 3d - 2$. Let the word v_2 be the label of the path $p_2 = p_1'' \alpha(\Gamma, T)$. Now the path $p_1' p_2 = p_1' p_1'' \alpha(\Gamma, T)$ is reduced.

Notice that $|v_2| = |p_2| \leq 2d^2(N-1) + 7d - 1$. We now define inductively a sequence of words and paths as follows: Suppose we have already defined our words $v_1, v_2, \ldots, v_{i-1}$ which are respectively the labels of reduced paths $p_1, \ldots, p_{i-1}$. Starting at vertex x_i we follow the path p'_{i-1} labeled by the word $v_1 v_2 \ldots v_{i-1}$. Let e_i be the terminal edge of the path p'_{i-1}. Then the path $p(e_i, e) = e_i p''_{i-1} e$ is reduced with $|p''_{i-1}| \leq 3d-2$. Let the word v_i be the label of the reduced path $p_i = p''_{i-1}\alpha(\Gamma, T)$. Now the path $p'_{i-1}p_i = p'_{i-1}p''_{i-1}\alpha(\Gamma, T)$ is reduced. Let the word $v := v_1 v_2 \ldots v_d$. Then notice that at any vertex x_i with $1 \leq i \leq d$, the path $p'_{i-1}p_i$ is a reduced path labeled by $v_1 \ldots v_i$ that already contains the subpath $\alpha(\Gamma, T)$. Thus for $i = 1, \ldots, d$ the path starting at x_i labeled by the word $v_1 \ldots v_d$ also contains the subpath $\alpha(\Gamma, T)$. Since for all $2 \leq i \leq d$, $|v_i| \leq 2d^2(N-1) + 7d - 1$, we have that $|v| \leq 2d^3(N-1) + 7d^2 - 2 \leq (2N+5)d^3$. Thus with $c_0 = 2N + 5$, we are done.

□

The freely reduced word $v = v(\Gamma, T)$ in $F(A)$ can be viewed as a "simplicity blocking" word for the elements of the fundamental group of a d-fold cover Γ of R_N.

Corollary 4.1.8. *Let $N \geq 2$ and let $c_0 = c_0(N) > 0$ be the constant provided by Proposition 4.1.7.*

Let $d \geq 1$, let Γ be a connected d-fold cover of the N-rose R_N and let $T \subseteq \Gamma$ be a maximal tree in Γ. Let $* \in V\Gamma$, let γ be a reduced edge-path from $*$ to $*$ in Γ and let γ' be the cyclically reduced form of the path γ (so that the label of γ' is a cyclically reduced word in $F(A)$). Suppose that the label of γ' contains as a subword the word $v = v(\Gamma, T)$ with $|v| \leq c_0 d^3$ provided by Proposition 4.1.7.

Then $\gamma \in \pi_1(\Gamma, *)$ does not belong to a proper free factor of $\pi_1(\Gamma, *)$.

Proof. From definitions $\gamma \in \pi_1(\Gamma, *)$. Using the tree T we can obtain a free basis $S_T = \{b_1, \ldots, b_r\}$ of $\pi_1(\Gamma, T)$. Then Proposition 2.1.5 tells us how to rewrite γ in terms of the basis S_T, both as freely reduced word and as a cyclically reduced word. Let $\alpha(\Gamma, T)$ be as before. Then for the label of γ' to contain the word v, we must have that the cyclically reduced form of γ' in terms of S_T contains $b_r^2 b_1^2 \ldots b_r^2$ as a subword. Now from Corollary 2.4.6 we know that γ' is not simple in $\pi_1(\Gamma, T)$. Finally from Lemma 3.1.6 γ is not simple in $\pi_i(\Gamma, T)$, that is, $\gamma \in \pi_1(\Gamma, *)$ does not belong to a proper free factor of $\pi_1(\Gamma, *)$.

□

4.1.2 Filling forcing words and finite covers

To proceed further we will once again adapt the idea of proof of Lemma 4.1.1 to produce a "filling-forcing" path $\beta(\Gamma, T)$ of controlled length.

Convention 4.1.9. If $F(B)$ is a free group with $|B| = r \geq 2$, then the total number of freely reduced words of length 3 in $F(B)$ are $L = 2r(2r-1)^2$. Let $\{u_1, \dots, u_L\}$ be the set of all freely reduced words of length 3 in $B^{\pm 1}$. Define a freely reduced word $u_B := u_1 y_1 u_2 y_2 u_3 y_3 \dots y_{L-1} u_L$ where each y_i is either the empty word, or $y_i \in B^{\pm 1}$. Namely, whenever the concatenation $u_j u_{j+1}$ is reduced to begin with, we define y_j to be the empty word. If this concatenation is not reduced, then we can always choose $y_j \in B^{\pm 1}$ so that $u_j y_j u_{j+1}$ is reduced in $F(B)$. Note that the $|u|_B \leq 3L + L - 1 = 4L - 1$.

We now define the path $\beta(\Gamma, T)$ as follows:

Definition 4.1.10. Let $(\Gamma, *)$ be a finite connected folded core graph with a base-vertex $*$. Let $T \subseteq \Gamma$ be a maximal subtree in Γ with $E_+(\Gamma - T) = \{e_1, \dots, e_r\}$, and let $S_T = \{b_1, \dots, b_r\}$ be the basis of $\pi_1(\Gamma, *)$ dual to T. We put

$$\beta(\Gamma, T) := \delta(u_{S_T}).$$

Thus $\beta(\Gamma, T)$ is a reduced edge-path from $*$ to $*$ in Γ representing the element u_{S_T} in $\pi_1(\Gamma, *)$. Recall that u_{S_T} has length $4L - 1 = 8r(2r-1)^2 - 1$. Therefore

$$|\beta(\Gamma, T)| \leq 8rd(2r-1)^2 - 1$$

where $d = \#V\Gamma$. In particular, if Γ is a d-fold cover of R_N then $r = d(N-1) + 1$ and

$$|\beta(\Gamma, T)| \leq 8d(d(N-1)+1)(2d(N-1)+1)^2 - 1 \leq 500d^4N^3.$$

The following proposition demonstrates the a "filling-forcing" property of the path $\beta(\Gamma, T)$

Proposition 4.1.11. *Let Γ be as in Definition 4.1.10 with T a maximal tree. Let S_T and $\beta(\Gamma, T)$ be as before. Let $\gamma \in \pi_1(\Gamma, *)$ be such that γ is represented by a cyclically reduced circuit in Γ containing $\beta(\Gamma, T)$ as a subpath. Then γ is filling in $\pi_1(\Gamma, *)$.*

Proof. We first use Proposition 2.1.5 to rewrite γ as a cyclically reduced word w in $S_T = \{b_1, \dots, b_r\}$. Then the occurrence of $\beta(\Gamma, T)$ in γ produces an occurrence of the reduced word $u_1 y_1 u_2 y_2 u_3 y_3 \dots y_{l-1} u_l$ in w. Since every reduced word of length 3 now occurs in w, by Proposition 2.4.7 γ is filling in $F(b_1, \dots, b_r) = \pi_1(\Gamma, *)$. ∎

We are now in a position to prove a key proposition that is used in the proofs of Theorem A and Theorem B:

Proposition 4.1.12. *Let $N \geq 2$. Then there exists a constant $c_1 = c_1(N) > 0$ with the following property. Let $(\Gamma, *)$ be a connected d-fold cover of the N-rose R_N, where $d \geq 1$ and let $T \subseteq \Gamma$ be a maximal subtree of Γ. Then there exists a freely reduced word $w = w(\Gamma, T)$ with $|w| \leq c_1 d^5$ such that for every vertex $x \in V\Gamma$ the path $p(x, w)$ from x labeled by w in Γ contains $\beta(\Gamma, T)$ as a subpath.*

Proof. Let us begin by enumerating the vertices of $V\Gamma = \{x_1, x_2, \dots, x_d\}$. Let $H \leq F_N$ be the subgroup of index d that is represented by $(\Gamma, *)$. We have seen above that $|\beta(\Gamma, T)| \leq 500d^4N^3$.

Let e be the first edge of the path $\beta(\Gamma, T)$. Starting at the vertex $x_1 \in V\Gamma$, there exists a unique path $[x_1, *]_T$ of length $\leq d - 1$ with terminal edge e_1 (say). Lemma 4.1.6 then gives us a reduced path $p(e_1, e) = e_1 p' e$ of length $\leq 3d$. Let the word w_1 be the label of the path $p_1 = [x_1, *]_T p' \beta(\Gamma, T)$. Note that $|w_1| \leq 500d^4N^3 + 3d$.

Starting at the vertex x_2 we follow a path p'_1 that has label w_1. Let e_2 be the terminal edge of the path p'_1. Then from Lemma 4.1.6, there is a reduced path $p(e_2, e) = e_2 p''_1 e$ with $|p(e_2, e)| \leq 3d$ and $|p''_1| \leq 3d - 2$. Let the word w_2 be the label of the path $p_2 = p''_1 \beta(\Gamma, T)$. Thus $|w_2| = |p_2| \leq 500d^4N^3 + 3d$.

Now the path $p'_1 p_2 = p'_1 p''_1 \beta(\Gamma, T)$ is reduced, starts at x_2, ends in $\beta(\Gamma, T)$, has label $w_1 w_2$ and has length

$$|p'_1 p_2| = |w_1 w_2| \leq 2(500d^4N^3 + 3d).$$

We proceed inductively as follows.

For $2 \leq i \leq d$ suppose that we have already constructed freely reduced words $w_1, \dots, w_{i-1} \in F_N = F(A)$ of length $|w_j| \leq 500d^4N^3 + 3d$ such that the word $w_1 \dots w_{i-1}$ is freely reduced and such that reading $w_1 \dots w_{i-1}$ from the vertex x_{i-1} gives a reduced path in Γ ending in $\beta(\Gamma, T)$.

Starting at vertex x_i we follow the path p'_{i-1} labeled by the word $w_1 w_2 \dots w_{i-1}$. Let e_i be the terminal edge of the path p'_{i-1}. Then the path $p(e_i, e) = e_i p''_{i-1} e$ is reduced with $|p''_{i-1}| \leq 3d - 2$. Let the word w_i be the label of the reduced path $p_i = p''_{i-1} \beta(\Gamma, T)$. We again have $|w_i| \leq 500d^4N^3 + 3d$. Now the path $p'_{i-1} p_i = p'_{i-1} p''_{i-1} \beta(\Gamma, T)$ is reduced, starts with x_i and ends in $\beta(\Gamma, T)$, completing the inductive step.

Finally let $w := w_1 w_2 \dots w_d$. Then w is freely reduced, has $|w| \leq 500d^5N^3 + 3d^2 \leq 1000N^3d^5$. By construction w has the property that for $i = 1, \dots, d$ reading w from x_i gives a path in Γ that contains $\beta(\Gamma, T)$ as a subpath.

We put $w(\Gamma, T) := w$ and $c_1 = 1000N^3$. The conclusion of the proposition now holds. □

The freely reduced word $w = w(\Gamma, T)$ in $F(A)$ can be viewed as a "fulling forcing" word for the elements of the fundamental group of a d-fold cover Γ of R_N.

4.2 A lower bound for the non-filling index function $f_{nfill}(n)$

Let $F_N = F(a_1, \dots, a_N)$ be free of rank $N \ge 2$, as before. It is well-known (see, for example, [44]) that for an integer $d \ge 1$ there are $\le (d!)^N$ subgroups of index d in F_N. Indeed, every subgroup of index d in F_N can be uniquely represented by a finite connected folded $2N$-regular A-graph on vertices $1, \dots, d$, where 1 is viewed as a base-vertex. Every such graph Γ is uniquely specified by choosing an ordered N-tuples of permutations in S_d. Indeed, if $\sigma_1, \dots, \sigma_N \in S_d$, we construct Γ with $V\Gamma = \{1, \dots, d\}$ by putting an edge from j to $\sigma_i(j)$ labeled by a_i for $1 \le i \le N$, and $1 \le j \le d$.

Thus indeed F_N has $\le (d!)^N$ subgroups of index d and it has $\le d(d!)^N$ subgroups of index $\le d$.

Theorem A . Let $N \ge 2$ and let $F_N = F(A)$ where $A = a_1, \dots, a_N$. Then there exists a constant $c > 0$ and an integer $M \ge 1$ such that for all $n \ge M$ we have

$$f_{prim}(n) \ge f_{simp}(n) \ge f_{nfill}(n) \ge c\frac{\log n}{\log \log n}.$$

Proof. Let $d \ge 1$ be an integer. Denote $m(d) = m := d(d)!^N$. Enumerate all the subgroups of F_N of index $\le d$ as $H_1, \dots, H_m$ (we do allow repetitions in this list since the actual number of such distinct subgroups is $< m(d)$. Let $\Gamma_1, \dots, \Gamma_m$ be the base-pointed finite covers of the rose R_N representing the subgroups $H_1, \dots, H_m$.

For $i = 1, \dots, m$ let $w_i \in F(A)$ be the freely reduced "filling forcing" word with $|w_i| \le c_1 d^5$ corresponding to Γ_i as provided by Proposition 4.1.12. We can now construct a freely reduced and cyclically reduced word

$$z_d := w_1 u_1 w_2 u_2 \dots u_{m-1} w_m u_m$$

where each u_i is either the empty word or $u_i \in \{a_1, \dots, a_N\}^{\pm 1}$. Then

$$||z_d|| \le c_1 m d^5 = c_1 d^6 (d!)^N.$$

We claim that $d_{nfill}(z_d) > d$. Indeed, suppose not, that is suppose that $d_{nfill}(z_d) \le d$. Then there exists $1 \le i \le m$ such that $z_d \in H_i$ and that z_d is a non-filling element of $H_i = \pi_1(\Gamma_i, *)$. Let γ be the path in Γ_i from $*$ to $*$ labeled by z_d. By Proposition 4.1.12 the fact that z_d is cyclically reduced and contains w_i as subword implies that γ contains the path $\beta(\Gamma_i, T)$ as a subword. Hence, by Proposition 4.1.11, γ is a filling element in $\pi_1(\Gamma_i, *)$, yielding a contradiction. Thus indeed $d_{nfill}(z_d) > d$.

Now for $d \ge 1$ let $n_d := c_1 d^6 (d!)^N$. We also put $n_0 = 1$. Then for every integer $d \ge 0$ we have

$f_{nfill}(n_d) > d$. By *Stirling's formula*, there is $C > 0$ such that for all sufficiently large $d \geq 1$ we have

$$d \geq C \frac{\log n_d}{\log \log n_d} \tag{$\dagger$}$$

Similarly, using a standard calculus argument we see that for all sufficiently large d we have

$$\frac{\log(n_{d-1})}{\log \log(n_{d-1})} \geq \frac{1}{2} \frac{\log(n_d)}{\log \log(n_d)}. \tag{$\ddagger$}$$

Let $d_0 \geq 2$ be such that for all $d \geq d_0$ the inequalities ($\dagger$) and ($\ddagger$) hold and that the function the function $\dfrac{\log x}{\log \log x}$ is monotone increasing on the interval $[n_{d_0-1}, \infty)$.

Now let $n \geq n_{d_0+1}$ be an arbitrary integer. There exists a unique $d \geq 0$ such that $n_{d-1} < n \leq n_d$. Since $f_{nfill}(n)$ is a non-decreasing function, we get that $f_{nfill}(n) \geq f_{nfill}(n_{d-1}) > d - 1$ and $d - 1 \geq d_0$.

Then

$$f_{nfill}(n) \geq f_{nfill}(n_{d-1}) > d - 1 \geq C \frac{\log(n_{d-1})}{\log \log(n_{d-1})} \geq \frac{C}{2} \frac{\log(n_d)}{\log \log(n_d)} \geq \frac{C}{2} \frac{\log n}{\log \log n},$$

and the conclusion of the theorem follows. □

4.3 Estimating the primitivity index function $f_{prim}(n)$ from below by the residual finiteness growth function

This section follows the appendix to [29] by Khalid Bou-Rabee. In the appendix to our paper [29], Bou-Rabee relates the primitivity index function $f_{prim}(n)$ to the residual finiteness growth function introduced in [8]. He applies some results of Kozma and Thom [42] to improve the lower bounds for the primitivity index function to almost linear.

Let G be a finitely generated, residually finite group. The *divisibility function* $\mathrm{D}_G(g) = \mathrm{D}(g; G)$ is the minimum $[G : H]$ where H varies over all subgroups of finite index in G with $g \notin H$. For a fixed finite generating set $A \subset G$ the *residual finiteness growth function* is $\mathrm{RF}_{G,A}(n) := \max\{\mathrm{D}(g; G) \ : \ g \in G, |g|_A \leq n, g \neq 1\}$. Here $|g|_A$ is the word-length of g with respect to the word metric on G corresponding to A. In the case where G is a nonabelian free group F_N with word-length $|\cdot|_A$ given by a free basis A, we simply use this basis and denote the function by $\mathrm{RF}_G(n)$. We state and prove Bou-Rabee's result for completeness:

Theorem 4.3.1. *Let $G = F_N$ be a free group of finite rank $N \geq 2$. Then $\mathrm{RF}_G(n) \leq f_{prim}(4n + 4)$ for all $n \geq 1$.*

Proof. For each $n \geq 1$ let w_n be an element in F_N with $|w_n|_A \leq n$ such that $\mathrm{D}_G(w_n) = \mathrm{RF}_G(n)$. In the free group F_N commutativity is a transitive relation on the set of all nontrivial elements, and therefore there exists $a \in A$ such that $[w_n, a] \neq 1$. Also, in a free group any two non-commuting elements freely generate a free subgroup of rank two. Thus the elements w_n and a freely generate a free subgroup of rank 2 in F_N, and hence $\gamma_n := [w_n, w_n^a] \neq 1$. (In [13, 15] the property, that for every nontrivial $w \in F_N$ there exists $a \in A$ such that $[w, w^a] \neq 1$, is referred to as F_N being 1-*malabelian*). Note that $|[w_n, w_n^a]|_A \leq 4n + 4$. Since γ_n is a nontrivial commutator in F_N, a result of Schützenberger [66] then implies that γ_n is not a proper power in F_N.

Let H be a finite-index subgroup of G with γ_n primitive in H. If $w_n \in H$ and $w_n^a \in H$, then $[w_n, w_n^a] \in [H, H]$, and thus $[w_n, w_n^a]$ cannot be primitive in H. Hence, w_n or w_n^a is not in H. In either case, it follows that $[G : H] \geq D_G(w_n) = \mathrm{RF}_G(n)$. Since H was an arbitrary finite-index subgroup for which $[w_n, w_n^a]$ is primitive, it follows that $\mathrm{RF}_G(n) \leq f_{prim}(4n + 4)$, as desired.

□

Bou-Rabee then uses a result of Kozma and Thom [42] about lower bounds for $\mathrm{RF}_{F_N}(n)$ to directly imply:

Corollary 4.3.2. *Let $G = F_N$ be free of finite rank $N \geq 2$. There exists a constant $C > 0$ such that for all sufficiently large n we have*

$$f_{prim}(4n + 4) \geq \exp\left(\left(\frac{\log(n)}{C \log\log(n)}\right)^{1/4}\right).$$

If we assume Babai's Conjecture on the diameter of Cayley graphs of permutation groups, then for all sufficiently large n we have $f_{prim}(4n + 4) \geq n^{\frac{1}{C \log\log(n)}}$.

As mentioned in Chapter 1, Bou-Rabee's methods for the proof of Theorem 4.3.1 do not work for the index functions $f_{simp}(n)$ and $f_{nfill}(n)$. Thus for these functions the lower bound given by Theorem A remains the best known bound.

Chapter 5

Lower bounds for Simplicity Index and Non-filling Index

In this chapter we prove Theorem B. Observe that Theorem B is a probabilistic result which essentially says that there exists a constant $c > 0$ such that for "long enough" random words $w_n \in F(A) = F_N$ (for $N \geq 2$) of length n, the simplicity index $d_{simp}(w_n)$ is at least $c \log^{1/3} n$, and the non-filling index $d_{nfill}(w_n)$ is at least $c \log^{1/5} n$. In the first section, we establish exactly what we mean by "random words", and chalk out some results that we need.

5.1 Non-backtracking simple random walk on F_N

Recall that we set for the free group $F_N = F(A) = F(a_1, \ldots, a_N)$ (where $N \geq 2$) a distinguished free basis $A = \{a_1, \ldots, a_N\}$. Put $\Upsilon = A \cup A^{-1}$.

Definition 5.1.1. We consider the following finite-state Markov chain $\mathcal{X}$. The set of states for $\mathcal{X}$ is Υ. For $x, y \in \Upsilon$, the transition probability $P_{x,y}$ from x to y is defined as:

$$P_{x,y} := \begin{cases} \frac{1}{2N-1}, & \text{if } y \neq x^{-1} \\ 0, & \text{if } y = x^{-1} \end{cases} .$$

Let M be the transition matrix of $\mathcal{X}$. That is, M is a $2N \times 2N$ matrix with columns and rows indexed by Υ where for $x, y \in \Upsilon$ the entry $m_{x,y}$ in M is equal to 1 if $y \neq x^{-1}$ and is equal to 0 if $y = x^{-1}$.

We summarize the following elementary properties of $\mathcal{X}$, which easily follow from the definitions:

Lemma 5.1.2. *Let $N \geq 2$ and $\mathcal{X}$ be as in Definition 5.1.1. Then:*

1. *$\mathcal{X}$ is an irreducible aperiodic finite-state Markov chain.*

2. *The uniform probability distribution μ_1 on Υ is stationary for $\mathcal{X}$.*

3. *The matrix M is an irreducible aperiodic nonnegative matrix with the Perron-Frobenius eigenvalue $\lambda = 2N - 1$.*

Proof. For any $x, y \in \Upsilon$ there exists $z \in \Upsilon$ such that xzy is a freely reduced word. Hence $P_{x,z}P_{z,y} > 0$, which means that $\mathcal{X}$ is an irreducible Markov chain. The fact that for every $x \in \Upsilon$, we have $P_{x,x} > 0$ implies that $\mathcal{X}$ is aperiodic. Thus (1) is verified.

Part (2) easily follows from the definition of $\mathcal{X}$ by direct verification.

Part (1) implies that M is an irreducible aperiodic nonnegative matrix. Therefore, by the basic Perron-Frobenius theory, the spectral radius $\lambda := \max\{|\lambda_*| : \lambda_* \in \mathbb{C} \text{ is an eigenvalue of } M\}$ is a positive real number which is itself an eigenvalue of M called the Perron-Frobenius eigenvalue of M. It is also known that λ admits an eigenvector with strictly positive coordinates, and that any other eigenvalue of M admitting such an eigenvector is equal to λ. It is easy to see from the definition of M that for the vector v with all entries equal to 1 we have $Mv = (2N-1)v$. Therefore $\lambda = 2N-1$, as claimed. □

Let $\Omega = \Upsilon^{\mathbb{N}} = \{\omega = x_1, x_2, \dots | x_i \in \Upsilon\}$. We put the discrete topology on Υ and the product topology on Ω so that Ω becomes a compact Hausdorff space. For every finite word $\sigma \in \Upsilon^*$ the *cylinder* $Cyl(\sigma) \subseteq \Omega$ consists of all sequences $\omega \in \Omega$ with σ as the initial segment. For each $\sigma \in \Upsilon^*$ the set $Cyl(\sigma)$ is compact and open in Ω and the sets $\{Cyl(\sigma) | \sigma \in \Upsilon^*\}$ provide a basis for the product topology on Ω.

By using the uniform distribution μ_1 on Υ as the initial distribution for $\mathcal{X}$, the Markov chain $\mathcal{X}$ defines a Borel probability measure μ on Ω via the standard convolution formula:

For $\sigma = x_1 \dots x_n \in \Upsilon^*$,

$$\mu(Cyl(\sigma)) = \mu_1(x_1)P_{x_1,x_2} \dots P_{x_{n-1},x_n}.$$

Note that the support of μ is exactly ∂F_N, that is, the set of all semi-infinite freely reduced words $\omega = x_1, x_2, \dots$ over Υ.

Convention 5.1.3. For $\sigma \in \Upsilon^*$ we denote $\mu(\sigma) := \mu(Cyl(\sigma))$. Also, for the remainder of this section we denote $\lambda := 2N-1$.

The following is a direct corollary of the definitions:

Lemma 5.1.4. *Let $\sigma = x_1 \dots x_n \in \Upsilon^*$, where $n \geq 1$. Then*

$$\mu(\sigma) = \begin{cases} \frac{1}{2N(2N-1)^{n-1}}, & \text{if } \sigma \text{ is freely reduced,} \\ 0, & \text{if } \sigma \text{ is not freely reduced.} \end{cases}$$

Notation 5.1.5. Let $v, w \in \Upsilon^*$. We denote by $\langle v, w \rangle$ the number of times the word v occurs as a subword of w.

For $n \geq 1$ let $S(n)$ be the set of all freely reduced words of length n in Υ^* (so that $\#(S(n)) = 2N(2N-1)^{n-1} = \frac{2N}{2N-1}\lambda^n$), and let μ_n be the uniform probability distribution on $S(n)$.

The following statement is a special case, when applied to $\mathcal{X}$, of Proposition 3.13 in [20].

Proposition 5.1.6. *Let $\epsilon > 0$ and $0 < \ell < 1$. Then there exist constants $C_1 > 1$ and $C_2 > 0$ with the following property. Let $n \geq 1$ and $\sigma \in \Upsilon^*$ be a freely reduced word be such that $|\sigma| = \ell \log_\lambda n = \ell \log n / \log \lambda$. Then for $w_n \in S(n)$ we have*

$$P_{\mu_n}(|\langle \sigma, w_n \rangle - n\mu(\sigma)| < n^{\epsilon+(1-\ell)/2}) = 1 - O(C_1^{-n^{C_2}}),$$

and therefore, since $\lambda = 2N-1$ and $\mu(\sigma) = \frac{2N-1}{2N}\lambda^{-|\sigma|} = \frac{2N-1}{2N}n^{-\ell}$,

$$P_{\mu_n}(\left|\langle \sigma, w_n \rangle - \frac{2N-1}{2N}n^{1-\ell}\right| < n^{\epsilon+(1-\ell)/2}) = 1 - O(C_1^{-n^{C_2}}),$$

Corollary 5.1.7. *Let $\epsilon > 0$ and $0 < \ell < 1$. Let constants $C_1 > 1$ and $C_2 > 0$ be the constants provided by Proposition 5.1.6.*

1. *Let $n \geq 1$ and let $E_n \subseteq S(n)$ consist of those $w_n \in S(n)$ such that for every freely reduced $\sigma \in \Upsilon^*$ with $|\sigma| = \ell \log_\lambda n = \ell \log n / \log \lambda$ we have*

$$\left|\langle \sigma, w_n \rangle - \frac{2N-1}{2N}n^{1-\ell}\right| < n^{\epsilon+(1-\ell)/2},$$

   Then

$$P_{\mu_n}(E_n) \geq 1 - O\left(n^\ell C_1^{-n^{C_2}}\right).$$

2. *Suppose that $\epsilon > 0, 0 < \ell < 1$ are chosen so that $\ell < 1 - 2\epsilon$, and thus $1 - \ell > \epsilon + (1-\ell)/2$. Let $H_n \subseteq S(n)$ consist of all $w_n \in S(n)$ such that for every freely reduced σ with $|\sigma| = \ell \log_\lambda n$ we have*

$$\langle \sigma, w_n \rangle \geq \frac{2N-1}{4N}n^{1-\ell}.$$

   Let $n_0 \geq 1$ be such that for all $n \geq n_0$ we have $\frac{2N-1}{4N}n^{1-\ell} \geq n^{\epsilon+(1-\ell)/2}$. Then for $n \geq n_0$ we have

$$P_{\mu_n}(H_n) \geq 1 - O(n^\ell C_1^{-n^{C_2}}).$$

Proof. For every freely reduced σ with $|\sigma| = \ell \log_\lambda n$ let $E'_{n,\sigma}$ consist of all $w_n \in S(n)$ such that $|\langle \sigma, w_n \rangle - n\mu(\sigma)| \geq$

$n^{\epsilon+(1-\ell)/2}$. Thus, by Proposition 5.1.6, for every such σ we have $P_{\mu_n}(E'_{n,\sigma}) \leq O(C_1^{-n^{C_2}})$.

Suppose $w_n \notin E_n$. Then there exists freely reduced $\sigma \in \Upsilon^*$ with $|\sigma| = \ell \log_\lambda n$ such that $w_n \in E'_{n,\sigma}$. Since there are $O(n^\ell)$ freely reduced words σ with $|\sigma| = \ell \log_\lambda n$, it follows that $P_{\mu_n}(S(n) \setminus E_n) \leq O\left(n^\ell C_1^{-n^{C_2}}\right)$. Hence $P_{\mu_n}(E_n) \geq 1 - O\left(n^\ell C_1^{-n^{C_2}}\right)$, as required, and part (1) of Corollary 5.1.7 is verified.

Part (2) now directly follows from part (1). □

Notation 5.1.8. For a freely reduced word $w \in \Upsilon^*$ let $\iota(w)$ be the maximal initial segment of w such that $(\iota(w))^{-1}$ is a terminal segment of w. Let $\tilde{w}$ be the word obtained by removing the initial and terminal segments of w of length $|\iota(w)|$. Thus $\tilde{w}$ is the cyclically reduced form of w.

The following facts are well-known and easy to check by a direct counting argument; see [3] for details:

Lemma 5.1.9. *The following hold:*

1. *For every $0 < \epsilon_0 < 1$ there exists $C_0 > 1$ such that for $w_n \in S(n)$*

$$P_{\mu_n}(|\iota(w_n)| \leq \epsilon_0 n) \geq 1 - O(C_0^{-n}).$$

2. *There is $C > 1$ such that for $w_n \in S(n)$*

$$P_{\mu_n}(w_n \textit{ is not a proper power in } F_N) \geq 1 - O(C^{-n}).$$

5.2 Bounding below the simplicity index $d_{simp}(g)$ and the non-filling index $d_{nfill}(g)$ for random elements

Recall that for a non-trivial element $g \in F_N$ we denote by $d_{prim}(g)$ the smallest $d \geq 1$ such that there exists a subgroup $H \leq F_N$ with $[F_N : H] \leq d$ such that $g \in H$ and, moreover, that g is primitive in H. Similarly, for $g \neq 1 \in F_N$ we denote by $d_{simp}(g)$ the smallest $d \geq 1$ such that there exists a subgroup $H \leq F_N$ with $[F_N : H] \leq d$ such that $g \in H$ and, moreover, that g belongs to a proper free factor of H. Finally, for $g \in F_N - \{1\}$ we denote by $d_{nfill}(g)$ the smallest $d \geq 1$ such that there exists a subgroup $H \leq F_N$ with $[F_N : H] \leq d$ such that $g \in H$ and, moreover, that g is not filling in H. As we have seen, for every $g \in F_N - \{1\}$ we have $d_{nfill}(g) \leq d_{simp}(g) \leq d_{prim}(g) \leq ||g||_A$, where $A = \{a_1, \dots, a_n\}$ is a free basis of F_N.

Recall that for $n \geq 1$ we denote by μ_n the uniform probability distribution on the sphere $S(n) \subseteq F(A) = F_N$. We can now prove Theorem B and re-state it here for convenience:

Theorem B . Let $N \geq 2$ and let $F_N = F(A)$ where $A = \{a_1, \ldots, a_N\}$.

Then there exist constants $c > 0$, $D_1 > 1$, $1 > D_2 > 0$ such that for $n \geq 1$ and for a freely reduced word $w_n \in F(A)$ of length n chosen uniformly at random from the sphere $S(n)$ of radius n in $F(A)$ we have

$$P_{\mu_n}\left(d_{simp}(w_n) \geq c \log^{1/3} n\right) \geq_{n\to\infty} 1 - O\left((D_1)^{-n^{D_2}}\right)$$

and

$$P_{\mu_n}\left(d_{nfill}(w_n) \geq c \log^{1/5} n\right) \geq_{n\to\infty} 1 - O\left((D_1)^{-n^{D_2}}\right)$$

so that

$$\lim_{n\to\infty} P_{\mu_n}\left(d_{simp}(w_n) \geq c \log^{1/3} n\right) = 1$$

and

$$\lim_{n\to\infty} P_{\mu_n}\left(d_{nfill}(w_n) \geq c \log^{1/5} n\right) = 1$$

Proof. Choose $\epsilon > 0$ and $0 < \ell < 1$ such that $\ell < 1 - 2\epsilon$ (for concreteness we can take $\ell = 1/2$ and $\epsilon = 1/5$). Thus $1 - \ell > \epsilon + (1-\ell)/2 > 0$. Let $n_0 \geq 1$ be such that for all $n \geq n_0$ we have

$$\frac{2N-1}{4N}(0.99n)^{1-\ell} \geq (0.99n)^{\epsilon+(1-\ell)/2} \geq 1.$$

Let $C_1 > 1$ and $C_2 > 0$ be the constants provided by Corollary 5.1.7. Note that we can assume that $0 < C_2 < 1$ since decreasing C_2 preserves the validity of the conclusion of Corollary 5.1.7.

For $w_n \in S(n)$ denote by w_n' the subword of w_n obtained by removing the initial and terminal segments of length $0.005n$ from w_n. Then $|w_n'| = 0.99n$ so that $w_n' \in S(0.99n)$. Since the uniform distribution on $A^{\pm 1}$ is stationary for the Markov chain $\mathcal{X}$, it follows that under the map $S(n) \to S(0.99n)$, $w_n \mapsto w_n'$ the uniform distribution μ_n on $S(n)$ projects to the uniform distribution $\mu_{0.99n}$ on $S(0.99n)$.

Let H_n' be the event that for $w_n \in S(n)$ the word w_n' satisfies the property that for every freely reduced word $\sigma \in F(A)$ with $|\sigma| = \ell \log_\lambda(0.99n)$ we have

$$\langle \sigma, w_n' \rangle \geq 1.$$

Since for $n \geq n_0$ we have $\frac{2N-1}{4N}(0.99n)^{1-\ell} \geq (0.99n)^{\epsilon+(1-\ell)/2} \geq 1$, Corollary 5.1.7 implies that

$$P_{\mu_n}(H_n') \geq 1 - O((0.99n)^\ell C_1^{-(0.99n)^{C_2}}) = 1 - O\left(n^\ell (C_1)^{-0.99^{C_2} n^{C_2}}\right) \geq 1 - O\left((C_1')^{-n^{C_2'}}\right)$$

where $C_1' = (C_1+1)/2$ and $C_2' = C_2/2$ (for the last inequality we use the fact that $0 < C_2 < 1$). Note that $C_1' > 1$ and $1 > C_2' > 0$.

Let $Q_n \subseteq S(n)$ be the event that for $w_n \in S(n)$ we have $\iota(w_n) \le 0.001n$. Lemma 5.1.9 implies that $P_{\mu_n}(Q_n) \ge 1 - O(C_0^{-n})$ for some constant $C_0 > 1$. Now let H_n'' be the set of all $w_n \in H_n'$ such that $\iota(w_n) \le 0.001n$, that is, $H_n'' = H_n' \cap Q_n$.

Then

$$P_{\mu_n}(H_n'') \ge 1 - O\left((C_1')^{-n^{C_2'}}\right) - O(C_0^{-n}) \ge_{n\to\infty} 1 - O\left((D_1)^{-n^{D_2}}\right)$$

where $D_1 = \min\{C_0, C_1'\}$ and $D_2 = \min\{C_2', 1\} = C_2'$, so that $D_1 > 1$ and $1 > D_2 > 0$.

We choose $c > 0$ such that $c_0 c^3 \le \frac{\ell}{2\log(2N-1)}$, where $c_0 > 0$ is the constant provided by Proposition 4.1.7.

Let $n \ge n_0$ and let $w_n \in S(n)$ be such that $w_n \in H_n''$.

Since $\iota(w_n) \le 0.001n$ and since w_n' is the subword of w_n obtained by removing the initial and terminal segments of length $0.005n$ from w_n, it follows that w_n' is a subword of the cyclically reduced form $\tilde{w}_n$ of w_n.

Let $d = d_{simp}(w_n) = d_{simp}(\tilde{w}_n)$. We claim that $d \ge c\log^{1/3} n$.

Indeed, suppose not, that is, suppose that $d < c\log^{1/3} n$. Let (Γ, x_0) be a d-fold cover of the N-rose $(R_N, *)$ such that $\tilde{w}_n$ lifts to a loop γ_n from x_0 to x_0 in Γ such that γ_n belongs to a proper free factor of $\pi_1(\Gamma, x_0)$. Note that since $\tilde{w}_n$ is cyclically reduced, the closed path γ_n is also cyclically reduced.

Let T be a maximal subtree of Γ and let $v = v(\Gamma, T)$ be the freely reduced word in $F(A)$ with $|v| \le c_0 d^3$ provided by Proposition 4.1.7. Thus $|v| \le c_0 d^3 \le c_0 c^3 \log n$.

By definition of H_n'', the fact that $w_n \in H_n''$ implies that the word w_n' contains as subwords all freely reduced words in $F(A)$ of length

$$\ell \log_\lambda(0.99n) = \frac{\ell}{\log(2N-1)}(\log n - |\log 0.99|)$$

There is $n_1 \ge n_0$ such that for all $n \ge n_1$ we have

$$\frac{\ell}{\log(2N-1)}(\log n - |\log 0.99|) \ge \frac{\ell}{2\log(2N-1)} \log n.$$

Hence for $n \ge n_1$ the word w_n' contains as subwords all freely reduced words of length $\frac{\ell}{2\log(2N-1)}\log n$. Since $|v| \le c_0c^3 \log n \le \frac{\ell}{2\log(2N-1)}\log n$, it follows that w_n' contains v as a subword.

Recall that w_n' is a subword of the cyclically reduced form $\tilde{w}_n$ of w_n.

Therefore, by Proposition 4.1.7, the path γ_n in Γ, labeled by $\tilde{w}_n$, contains $\alpha(\Gamma, T)$ as a subpath. By

Corollary 4.1.8 this implies that γ_n does not belong to a proper free factor of $\pi_1(\Gamma, x_0)$, yielding a contradiction.

Thus $d = d_{simp}(w_n) \geq c \log^{1/3} n$, as claimed.

We have verified that for every $w_n \in H''_n$, where $n \geq n_1$, we have $d_{simp}(w_n) \geq c \log^{1/3} n$, and we also know that

$$P_{\mu_n}(H''_n) \geq 1 - O\left((D_1)^{-n^{D_2}}\right).$$

The conclusion of Theorem B regarding $d_{simp}(w_n)$ is established.

The proof of the conclusion of Theorem B regarding $d_{nfill}(w_n)$ is identical, with Proposition 4.1.11 and Proposition 4.1.12 used instead of Proposition 4.1.7 and Corollary 4.1.8.

□

Chapter 6

Untangling closed geodesics on hyperbolic surfaces

Having established lower bounds for our free group functions, we can now focus on our motivating problem relating to untangling closed curves on surfaces. In this chapter we will apply the free group bounds to the surface functions defined in Chapter 1, and thereby prove Theorem C and Theorem D.

6.1 Lower bounds for $\deg_{\Sigma,\rho}$ and $f_{\Sigma,\rho}$ for hyperbolic surfaces.

We need the following well-known fact:

Lemma 6.1.1. *Let S be a compact connected surface with $b \geq 2$ boundary components such that $\pi_1(S)$ is free of rank ≥ 2. Let γ be an essential simple closed curve (possible peripheral) on S and let $x \in S$ be a base-point for S. Then the loop at x corresponding to γ belongs to a proper free factor of $\pi_1(S, x)$.*

Proof. Without loss of generality we may assume that $x \in \gamma$.

By assumption, we have $\pi_1(S, x) = F_m$ with $m \geq 2$. Since S has $b \geq 2$ boundary components, it follows that every boundary component (when realized as a loop at x) represents a primitive element of F_m.

Let γ be an essential simple closed curve on S. If γ is peripheral, then γ is a primitive element of F_m and thus belongs to a proper free factor of F_m.

Suppose now that γ is non-peripheral. Then cutting S along γ yields a nontrivial splitting of $F_m = \pi_1(S)$ as an amalgamated product (if γ is separating) or as an HNN-extension (if γ is non-separating) over $\langle\gamma\rangle = \mathbb{Z}$. Suppose that γ is separating, and it cuts S into two compact surfaces S_1 and S_2 with $S_1 \cap S_2 = \gamma$ and $S_1 \cup S_2 = S$, each of $\pi_1(S_1), \pi_1(S_2)$ is free of rank ≥ 2. Thus $F_m = \pi_1(S, x) = \pi_1(S_1, x) *_\gamma \pi_1(S_2, x)$. The fact that $b \geq 2$ means that at least one of S_1, S_2 has ≥ 2 boundary components. Assume for concreteness that S_1 has ≥ 2 boundary components. Then γ is primitive in $\pi_1(S_1, x)$. Thus we can find a free basis $a_1, \ldots, a_m$ of $\pi_1(S_1, x)$ such that $m \geq 2$ and $\gamma = a_m$. Also choose a free basis $b_1, \ldots, b_k$ of $\pi_1(S_2, x)$, where $k \geq 2$. Let $v \in F(b_1, \ldots, b_k) = \pi_1(S_2, x)$ be the freely reduced word equal to γ in $\pi_1(S_2, x)$. Then the above splitting of $\pi_1(S, x)$ can be written as $\pi_1(S, x) = F(a_1, \ldots, a_m) *_{a_m = v} F(b_1, \ldots, b_k)$. By eliminating the generator a_m

from this presentation, we see that $\pi_1(S,x) = F(a_1,\dots,a_{m-1},b_1,\dots,b_k)$. Thus $\gamma = v(b_1,\dots,b_k)$ belongs to a proper free factor $F(b_1,\dots,b_k)$ of $\pi_1(S,x)$, as required. The case where γ is non-separating is similar.

Note that there is a general result (see, for example, [5, Lemma 4.1] and [62, Proposition 5.1]) which says that whenever the free group F_N (with $N \geq 2$) splits nontrivially as an amalgamated free product or an HNN-extension over a maximal infinite cyclic subgroup $\langle g\rangle$, then g belongs to a proper free factor of F_N. $\square$

The following proposition relates the degree function $\deg_{\Sigma,\rho}(\gamma)$ for curves in hyperbolic surfaces discussed in Chapter 1, with the simplicity index d_{simp} in free groups for curves contained in suitable subsurfaces:

Proposition 6.1.2. *Let (Σ,ρ) be a compact connected hyperbolic surface with (possibly empty) geodesic boundary. Let $\Sigma_1 \subseteq \Sigma$ be a compact connected subsurface with ≥ 3 boundary components, each of which is a geodesic in (Σ,ρ). Let $x \in \Sigma_1$ be a base-point. Then for every nontrivial element $g \in \pi_1(\Sigma_1,x)$ represented by a closed geodesic γ_g on Σ we have*

$$\deg_{\Sigma,\rho}(\gamma_g) \geq d_{simp}(g;\pi_1(\Sigma_1,x)).$$

Proof. By assumption $\pi_1(\Sigma_1,x) \cong F_m$ is free of rank $m \geq 2$. The fact that Σ_1 is a subsurface of Σ with geodesic boundary implies that if $g \in \pi_1(\Sigma_1,x)$ is a nontrivial element, then the shortest geodesic in Σ in the free homotopy class of g is contained in Σ_1. Indeed, the universal cover $X := \widetilde{(\Sigma_1,x)}$ is a convex $\pi_1(\Sigma_1,*)$-invariant subset of $\widetilde{(\Sigma,x)} = \mathbb{H}^2$. Therefore for every nontrivial element $g \in \pi_1(\Sigma_1,x)$ the axis $Axis(g)$ of g in $\mathbb{H}^2$ is contained in X. The image of $Axis(g)$ in Σ is the unique closed geodesic in the free homotopy class of g; the fact that $Axis(g) \subseteq X$ implies that this closed geodesic is contained in Σ_1, as claimed.

Now let $1 \neq g \in \pi_1(\Sigma',x)$ and γ_g be as in the assumptions of the proposition. Thus γ_g is contained in Σ_1.

Let $d = \deg_{\Sigma,\rho}(\gamma_g)$. Let $p : \hat{\Sigma} \to \Sigma$ be a d-fold cover of Σ such that γ_g lifts to a simple closed geodesic $\hat{\gamma}_g$ in $\hat{\Sigma}$. Let $\hat{\Sigma}_1 \subseteq \hat{\Sigma}$ be the connected component of the full preimage $p^{-1}(\Sigma_1)$ of Σ_1 containing $\hat{\gamma}_g$. Then $p : \hat{\Sigma}_1 \to \Sigma_1$ is a d'-fold cover of Σ_1 with $d' \leq d$. Pick a base-point $x' \in \hat{\Sigma}_1$ such that $p(x') = x$.

The cover $p : (\hat{\Sigma}_1,x') \to (\Sigma_1,x)$ corresponds to a subgroup $H \leq \pi_1(\Sigma_1,x)$ of index d', such that $p_{\#}(\pi_1(\hat{\Sigma}_1,x)) = H$, and that $p_{\#}$ maps $\pi_1(\hat{\Sigma}_1,x')$ isomorphically to H.

Since $\hat{\Sigma}_1$ is a cover of Σ_1, the surface $\hat{\Sigma}_1$ has ≥ 2 boundary components and $\pi_1(\hat{\Sigma}_1)$ is free of rank ≥ 2. By Lemma 6.1.1, the fact that $\hat{\gamma}_g$ is an essential simple closed curve on $\hat{\Sigma}_1$ implies that $\hat{\gamma}_g$ corresponds an element $w \in \pi_1(\hat{\Sigma}_1,x')$ which belongs to a proper free factor of $\pi_1(\hat{\Sigma}_1,x')$. Since $p(\hat{\gamma}_g) = \gamma_g$, we have $p_{\#}(w) = g \in H$. Since $p_{\#}$ maps $\pi_1(\hat{\Sigma}_1,x')$ isomorphically to H, we conclude that g belongs to a proper free factor of H. Thus $H \leq \pi_1(\Sigma_1,x)$, $[\pi_1(\Sigma_1,x):H] = d'$ and g belongs to a proper free factor of H. Therefore

$d' \geq d_{simp}(g; \pi_1(\Sigma_1, x))$. Therefore

$$\deg_{\Sigma,\rho}(\gamma_g) = d \geq d' \geq d_{simp}(g; \pi_1(\Sigma_1, x)),$$

as required.

□

We now prove Theorem D which we re-state here for completeness:

Theorem D . Let Σ be a compact connected surface with a hyperbolic structure ρ and with (possibly empty) geodesic boundary. Let $\Sigma_1 \subseteq \Sigma$ be a compact connected subsurface with ≥ 3 boundary components, each of which is a geodesic in (Σ, ρ). Let $x \in \Sigma_1$ and let A be a free basis of $\pi_1(\Sigma_1, x)$.

Let $w_n \in F(A) = \pi_1(\Sigma_1, x)$ be a freely reduced word of length n over $A^{\pm 1}$ generated by a simple non-backtracking random walk on $F(A) = \pi_1(\Sigma_1, x)$. Let γ_n be the closed geodesic on (Σ, ρ) in the free homotopy class of w_n.

Then there exist constants $c > 0, K' \geq 1$ such that

$$\lim_{n\to\infty} Pr(\deg_{\Sigma,\rho}(\gamma_n) \geq c \log^{1/3} n) = 1$$

and such that with probability tending to 1 as $n \to \infty$ we have that $w_n \in \pi_1(\Sigma, x)$ is not a proper power and that $n/K' \leq \ell_\rho(\gamma_n) \leq K'n$.

Proof. As we have seen in the proof of Proposition 6.1.2, the fact that Σ_1 is a subsurface of Σ with geodesic boundary implies that if $g \in \pi_1(\Sigma_1, *)$ is a nontrivial element, then the shortest geodesic in Σ in the free homotopy class of g is contained in Σ_1.

By Theorem B and Lemma 5.1.9, there exist an integer $n_0 \geq 1$ such that for $n \geq n_0$, with probability tending to 1 as $n \to \infty$ we have that w_n is not a proper power in $F(A)$, that $0.99n \leq ||w_n||_A \leq n = |w_n|_A$ and $d_{simp}(w_n; F(A)) \geq c \log^{1/3} n$, where $c = c(A) > 0$ is the constant provided by Theorem B for the free group $F_m = F(A)$.

Proposition 6.1.2 now implies that with probability tending to 1 as $n \to \infty$ we have

$$\deg_{\Sigma,\rho}(\gamma_n) \geq d_{simp}(w_n; F(A)) \geq c \log^{1/3} n.$$

Finally, the fact that Σ_1 has geodesic boundary in (Σ, ρ) also implies that there exists a constant $K \geq 1$ such that for every nontrivial element $g \in \pi_1(\Sigma_1, x)$ represented by a closed geodesic γ on (Σ, ρ) we have

$||g||_A/K \leq \ell_\rho(\gamma) \leq K||g||_A$. Since with probability tending to 1 as $n \to \infty$ we have $0.99n \leq ||w_n||_A \leq n = |w_n|_A$, it follows that for all sufficiently large n with with probability tending to 1 as $n \to \infty$ we have $0.99n/K \leq \ell_\rho(\gamma_n) \leq Kn$, as required.

□

Remark 6.1.3. Theorem D directly implies (e.g. by taking Σ_1 to be a suitable pair-of-pants subsurface) that if (Σ, ρ) is a compact connected hyperbolic surface of genus ≥ 2 with (possibly empty) geodesic boundary, then there exists $c' = c'(\Sigma) > 0$ such that for every $L \geq sys(\rho)$ we have $f_\rho(L) \geq c'(\log L)^{1/3}$.

6.2 Lower bounds for $\deg_{\Sigma,\rho}^{nfill}$ and $f_{\Sigma,\rho}^{nfill}$ for hyperbolic surfaces.

Our results about the behavior of d_{nfill} in free groups can also be used to obtain information about $\deg_{\Sigma,\rho}^{nfill}$ for compact hyperbolic surfaces.

Lemma 6.2.1. *Let (Σ, ρ) be a compact connected hyperbolic surface with $b \geq 1$ geodesic boundary components. Then the following hold:*

1. *If γ is a non-filling closed geodesic on (Σ, ρ) , then γ represents a non-filling element of the free group $\pi_1(\Sigma)$.*
2. *For any closed geodesic γ on (Σ, ρ) we have $\deg_{\Sigma,\rho}^{nfill}(\gamma) \geq d_{nfill}(\gamma, \pi_1(\Sigma))$.*

Proof. To see that (1) holds, let γ be a non-filling closed geodesic on (Σ, ρ). Then γ is contained in a proper compact connected subsurface Σ_1 of (Σ, ρ) with geodesic boundary. Cutting Σ open along the boundary of Σ_1 provides a nontrivial graph-of-groups decomposition of $\pi_1(\Sigma)$ with maximal cyclic edge groups and such that γ belongs to a vertex group of this decomposition. Hence γ is non-filling in $\pi_1(\Sigma)$. Thus (1) holds.

For (2), let γ be a closed geodesic on (Σ, ρ). Let $d = \deg_{\Sigma,\rho}^{nfill}(\gamma)$ and let $\hat{\Sigma} \to \Sigma$ be a degree-d cover such that γ lifts to a closed non-filling geodesic $\hat{\gamma}$ on $\hat{\Sigma}$. This cover corresponds to a subgroup $H = \pi_1(\Sigma_1) \leq \pi_1(\Sigma)$ of index d containing the element γ. The fact that $\hat{\gamma}$ is a non-filling curve in Σ_1 implies, by part (1) of this lemma, that γ is a non-filling element of $H = \pi_1(\Sigma_1)$. Therefore, by definition, $d_{nfill}(\gamma, \pi_1(\Sigma)) \leq d = \deg_{\Sigma,\rho}^{nfill}(\gamma)$, as required. □

We can now prove Theorem C which we re-state here for convenience:

Theorem C . Let (Σ, ρ) be a compact connected hyperbolic surface with $b \geq 1$ geodesic boundary components. Then there exists $C' > 0$ such that for all sufficiently large L we have

$$f_{\Sigma,\rho}(L) \geq f_{\Sigma,\rho}^{nfill}(L) \geq C' \frac{\log L}{\log \log L}.$$

Proof. Let $\pi_1(\Sigma) = F_N = F(A)$ where $A = \{a_1, \dots, a_N\}$ with $N \geq 2$. The universal cover $X = (\tilde{\Sigma}, \tilde{\rho})$ is a convex $\pi_1(\Sigma)$-invariant subset of $\mathbb{H}^2$. Therefore the orbit map $F(A) \to \mathbb{H}^2$, $w \mapsto w*$ (where $* \in \mathbb{H}^2$ is some basepoint) is a $\pi_1(\Sigma)$-equivariant quasi-isometry. Hence there exists $K \geq 1$ such that for every closed geodesic γ on (Σ, ρ) representing an element $w \in \pi_1(\Sigma)$ we have $||w||_A/K \leq \ell_\rho(\gamma) \leq K||w||_A$.

By Theorem A there exists a sequence of nontrivial cyclically reduced elements $w_n \in F(A)$ such that $||w_n||_A = n$ and that for all sufficiently large n we have

$$d_{nfill}(w_n, F(A)) \geq C\frac{\log n}{\log\log n},$$

where $C > 0$ is the constant provided by Theorem A. By Lemma 6.2.1, it follows that for all sufficiently large n we have

$$\deg^{nfill}_{\Sigma,\rho}(\gamma) \geq d_{nfill}(w_n, F(A)) \geq C\frac{\log n}{\log\log n}.$$

Since $||w||_A/K \leq \ell_\rho(\gamma) \leq K||w||_A$, the statement of the theorem now follows. □

Theorem 6.2.2. *Let (Σ, ρ) be a compact connected hyperbolic surface with $b \geq 1$ geodesic boundary components. Let $A = \{a_1, \dots, a_N\}$ be a free basis of $\pi_1(\Sigma, x)$, so that $\pi_1(\Sigma) = F(A)$. Let $w_n \in F(A) = \pi_1(\Sigma, x)$ be a freely reduced word of length n over $A^{\pm 1}$ generated by a simple non-backtracking random walk on $F(A)$. Let γ_n be the closed geodesic on (Σ, ρ) in the free homotopy class of w_n.*

Then there exist constants $c_1 > 0, K_1 \geq 1$ such that

$$\lim_{n\to\infty} Pr(\deg^{nfill}_{\Sigma,\rho}(\gamma_n) \geq c_1 \log^{1/5} n) = 1$$

and such that with probability tending to 1 as $n \to \infty$ we have that $w_n \in \pi_1(\Sigma, x)$ is not a proper power and that $n/K_1 \leq \ell_\rho(\gamma_n) \leq K_1 n$.

Proof. The proof is essentially identical to the proof of Theorem D. □

6.3 Degree and index functions based on the geometric intersection number

Let Σ be a compact connected surface admitting some hyperbolic structure (so that $\pi_1(\Sigma)$ is free of rank ≥ 2). Denote by $\mathcal{C}_\Sigma$ the set of free homotopy classes of essential closed curves on Σ that are not proper powers in $\pi_1(\Sigma)$. For $[\gamma] \in \mathcal{C}_\Sigma$ denote by $d_\Sigma([\gamma])$ the smallest degree of a finite cover of Σ such that a representative of $[\gamma]$ lifts to a simple closed curve in that cover. Note that if ρ is a hyperbolic metric on Σ, then for every $[\gamma] \in \mathcal{C}_\Sigma$ there exists a unique closed ρ-geodesic $\gamma \in [\gamma]$ and $d_\rho(\gamma) = d_\Sigma([\gamma])$. Moreover, in this case $i([\gamma],[\gamma]) = i(\gamma,\gamma)$, where $i(-,-)$ is the geometric intersection number. For an integer $m \geq 1$ we define $f_\Sigma(m)$ as the maximum of $d_\Sigma([\gamma])$ where $[\gamma]$ varies over all elements of $\mathcal{C}_\Sigma$ with $i([\gamma],[\gamma]) \leq m$. Similarly, for $[\gamma] \in \mathcal{C}_\Sigma$ denote by $d_\Sigma^{nfill}([\gamma])$ the smallest degree of a finite cover of Σ such that a representative of $[\gamma]$ lifts to a non-filling closed curve in that cover. Then define $f_\Sigma^{nfill}(m)$ as the maximum of $d_\Sigma^{nfill}([\gamma])$ where $[\gamma]$ varies over all elements of $\mathcal{C}_\Sigma$ with $i([\gamma],[\gamma]) \leq m$. Since simple curves are non-filling, we always have $d_\Sigma([\gamma]) \geq d_\Sigma^{nfill}([\gamma])$ and hence $f_\Sigma(m) \geq f_\Sigma^{nfill}(m)$.

A result of Basmajian [4, Theorem 1.1] (which also can be derived from the results of Bonahon about geodesic currents [7]) states that:

Proposition 6.3.1. *Let (Σ,ρ) be a connected compact hyperbolic surface with a (possibly empty) geodesic boundary. Then there exists a constant $K = K(\Sigma,\rho) \geq 1$ such that for every closed geodesic γ on (Σ,ρ) we have*

$$i([\gamma],[\gamma]) \leq K\ell_\rho(\gamma)^2.$$

Theorem C can be used to derive a lower bound for f_Σ:

Theorem 6.3.2. *Let Σ be a compact connected surface admitting some hyperbolic structure. Then there exist a constant $c = c(\Sigma) > 0$ and an integer $m_0 \geq 1$ such that for all $m \geq m_0$ we have*

$$f_\Sigma(m) \geq f_\Sigma^{nfill}(m) \geq c\frac{\log m}{\log\log m}.$$

Proof. Fix a hyperbolic metric ρ on Σ. By Proposition 6.3.1, there exists a constant $K = K(\rho) > 0$ such that for every $[\gamma] \in \mathcal{C}_\Sigma$ we have $i([\gamma],[\gamma]) \leq K\ell_\rho([\gamma])^2$. Let $C' = C'(\Sigma,\rho) > 0$ be the constant provided by Theorem C. Then Theorem C implies that there exist a sequence of closed geodesics γ_n on (Σ,ρ) and an integer $n_0 \geq 1$ such that for every $n \geq n_0$ we have $\ell_\rho(\gamma_n) \leq n$ and $d_\Sigma^{nfill}([\gamma_n]) \geq C'\frac{\log n}{\log\log n}$. Therefore $i(\gamma_n,\gamma_n) \leq K\ell_\rho(\gamma_n)^2 \leq Kn^2$ for all $n \geq n_0$.

Fix an integer $n_1 \geq n_0$ such that for all integers $n \geq n_1$ we have $(n+1)^2 \leq 2n^2$.

Let $m \geq Kn_1^2$ be an integer. Choose an integer $n \geq n_1$ such that $Kn^2 \leq m \leq K(n+1)^2$. Then

$$i([\gamma_n],[\gamma_n]) = i(\gamma_n,\gamma_n) \leq Kn^2 \leq m \leq K(n+1)^2 \leq 2Kn^2$$

and $n \geq \frac{\sqrt{m}}{\sqrt{2K}}$.

Therefore $i([\gamma_n],[\gamma_n]) \le m$ and

$$d_{\Sigma}^{nfill}([\gamma_n]) \ge C' \frac{\log n}{\log\log n} \ge C' \frac{\log \frac{\sqrt{m}}{\sqrt{2K}}}{\log\log \frac{\sqrt{m}}{\sqrt{2K}}} = C' \frac{\frac{1}{2}\log m - \log\sqrt{2K}}{\log\left(\frac{1}{2}\log m - \log\sqrt{2K}\right)},$$

and the statement of Theorem 6.3.2 follows. □

Remark 6.3.3. Although upper bounds for the various degree and under functions are usually easier to obtain than the lower bounds, for the moment no upper bounds for $f_{\Sigma}(m)$ are available. The reason is that on a fixed hyperbolic surface there are arbitrarily long simple closed geodesics (which thus have self-intersection number 0). Thus the linear upper bound for $f_{\Sigma,\rho}(m)$, obtained by Patel [51] does not directly imply any upper bound for $f_{\Sigma}(m)$. However, we conjecture that the following statement should be true:

Let (Σ,ρ) be a connected compact hyperbolic surface with a (possibly empty) geodesic boundary. Then there exists a constant $B \ge 1$ with the following property. Whenever γ is a closed geodesic on (Σ,ρ) such that $\ell_{\rho}([\gamma] \le \ell_{\rho}(g([\gamma]))$ for every $g \in Mod(\Sigma)$, then $\ell_{\rho}(\gamma) \le Bi([\gamma],[\gamma]) + B$.

If true, this statement, together with Patel's theorem, would imply that $f_{\Sigma}(m)$ has a linear upper bound for all sufficiently large m.

Chapter 7

Nielsen Equivalence Classes in a Class of Random Groups

In this chapter we discuss Nielsen Equivalence Classes in random groups. We show that for every $k \geq 2$ and every $n \geq 2$ there exists a torsion free, one ended, word hyperbolic group G of rank n with exactly k Nielsen equivalence classes of generating n- tuples.

7.1 Introduction

Jakob Nielsen defined the notion of Nielsen Equivalence in the 1920s [49, 50]. The following definition is as in [39].

Definition 7.1.1. If G is a group, $n \geq 1$, and $\tau = (g_1, \ldots, g_n)$ is an ordered n-tuple of elements in G, an *elementary Nielsen transformation* on τ is one of the following three types of moves:

1. For some $i \in \{1, \ldots, n\}$ replace g_i in τ by g_i^{-1}
2. For some $i \neq j$, $i, j \in \{1, \ldots, n\}$ interchange g_i and g_j in τ
3. For some $i \neq j$, $i, j \in \{1, \ldots, n\}$ replace g_i in τ by $g_i g_j^{\pm 1}$

Two n-tuples $\tau = (g_1, \ldots, g_n)$ and $\tau' = (g_1', \ldots, g_n')$ are called *Nielsen equivalent*, denoted $\tau \sim_{\text{NE}} \tau'$, if there exists a finite chain of elementary Nielsen transformations taking τ to τ'.

Since elementary Nielsen transformations are invertible it follows that Nielsen equivalence is an equivalence relation on the set G^n of ordered n-tuples of elements in G for every $n \geq 1$.

Let F_n be a free group of rank $n \geq 1$ with a distinguished free basis $(x_1, \ldots, x_n)$. Nielsen showed that an n-tuple $(y_1, \ldots, y_n)$ of elements in F_n is a free basis of F_n if and only if $(x_1, \ldots, x_n) \sim_{\text{NE}} (y_1, \ldots, y_n)$ in F_n. Note also that the definition of Nielsen equivalence directly implies that if $(x_1, \ldots, x_n) \sim_{\text{NE}} (y_1, \ldots, y_n)$ in G, then $< x_1, \ldots, x_n > = < y_1, \ldots, y_n > \leq G$; that is, if two n-tuples are Nielsen equivalent, they generate the same subgroup of G. This also shows that if G is a group with $(g_1, \ldots, g_n), (g_1', \ldots, g_n') \in G^n$, then $(g_1, \ldots, g_n) \sim_{\text{NE}} (g_1', \ldots, g_n')$ if and only if there exists an automorphism $\phi \in Aut(F_n)$ with $\phi(x_i) = w_i(x_1, \ldots, x_n)$ such that $g_i' =_G w_i(g_1, \ldots, g_n)$ for each $i \in \{1, \ldots, n\}$.

In general, it is quite hard to distinguish Nielsen equivalence classes of n-tuples that generate the same subgroup. The only exception is the case $n = 2$. Here, a classic result due to Nielsen says that if $(g_1, g_2) \sim_{\mathrm{NE}} (h_1, h_2)$ in G, then $[g_1, g_2]$ is conjugate to $[h_1, h_2]$ or $[h_1, h_2]^{-1}$ in G. The result uses the fact that primitive elements in F_2 are well understood. However very few results exist for $n \geq 3$. In fact, even in the algorithmically nice setting of torsion-free, word hyperbolic groups the problem of deciding if two tuples are Nielsen equivalent is algorithmically undecidable.

Further observe that the subgroup membership problem is a special case of this problem since two tuples $(g_1, ..., g_n, h)$ and $(g_1, ..., g_n, 1)$ are Nielsen equivalent if and only if $h \in < g_1, ..., g_n >$. In particular, this implies that Nielsen equivalence is not decidable for finitely presented torsion-free small cancellation groups as they do not have decidable subgroup membership problem as shown by Rips [55].

In [39] Kapovich and Weidmann use very sophisticated methods to show that:

Theorem 7.1.2. *[39] Let $n \geq 2$ be arbitrary integer. There exists a torsion-free word-hyperbolic one-ended group G of rank n admitting generating n-tuples $(a_1, \ldots, a_n)$, and $(b_1, \ldots, b_n)$ such that the $(2n-1)$-tuples $(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}})$ and $(b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n-1 \text{ times}})$ are not Nielsen-equivalent.*

The operation of making an n-tuple of elements of G into an $(n+1)$-tuple by appending the $(n+1)$-st entry equal to $1 \in G$ is sometimes called a *stabilization move.* In the result above $(n-1)$ stabilizations are done. Note that since $(a_1, \ldots, a_n)$, and $(b_1, \ldots, b_n)$ are generating n-tuples, if we had instead done n stabilizations, then $(a_1, \ldots, a_n, \underbrace{1, \ldots, 1}_{n \text{ times}}) \sim_{\mathrm{NE}} (b_1, \ldots, b_n, \underbrace{1, \ldots, 1}_{n \text{ times}})$ in G since we could express each a_i in terms of $\{b_1, \ldots, b_n\}$ for instance. It is also clear from the result that the n-tuples $(a_1, \ldots, a_n)$ and $(b_1, \ldots, b_n)$ are not Nielsen equivalent in G since if they were, doing the stabilizations would not have mattered. Hence, the result shows that one *needs* n stabilizations in order to make the tuples $(a_1, \ldots, a_n)$, and $(b_1, \ldots, b_n)$ Nielsen equivalent.

On the other hand the above result implies that there are at least 2 Nielsen equivalence classes of generating n-tuples in G. Using much simpler methods we can say more about the number of Nielsen equivalence classes in G. We can show that in the case above there are in fact *precisely* two Nielsen equivalence classes of generating n-tuples. Further we can generalize to show the existence of groups G with precisely k Nielsen equivalence classes of generating n-tuples where $k \geq 2$ is also an arbitrary integer. We should observe here that even though our methods are simpler than those employed by Kapovich and Weidmann in [39], our goals are different. They wanted to show that the two generating n-tuples above were as "far from each other" as possible in terms of Nielsen equivalence, whereas our goal is to consider the precise number of Nielsen equivalence classes.

We can essentially say that for every $k \geq 2, n \geq 2$ there exists a torsion free, one-ended, word hyperbolic group G of rank n with exactly k Nielsen equivalence classes of generating n-tuples. More precisely:

Theorem E . Let $k \geq 2$, $n \geq 2$ be arbitrary integers. Then there exists a generic set $\mathcal{R}$ of kn-tuples

$$\tau = (u_{11}, \ldots, u_{1n}, u_{21}, \ldots, u_{2n}, \ldots, u_{k1}, \ldots, u_{kn})$$

where for each $\tau \in \mathcal{R}$, $i \in \{1, \ldots, k\}$, $j \in \{1, \ldots, n\}$, u_{ij} is a cyclically reduced word in $F(a_{i1}, \ldots, a_{in})$. Further, $|u_{11}| = \ldots = |u_{1n}| = \ldots = |u_{k1}| = \ldots = |u_{kn}|$ and such that the following holds for each $\tau = (u_{11}, \ldots, u_{1n}, u_{21}, \ldots, u_{2n}, \ldots, u_{k1}, \ldots, u_{kn})$:
Let G be a group given by the presentation:

$$G =< a_{11}, \ldots, a_{1n}, \ldots, a_{k1}, \ldots, a_{kn} | a_{1j} = u_{2j}(\underline{a_2}), a_{2j} = u_{3j}(\underline{a_3}), \ldots, a_{(k-1)j} = u_{kj}(\underline{a_k}),$$
$$a_{kj} = u_{1j}(\underline{a_1}), \text{ for } 1 \leq j \leq n >, \qquad (*)$$

where for $i \in \{1, \ldots, k\}$, $(\underline{a_i}) = (a_{i1}, \ldots, a_{in})$. Then G is a torsion-free word-hyperbolic one-ended group of rank n admitting precisely k Nielsen equivalence classes of generating n-tuples.

The precise meaning of $\mathcal{R}$ being generic is explained in Section 7.2.2. There are essentially two parts to the proof of Theorem E: first showing that our group G is torsion-free word-hyperbolic one-ended of rank n; and second showing that G admits precisely k Nielsen equivalence classes of generating n-tuples. The proofs for the first part are more or less identical to the proofs presented in [39]. We will now give a sketch for the proof of the second part. Observe that each of the tuples $(a_{11}, \ldots, a_{1n}), \ldots, (a_{k1}, \ldots, a_{kn})$ are generating n-tuples for G since we can see for example

$$G =< a_{11}, \ldots, a_{1n} | U_1, \ldots, U_n > \qquad (**)$$

where each U_j is obtained by freely and cyclically reducing the word formed by modifying the relation $a_{1j}^{-1} u_{2j}(\underline{a_2})$ as follows: for $j \in \{1, \ldots, n\}$, each $u_{2j} \in F(a_{21}, \ldots, a_{2n})$ but each $a_{2j} = u_{3j}(\underline{a_3})$. Hence we replace each occurrence of a_{2j} by $u_{3j} \in F(a_{31}, \ldots, a_{3n})$. Proceeding in a this manner by successively replacing a_{ij} by $u_{(i+1)j}$ for $i \in \{1, \ldots, k-1\}$, and finally replacing each a_{kj} by $u_{1j} \in F(a_{11}, \ldots, a_{1n})$ we get that each $a_{1j}^{-1} u_{2j}(\underline{a_2})$ can be replaced by a word in $F(a_{11}, \ldots, a_{1n})$.

It is easily shown using Tietze transformations that the presentations (*) and (**) represent the same group G. The genericity of (*) shows that the presentation (**) satisfies arbitrarily strong small cancellation condition $C'(\lambda)$ where $0 < \lambda < 1$ is an arbitrarily small fixed number. We show that the k Nielsen equivalence classes of generating n-tuples are precisely the classes of the n-tuples $(a_{11}, \ldots, a_{1n}), \ldots, (a_{k1}, \ldots, a_{kn})$. We first show that no two of these are Nielsen equivalent to each other, and then that any other generating n-tuple must be Nielsen equivalent to one of these. We start by picking two n-tuples $(a_{l1}, \ldots, a_{ln}), (a_{m1}, \ldots, a_{mn})$ with $m \neq l$, and $m, l \in \{1, \ldots, k\}$. Suppose $(a_{l1}, \ldots, a_{ln}) \sim_{\mathrm{NE}} (a_{m1}, \ldots, a_{mn})$ in G^n. Then $(a_{l1}, \ldots, a_{ln}) \sim_{\mathrm{NE}} (w_1, \ldots, w_n)$ in $F(a_{l1}, \ldots, a_{ln})$ with $w_j = a_{mj}$ in G for all $1 \leq j \leq n$. We show that, in fact, even $w_1 \neq a_{m1}$ in G. We use properties of primitive words and the results of Schupp on conjugacy diagrams [45] to note that w_1 must contain a large part of some relator. Since $w_1 \in F(a_{l1}, \ldots, a_{ln})$, the relator would have to be some $a_{(l-1)j}^{-1} u_{lj}$ (for $l \neq 1$). Using genericity we will show that this must imply that w_1 contains all freely reduced words of length 2 in $F(a_{l1}, \ldots, a_{ln})$. We will construct a sequence of graphs and use the notion of Stallings foldings to show that this leads to a contradiction.

Now suppose $(g_1, \ldots, g_n) \in G^n$ is a generating n-tuple. We will show that $(g_1, \ldots, g_n)$ must be Nielsen equivalent of one of the n-tuples $(a_{11}, \ldots, a_{1n}), \ldots, (a_{k1}, \ldots, a_{kn})$. We construct a sequence of graphs using essentially folds. We then use the Arzhantseva-Ol'shanskii "surgery trick" in the form described in [37] to show that the sequence eventually terminates in an n-rose. We then once again use a condition imposed on the generic set $\mathcal{R}$ to conclude that this n-rose must be one consisting of some $(a_{p1}, \ldots, a_{pn})$, $p \in \{1, \ldots, k\}$ with $a_{p1}, \ldots, a_{pn}$ respectively labeling the n-petals.

7.2 Preliminaries

7.2.1 Small Cancellation Theory

Within group theory, small cancellation theory studies group presentations where relators have "small overlaps". It was the thought child of Martin Greendlinger [25, 26, 27] who developed it in its most commonly used form. See Chapter V of [45] for details.

Recall that a set R of cyclically reduced words in $F = F(a_1, \ldots, a_n)$ is called *symmetrized* if for every $r \in R$ all cyclic permutations of $r^{\pm 1}$ belong to R. For a symmetrized set $R \subseteq F(a_1, \ldots, a_n)$, a freely reduced word v is called a *piece* with respect to R if there exist $r_1, r_2 \in R$ such that $r_1 \neq r_2$ and v is an initial segment of both r_1 and r_2.

Definition 7.2.1. (Small Cancellation Condition) Let $R \subseteq F(a_1, \ldots, a_n)$ be a symmetrized set of cyclically reduced words. Let $0 < \lambda < 1$. We say that R is a $C'(\lambda)$- *set* if, whenever v is a piece with respect to R and $r \in R$ is such that r contains v as a subword, then $|v| < \lambda |r|$. We say that the presentation $G = < a_1, \ldots, a_n | R >$ satisfies the $C'(\lambda)$- *small cancellation condition* if $R \subseteq F(a_1, \ldots, a_n)$ is a $C'(\lambda)$ set.

The following proposition is a well-know property of small cancellation groups [45]:

Proposition 7.2.2. *(Greendlinger's Lemma) Let $G = < a_1, \ldots, a_n | R >$ be a $C'(\lambda)$-presentation with $\lambda \leq 1/6$. Let $w \in F(a_1, \ldots, a_n)$ be a freely reduced word such that $w =_G 1$. Then w contains a subword v such that v is also a subword of some $r \in R$ satisfying $|v| > (1 - 3\lambda)|r| \geq |r|/2$.*

Definition 7.2.3. Let $G = < a_1, \ldots, a_n | R >$ be a $C'(\lambda)$-presentation with $\lambda \leq 1/100$. For a freely reduced word $w \in F(a_1, \ldots, a_n)$ we say that w is *Dehn reduced with respect to R* if w does not contain a subword v such that v is also a subword of some $r \in R$ with $|v| > |r|/2$.

We say that a freely reduced word w is *λ-reduced with respect to R* if w does not contain a subword v such that v is also a subword of some $r \in R$ with $|v| > (1 - 3\lambda)|r|$.

Similarly, we say that a cyclically reduced word $w \in F(a_1, \ldots, a_n)$ is λ- *cyclically reduced with respect to R* if every cyclic permutation of w is λ-reduced with respect to R. We also say that a cyclically reduced word $w \in F(a_1, \ldots, a_n)$ is *cyclically Dehn-reduced with respect to R* if every cyclic permutation of w is Dehn-reduced with respect to R.

Observe that for $\lambda \leq 1/6$, Dehn-reduced words are also λ-reduced.

The following is a consequence of studying Conjugacy diagrams in $C'(\lambda)$ groups that arise from results in Chapter V of [45]. It is Corollary 2.6 in [40]:

Lemma 7.2.4. *Let $G = < a_1, \ldots, a_n | R >$ be a $C'(\lambda)$ presentation with $\lambda \leq 1/100$. Suppose that for every $r \in R$ we have $|r| \geq 2/\lambda + 1$. Then if w is a cyclically reduced word in $F(a_1, \ldots, a_n)$ that is conjugate to a_1 in G, we must have that either $w = a_1$ in $F(a_1, \ldots, a_n)$ or that w is not λ-cyclically reduced with respect to R.*

7.2.2 Genericity

Roughly speaking a property $\mathcal{P}$ for groups is *generic* if a "random" group satisfies $\mathcal{P}$. There are plenty of different models. See [37, 39, 40] for details.

For $n \geq 2$, $A_i = \{a_{i1}, \ldots, a_{in}\}$, we say $F_i = F(A_i)$ is the set of freely reduced words in $A_i^{\pm}$. Let $s \geq 1$ and let $U \subseteq F_i^s$ be some set of s-tuples with entries from $F(A_i)$. Thus for $l \geq 0$ we denote by $\gamma_{A_i}(l, U)$ the number of all s-tuples $(u_1, \ldots, u_s) \in U$ such that $|u_1| = \ldots = |u_s| = l$. Thus for $l \geq 1$, we have that $\gamma_{A_i}(l, F_i^s) = (2n(2n-1)^{l-1})^s$.

Definition 7.2.5. A subset $S \subseteq F(A_i)$ is generic in $F(A_i)$ if

$$\lim_{N \to \infty} \frac{\#\{v \in S : |v| = N\}}{\#\{v \in F(A_i) : |v| = N\}} = \lim_{N \to \infty} \frac{\#\{v \in S : |v| = N\}}{2n(2n-1)^{N-1}} = 1$$

If in addition, the convergence to 1 in the above limit is exponentially fast, we say that S is *exponentially generic* in $F(A_i)$.

Definition 7.2.6. Let $U \subseteq F_i^s$ be some set of s-tuples with entries from A_i such that for every $(u_1, \ldots, u_s) \in U$ we have that $|u_1| = \ldots = |u_s|$. Let $U' \subseteq U$. We say that U' is generic in U if

$$\lim_{l \to \infty} \frac{\gamma_{A_i}(l, U')}{\gamma_{A_i}(l, U)} = 1.$$

If in addition, the convergence to 1 is exponentially fast, we say that U' is *exponentially generic in U*.

Let $C_i = C_{A_i}$ be the set of cyclically reduced words in $F(A_i)$. Let $C(i) = \{(u_1, \ldots, u_s) \in C_i^s \| |u_1| = \ldots = |u_s|\}$. Then the definitions immediately show that:

Lemma 7.2.7. *For any i: The intersection of a finite number of (exponentially) generic subsets of $C(i)$ is (exponentially) generic in $C(i)$.*

The following properties of s-tuples of cyclically reduced elements in a free group are well known to be exponentially generic:

Lemma 7.2.8. *For $n \geq 2$, for any i:*

(1) The property that no element of an s-tuple is a proper power in $F(A_i)$ is exponentially generic in $C(i)$.

(2) Let $0 < \lambda < 1$ be arbitrary. Then the property that an s-tuple $(u_1, \ldots, u_s)$ after cyclic reduction and symmetrization satisfies the $C'(\lambda)$ small cancellation condition is exponentially generic in $C(i)$.

(3) The property that for an s-tuple $(u_1, \ldots, u_s)$ for every $i \neq j$ the element u_i is not conjugate to $u_j^{\pm 1}$ in $F(A_i)$, is exponentially generic in $C(i)$.

(4) Let $K \geq 1$ be an integer and let $0 < \lambda < 1$. Then the property of as s-tuple $(u_1, \ldots, u_s)$ that every subword v of some u_j of length $\geq \lambda |u_j|$ contains as a subword every freely reduced word of length $\leq K$ in $F(A_i)$, is exponentially generic in $C(i)$.

7.2.3 The Genericity Condition

This subsection largely follows [37]. Recall that $1 \leq i \leq k$, and that for $n \geq 2$, we have fixed the sets $A_i = (a_{i1}, \ldots, a_{in})$. The following definition is from [3].

Definition 7.2.9. [3] (μ_i-readable) Let $0 < \mu < 1$ be a real number. A reduced word w in $F(A_i) = F(a_{i1}, \ldots, a_{in})$ of length $p > 0$ is called *μ_i-readable* if there exists a connected folded A_i-graph Γ where every edge is labeled by a letter of A_i such that:

(1) The number of edges in Γ is $\leq \mu p$

(2) The free group $\pi_1(\Gamma)$ has rank $\leq n-1$.

(3) There is a reduced path in Γ with label w.

The following is a specific version of a definition from [1]

Definition 7.2.10. [1] ($(\mu, n)_i$-readable) Let $0 < \mu < 1$ be a real number. A reduced word w in $F(A_i) = F(a_{i1}, \ldots, a_{in})$ of length $p > 0$ is called $(\mu, n)_i$*-readable* if there exists a connected folded A_i-graph Γ such that:

(1) The number of edges in Γ is $\leq \mu p$

(2) The free group $\pi_1(\Gamma)$ has rank $\leq n$.

(3) There is a reduced path in Γ with label w.

(4) The graph Γ has at least one vertex of degree $< 2n$.

It is not too difficult to see that words with no long "readable" subwords are generic i.e. non-$(\mu, n)_i$-readability is generic. Consider what happens when $n = 3$, and take $F = F(a, b, c)$. Look at the picture below:

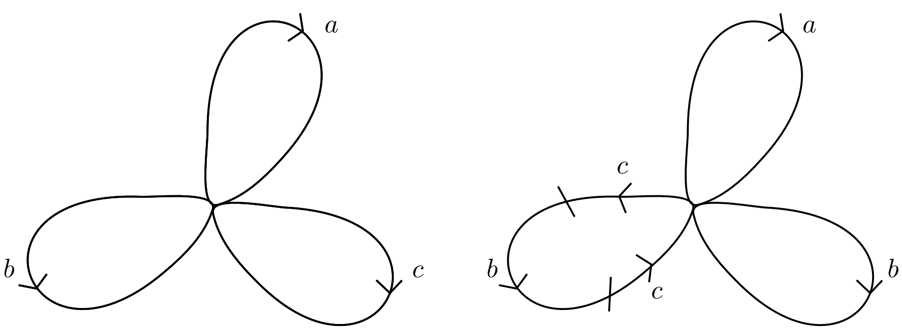

Figure 7.1: Why non-readability is generic

On the left part of Figure 7.1, the graph has rank 3, and the only vertex has degree 6. In this graph, all reduced words of length p can be obtained as labels of reduced paths. On the other hand for the graph on the right side of Figure 7.1, while the rank is still 3, there exists a vertex of degree $2 < 6$. One can see that in this case, the ratio of the number of reduced words that are labels of reduced paths of length p over the total number of reduced words of length p goes to 0 as $p \to \infty$.

The following is an appropriate "small-cancellation condition" that exploits the above:

Definition 7.2.11. [37] ($(\lambda, \mu, n)_i$-condition) Let $0 < \mu < 1$ be a real number, let $n \geq 2, s \geq 2$ be integers, and let $\lambda > 0$ be a real number such that

$$\lambda \leq \frac{\mu}{15n + 3\mu} < 1/6.$$

We'll say that a tuple of nontrivial cyclically reduced words $(r_1, \ldots, r_s)$ in $F(A_i) = F(a_{i1}, \ldots, a_{in})$ satisfies the $(\lambda, \mu, n)_i$-*condition* if:

(1) The tuple $(r_1, \ldots, r_s)$ satisfies the $C'(\lambda)$ small cancellation condition.

(2) The words r_j are not proper powers in $F(A_i)$.

(3) If w is a subword of a cyclic permutation of some r_j and $|w| \geq |r_j|/2$ then w is not $(\mu, n)_i$-readable and is not μ_i-readable.

The following follows from results of Arzhantseva and Ol'shanskii [1, 3]:

Theorem 7.2.12. *[37] For any $n \geq 2, s \geq 2$ and for any λ, μ as in Definition 7.2.11, let U_i consist of the set of all s-tuples of nontrivial cyclically reduced words $(r_1, \ldots, r_s)$ in $F(A_i) = F(a_{i1}, \ldots, a_{in})$ that satisfy the $(\lambda, \mu, n)_i$-condition. Then U_i is exponentially generic in $C(i)$.*

We now construct a generic set $\mathcal{R}$ that we need in the construction of the group G in our main result Theorem E.

Proposition 7.2.13. *Let $k \geq 2$, $n \geq 2$ be arbitrary integers. Then there exists an exponentially generic set $\mathcal{R}$ of kn-tuples*

$$\tau = (u_{11}, \ldots, u_{1n}, u_{21}, \ldots, u_{2n}, \ldots, u_{k1}, \ldots, u_{kn})$$

where for each $\tau \in \mathcal{R}$, $i \in \{1, \ldots, k\}$, $j \in \{1, \ldots, n\}$, u_{ij} is a cyclically reduced word in $F(a_{i1}, \ldots, a_{in})$. Further, $|u_{11}| = \ldots = |u_{1n}| = \ldots = |u_{k1}| = \ldots = |u_{kn}|$ and such that:

(1) For every $(u_{11}, \ldots, u_{1n}) \in C(1), (u_{21}, \ldots, u_{2n}) \in C(2), \ldots$ and $(u_{k1}, \ldots, u_{kn}) \in C(k)$, the presentation (*) that defines the group G in Theorem E satisfies $C'(1/100)$-small cancellation condition.

(2) For every $(u_{11}, \ldots, u_{1n}) \in C(1), (u_{21}, \ldots, u_{2n}) \in C(2), \ldots$ and $(u_{k1}, \ldots, u_{kn}) \in C(k)$, $|u_{ij}|_{A_i} \geq 10^{10}$

(3) For every $(u_{11}, \ldots, u_{1n}) \in C(1), (u_{21}, \ldots, u_{2n}) \in C(2), \ldots$ and $(u_{k1}, \ldots, u_{kn}) \in C(k)$, every subword v of some u_{ij} of length $\geq |u_{ij}|/100$ contains as a subword every freely reduced word of length 2 in $F(a_{i1}, \ldots, a_{in})$.

(4) For every $(u_{11},\ldots,u_{1n}) \in C(1), (u_{21},\ldots,u_{2n}) \in C(2), \ldots$ and $(u_{k1},\ldots,u_{kn}) \in C(k)$, we have that the n-tuple $(u_{i1},\ldots,u_{in})$ satisfies the $(\lambda,\mu,n)_i$-condition for the same μ,λ.

Proof. The proof follows directly from Lemma 7.2.7, Theorem 7.2.12, and Lemma 7.2.8. □

7.3 Proof of Theorem E

In our quest to prove Theorem E we want to start by showing that in the terminology of Theorem E no two n-tuples amongst $(a_{11},\ldots,a_{1n}),\ldots,(a_{k1},\ldots,a_{kn})$ are Nielsen equivalent to each other. The following is a slightly more technical version of the proof in Theorem 4.4 in [40].

Theorem 7.3.1. *Let $k \geq 2$, $n \geq 2$ be arbitrary integers. Then there exists a generic set $\mathcal{R}$ of kn-tuples*

$$\tau = (u_{11},\ldots,u_{1n},u_{21},\ldots,u_{2n},\ldots,u_{k1},\ldots,u_{kn})$$

where for each $\tau \in \mathcal{R}$, $i \in \{1,\ldots,k\}$, $j \in \{1,\ldots,n\}$, u_{ij} is a cyclically reduced word in $F(a_{i1},\ldots,a_{in})$. Further, $|u_{11}| = \ldots = |u_{1n}| = \ldots = |u_{k1}| = \ldots = |u_{kn}|$ and such that the following holds for each $\tau = (u_{11},\ldots,u_{1n},u_{21},\ldots,u_{2n},\ldots,u_{k1},\ldots,u_{kn})$:
Let G be a group given by the presentation:

$$G = < a_{11},\ldots,a_{1n},\ldots,a_{k1},\ldots,a_{kn} | a_{1j} = u_{2j}(\underline{a_2}), a_{2j} = u_{3j}(\underline{a_3}),\ldots,a_{(k-1)j} = u_{kj}(\underline{a_k}),$$
$$a_{kj} = u_{1j}(\underline{a_1}),\ \text{for } 1 \leq j \leq n >, \quad (*)$$

where for $i \in \{1,\ldots,k\}$, $(\underline{a_i}) = (a_{i1},\ldots,a_{in})$. For some $m,l \in \{1,\ldots,k\}$ with $m \neq l$, consider the tuples $(a_{l1},\ldots,a_{ln}),(a_{m1},\ldots,a_{mn})$. Then $(a_{l1},\ldots,a_{ln}) \nsim_{NE} (a_{m1},\ldots,a_{mn})$ in G^n.

Proof. We construct an exponentially generic set $\mathcal{R}$ as in Proposition 7.2.13. Let $\tau \in \mathcal{R}$ be an kn-tuple $\tau = (u_{11},\ldots,u_{1n},u_{21},\ldots,u_{2n},\ldots,u_{k1},\ldots,u_{kn})$, and G be the group described above using τ. Now we want to show that no two n-tuples amongst $(a_{11},\ldots,a_{1n}),\ldots,(a_{k1},\ldots,a_{kn})$ are Nielsen equivalent to each other. For some $m,l \in \{1,\ldots,k\}$ with $m \neq l$, consider the tuples $(a_{l1},\ldots,a_{ln}),(a_{m1},\ldots,a_{mn})$.

We proceed by contradiction. Suppose $(a_{l1},\ldots,a_{ln}) \sim_{\mathrm{NE}} (a_{m1},\ldots,a_{mn})$ in G^n. Then $(a_{l1},\ldots,a_{ln}) \sim_{\mathrm{NE}} (w_1,\ldots,w_n)$ in $F(a_{l1},\ldots,a_{ln})$ with $w_j = a_{mj}$ in G for all $1 \leq j \leq n$. In particular $w_1 = a_{m1}$ in G. Note that it may well be that $w_1 \neq u_{l1}$ in $F(a_{l1},\ldots,a_{ln})$. After a possible conjugation of $(w_1,\ldots,w_n)$ in $F(a_{l1},\ldots,a_{ln}) = F(A_l)$, we may assume that w_1 is cyclically reduced in $F(A_l)$ and conjugate to a_{m1} in G.

By Lemma 7.2.4, we observe that $w_1(\underline{a_l})$ must contain at-least a fourth of a cyclic permutation of some relator $(u_{lj}(\underline{a_l})a_{(l-1)j}^{-1})^{\pm 1}$ (note that this is true if $l \neq 1$; the case $l = 1$ is similar). Now using condition (3) in Proposition 7.2.13, we conclude that w_1 contains every freely reduced word of length 2 in $F(A_l)$ as a subword.

Now as in Figure 2.2 and Section 2.2, we construct the wedge of circles graph S_σ for $\sigma = (w_1, \ldots, w_n)$. Observe that since $(a_{l1}, \ldots, a_{ln}) \sim_{\mathrm{NE}} (w_1, \ldots, w_n)$, we have that $F(A_l) =< w_1, \ldots, w_n >$. Hence the map $\mu_\# : \pi_1(S_\sigma, x) \to \pi_1(R_n, x_0)$ is surjective. Thus by Lemma 2.2.8, we get a sequence of A_l-graphs

$$S_\sigma = \Gamma_0, \Gamma_1, \ldots, \Gamma_p = R_n$$

such that Γ_i can be obtained from Γ_{i-1} by a single fold for $1 \leq i \leq p$. Note that each Γ_i is a core graph. This holds as Γ_i is the image of loops labeled with freely reduced words and the base point must lie in the core as the cyclically reduced word w_i can be read in each Γ_i by a closed path based at the base vertex. Thus Γ_{p-1} is a core graph that folds onto R_n with a single fold and w_1 can be read by a closed path in Γ_{p-1}. The following claim will give us our contradiction:

Claim 1. We claim that since Γ_{p-1} is a core graph that folds onto R_n with a single fold, there exists a reduced word of length 2 that cannot be read as the label of an edge-path in Γ_{p-1}.

Observe that by Remark 2.2.9 the graphs $\Gamma_0, \Gamma_1, \ldots, \Gamma_p$ are all of rank n. Thus the fold to get from Γ_{p-1} to R_n must identify a loop edge with a non-loop edge. This is because identifying two loop edges decreases the rank and identifying two non-loop edges yields a graph with a non-loop edge. Hence w.l.o.g. we can assume that Γ_{p-1} has two vertices x and y and the fold identifies a loop edge at x and an edge from x to y both with label a_{l1}. Γ_{p-1} has $n-1$ more edges that are labeled by $a_{l2}, \ldots, a_{ln}$. led by $a_{l2}, \ldots, a_{ln}$. Since Γ_{p-1} is a core graph, we must have that there exists either a loop edge at y or a second edge between x and y. Thus the graph Γ_{p-1} which is one fold away from R_n will look something like the following with possibly more loops at x or y or more edges between x and y:

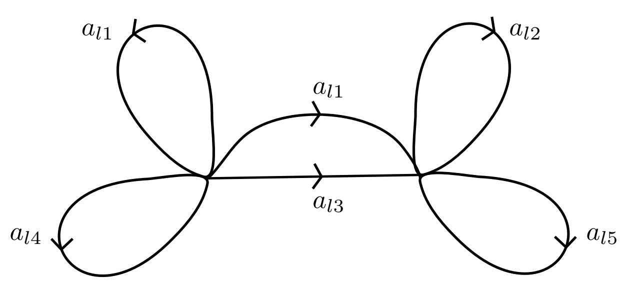

Figure 7.2: An example of the graph Γ_{p-1}

From Figure 7.2 we see that there will always be a two letter word that we can not read on this graph Γ_{p-1}. For instance, on the graph above, we can not read $a_{l3}a_{l1}$. Thus the claim holds.

However, we observed that w_1 can be read by a closed path in Γ_{p-1}, and since we saw above that w_1 contains every freely reduced word of length 2 in $F(A_l)$ as a subword, the above claim gives us our contradiction. Hence, $(a_{l1}, \dots, a_{ln}) \not\sim_{\text{NE}} (a_{m1}, \dots, a_{mn})$ in G^n. □

Observe that thus far we have only used conditions (1)-(3) from our generic set $\mathcal{R}$. Recall that our goal to show that there are *precisely* k Nielsen equivalence classes of generating n-tuples in our group G as defined in the above theorem. The next step will be to show that any generating n-tuple must be Nielsen equivalent to one of the n-tuples between $(a_{11}, \dots, a_{1n}), \dots, (a_{k1}, \dots, a_{kn})$. To do this, we will construct a sequence of graphs that ends with a Rose. However, for this we need the following Lemma. This Lemma is due to Arzhantseva[1]. We state it as a slightly modified version of how it appears in [37]:

Lemma 7.3.2. *Let $\mathcal{R}, \tau, G$ be as defined in the above theorem. Let $A = \{a_{11}, \dots, a_{1n}, \dots, a_{k1}, \dots, a_{kn}\}$. Let Γ be a finite connected folded A-graph with a base vertex x_0 with $rank(\pi_1(\Gamma, x_0)) \leq n < kn$. Then either $\phi : \pi_1(\Gamma, x_0) \to G$ is injective and hence the subgroup H of G represented by Γ is free or there exists an AO-move on Γ that reduces the number of edges of Γ.*

Proof. A version of the argument appears in the paper of Arzhantseva and Olshanskii[3] and also in a paper of Kapovich and Schupp [37] but we give it here for completeness. Suppose ϕ is not injective then there is a nontrivial reduced loop at x_0 in Γ whose label is equal to 1 in G. Since Γ is folded, by Lemma 2.2.2 the label of this loop is a reduced word. By Greendlingers Lemma since G satisfies the small cancellation condition $C'(\lambda)$ for $\lambda \leq 1/6$, we have that the label of this loop must contain a subword v which is a subword of a cyclic permutation r of a defining relator r' from the given presentation of G. Further $|v| > (1-3\lambda)|r|$. Let p be the path in Γ whose label is v. Since $rank(\pi_1(\Gamma, x_0)) \leq n$, a counting argument implies that Γ contains at most $3n-1$ maximal arcs. We can then write $p = p_1 \dots p_s$ where $p_2, \dots, p_{k-1}$ are maximal arcs while p_1, p_k lie on some maximal arcs. Let $v_1, \dots, v_s$ be the labels of $p_1, \dots, p_s$ so that $v = v_1 \dots v_s$. We now consider the different possibilities for the length of these arcs.

Case 1. Suppose that for some q, $|p_q| \geq 5\lambda|r|$.

Now suppose that $1 < q < k$ so that p_q is actually a maximal arc. By the $C'(\lambda)$ condition and since there are no proper powers amongst our relators, we see that the arc p_q can not overlap with the other arcs p_t $(1 < t < s, t \neq q)$. For the same reasons the overlap of p_q with either p_1 or p_s has length less than $\lambda|r|$. Thus there is a subpath of p_q of length $\geq 3\lambda|r|$ that does not overlap the rest of the path p. Suppose now that $q = 1$ (the case $q = s$ is symmetric) and so $|p_1| \geq 5\lambda|r|$. This immediately says that p_1 can not overlap any of p_t for $1 < t < s$ as otherwise we would be in the case above. Finally, with the same reasoning as before the overlap between p_1 and p_s is $< \lambda|r|$. Thus, if for any $1 \leq q \leq s$ we have that $|p_q| \geq 5\lambda|r|$, we must also have that there is a subpath p'_q of p_q of length $\geq 3\lambda|r|$ that does not overlap the rest of the path p.

Now we do the following AO-move on Γ. First we add an arc from $t(p_s) = t(p)$ to $o(p_1) = o(p)$ labeled by the part of the relator r (of length $< 3\lambda|r|$) which is missing in v. Next remove the arc p'_q (of length $\geq 3\lambda|r|$) that does not intersect with anything else in p. Clearly this AO-move has reduced the number of edges in Γ and so in this case, we are done.

Case 2. Suppose that $|p_q| < 5\lambda|r|$ for all $1 \leq q \leq s$.

Then we can read the word v as the label of a path in some connected subgraph Γ' of Γ where $rank(\pi_1(\Gamma')) \leq n$. Take Γ' to consist of the union of all the edges in p. Let us estimate the number of edges in Γ'. Recall that Γ has $\leq 3n-1$ distinct maximal arcs and each p_q either is or is contained in a maximal arc of Γ. Since $|p_q| < 5\lambda|r|$ for all $1 \leq q \leq s$, we then have that the number of edges in Γ' is $\leq (3n-1)5\lambda|r| \leq 3n5\lambda|r|$. But the conditions on μ show that then the number of edges in Γ' is $\leq \mu(1-3\lambda)|r| \leq \mu|v|$. This contradicts condition (4) of $\mathcal{R}$. □

Thus in particular, in the case above, one can also do an AO-move on the graph and reduce the number of edges in the graph. Note also that for the proof above, we only used condition (4) of our generic set $\mathcal{R}$. We need one more lemma for our purposes. This is Lemma 2 in [2], and Lemma 3.1 in [37]. The lemma holds in a more general setting than it is stated here.

Lemma 7.3.3. *[2, 37] Let Γ be a finite connected A-graph with base vertex x_0 and fundamental group being free of rank n. Let Γ' be obtained from Γ by a fold preserving the Euler characteristic, or by an AO-move. Let x'_0 be the image of x_0 under the fold or AO-move. Let $\phi : \pi_1(\Gamma, x_0) \to G$ and $\phi' : \pi_1(\Gamma', x'_0) \to G$ be the label maps for Γ and Γ' respectively. For any l-tuple ρ freely generating $\pi_1(\Gamma, x_0)$, let ρ' be the l-tuple that is the image of ρ under the map induced at the level of fundamental groups by the fold or by the AO-move. Then $\phi(\rho) \sim_{NE} \phi'(\rho')$ in G.*

The proof is quite straightforward and is given in [37]. However, exactly what this induced map at the level of fundamental groups is merits some explanation. Suppose first that Γ' is obtained from Γ by a fold P that preserves the Euler characteristic. Here $P : \Gamma \to \Gamma'$ is a homotopy equivalence and $x'_0 = P(x_0)$. The induced map then is just the natural map $P_\# : \pi_1(\Gamma, x_0) \to \pi_1(\Gamma', x'_0)$ which incidentally happens to be an isomorphism in this case.

On the other hand suppose Γ' is obtained from Γ by an AO-move. Using the notation of Definition 2.2.10, we see that Γ' is obtained from Γ by removing an arc p' and adding an arc q. We define a map $P : \Gamma \to \Gamma'$

to be identity on all edges and vertices of Γ that are unchanged by the AO-move. Note that in particular P acts as identity on the end-points of the arc p'. We define P to push p' to the path $p_1^{-1}qp_2^{-1}$. Once again the map $P : \Gamma \to \Gamma'$ is a homotopy equivalence, and induces a natural map $P_\# : \pi_1(\Gamma, x_0) \to \pi_1(\Gamma', x_0')$.

Now we can show that in fact there are exactly k Nielsen equivalence classes of generating n-tuples in G.

Theorem 7.3.4. *Let $\mathcal{R}, \tau, G$ be as defined in the above theorem. Let $(w_1, \ldots, w_n)$ be an n-tuple such that $< w_1, \ldots, w_n >= G$. Then $(w_1, \ldots, w_n) \sim_{NE} (a_{m1}, \ldots, a_{mn})$ for some $1 \leq m \leq k$.*

Proof. Let us assume for now that G is not free and has rank n. Call $\sigma = (w_1, \ldots, w_n)$ and let S_σ be the wedge of circles graph as above, and x_0 be the base vertex. Note that the rank of S_σ is n since if not, then G could have been generated by $< n$ elements. We now inductively define a sequence of graphs with $\Gamma_0 = S_\sigma$. Starting with the graph Γ_s, we fold till we can not any more to obtain the graph Γ_s'. Note that since the image under $\mu_\#$ does not change on folding a graph, we know that the rank of Γ_s' is still n. If Γ_s' is a n-rose, we stop and set $\Gamma_{s+1} = \Gamma_s'$. However, if Γ_s' is not a rose, then observe first that $\mu_\# : \pi_1(\Gamma_s') \to G$ is not injective because G is not free. Thus by Theorem 7.3.2 we can do an AO-move on Γ_s' to get a graph Γ_s'' and we set $\Gamma_{s+1} = \Gamma_s''$. Note that in this case also the rank of Γ_{s+1} is n as an AO-move does not change the Euler characteristic of a graph.

Since in the process described above, we decrease the number of edges, this sequence will terminate in a rose whose base vertex x_0' is the image of x_0 after the various folds or AO-moves. Since the rank of the graphs described above remains n throughout, we will thus end up with an n-rose. If the n labels of this rose are the set A_i for any $1 \leq i \leq k$, we are done. Suppose not. Let the labels be $e_1, \ldots, e_n \in \{a_{11}, \ldots, a_{1n}, \ldots, a_{k1}, \ldots, a_{kn}\}$. Note of course that for any $p \neq q$, we have that $e_p^{\pm 1} \neq e_q^{\pm 1}$. Now there must exist an a_{ij} for some $1 \leq i \leq k$, $1 \leq j \leq n$ which is not the label of a petal in our n-rose. But by Lemma 7.3.3, we have $(w_1, \ldots, w_n) \sim_{\mathrm{NE}} (e_1, \ldots, e_n)$, and thus $G =< e_1, \ldots, e_n >$. In particular $a_{ij} \in G =< e_1, \ldots, e_n >$ and so $a_{ij} = w(e_1, \ldots, e_n)$. However this then yields a contradiction in a manner similar to Theorem 7.3.1 and so we are done. ☐

All that remains to show is that the group G has all the nice qualities described in the main result. This has already been shown in its entirely by Kapovich and Weidmann in [39]. Hence we will only sketch the proof, and will refer the reader to [39] for details.

Theorem 7.3.5. *[39] Let $\mathcal{R}, \tau, G$ be as defined in the above theorem. Then G is a torsion-free word-hyperbolic one-ended group of rank n*

Proof. **Part (1)** We show that all subgroups of G generated by $\leq n-1$ elements are free.

Let $H =< g_1, \ldots, g_k >\leq G$ with $k < n$. We pick a graph Γ that represents H with minimal number of edges of rank $\leq k$. For instance, we could pick $\Gamma = S_\sigma$ where S_σ is the wedge of circles graph and $\sigma = (g_1, \ldots, g_k)$. If the labeling map $\phi : \pi_1(\Gamma) \to G$ is injective, then H is free and we are done. Suppose then that this map is not injective. Then as before, there exists a non-trivial loop in Γ whose label is 1 in G. By Greendlinger's Lemma the label of this path contains a subword v which is also the subword of a cyclic permutation of a defining relator r and $|v| > (1 - 3\lambda)|r|$. Let $p = p_1 \ldots p_s$ be a concatenation of paths in Γ such that $p_2, \ldots, p_{k-1}$ are maximal arcs, and the label of p is v. Then as in the proof of Lemma 7.3.2, we consider various possibilities for the lengths of p_q for $1 \leq q \leq s$, and we see that we can either contradict the fact that the graph Γ was minimal in terms of edges, or contradict one of the genericity conditions. Thus $\phi : \pi_1(\Gamma) \to G$ is injective and $H \cong \pi_1(\Gamma)$.

Part (2) We show that $rank(G) = n$

As we saw earlier in this chapter, $G =< a_{11}, \ldots, a_{1n} >$ for instance, and hence $rank(G) \leq n$. Suppose $s = rank(G) < n$. Then by Part (1) G is free. Hence there exists a graph Γ with edges labeled by elements of $A = A_1 \cup \ldots \cup A_k$ with $rank(\Gamma) \leq n$ such that the image of the labeling homomorphism $\phi : \pi_1(\Gamma) \to G$ is G. Since the rank of Γ is $< n$, there must exist an element $y \in \{a_{11}, \ldots, a_{1n}, \ldots, a_{k1}, \ldots, a_{kn}\}$ such that Γ has no loop labeled by y. W.l.o.g assume that $y = a_{11}$. However, the image of ϕ is all of G and thus there is a path in Γ of length > 1 whose label z is equal to a_{11} in G. Thus $z a_{11}^{-1} =_G 1$. Then once again by using Greendlinger's lemma and small cancellation, we get a contradiction in exactly the same way as Part (1). Thus $rank(G) = n$.

Part (3) We show that G is not free.

We saw earlier that the presentation (**) which was $G =< a_{11}, \ldots, a_{1n} | U_1, \ldots, U_n >$ represented the same group G as in presentation (*). However (**) is a proper quotient of F_n. Since F_n is not isomorphic to any of its proper quotients, F_n is not isomorphic to G. But by Part (2) since $rank(G) = n$, we have that G is not isomorphic to any F_s for $1 \leq s < n$. Thus G is not free.

Part (4) We show that G is torsion free, word-hyperbolic, and one-ended.

Since (*) is a $C'(1/6)$ small cancellation presentation where the defining relators are not proper powers, results from [45] show that G is torsion free, non-elementary, and word-hyperbolic.

It remains to show that G is one-ended. Suppose not. Then G is free product $G = A * B$ with A, B both non-trivial. Now by Grushko's Theorem, $n = rank(A) + rank(B)$ where $1 \leq rank(A), rank(B) \leq n - 1$. Now by Part (1) we get that A, B are both free. Hence $G = A * B$ is also free. This contradicts Part (4). Hence, G is one-ended.

□

Now Theorem E is a direct consequence of Theorem 7.3.1, Theorem 7.3.4, and Theorem 7.3.5.

Chapter 8

Asides and Open Questions

We start with a brief aside regarding the index of the image of a finite index subgroup of a free group in a proper quotient of the free group. We end by discussing some questions with regards to Chapters 1-6.

8.1 Finite Index Subgroups of Proper Quotients of $F(a, b)$

Let $F_2 = F(a, b)$ be the free group of rank 2. Let $r \in F(a, b)$ be a random word as obtained using the methods of Section 5.1. Consider the quotient $G =< a, b|r >= F(a, b)/N$. Let $\beta : F(a, b) \to G = F(a, b)/N$ be the natural surjective group homomorphism that takes any $g \in F(a, b)$ to $Ng \in G$. Let $H \leq F(a, b)$ be a finite index subgroup with index $[F(a, b) : H] = j < \infty$. Then it is clear that $K := \beta(H) \leq G$ is a finite index subgroup of G with $1 \leq [G : K] \leq j$. Let $[G : K] = k$. We want to investigate if in fact we can say more about k.

For the remainder of this section, when we refer to a graph B_m, we will mean the graph in Figure 8.1 on $2m$ vertices for any $m \geq 2$.

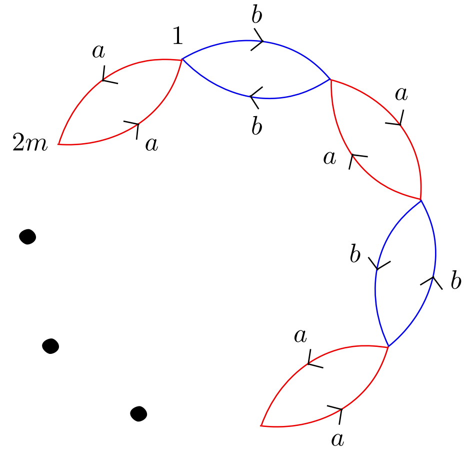

Figure 8.1: The graph B_m

Currently, we can show that:

Proposition 8.1.1. *With terminology as above,*

(1) $k \neq j-1$*.*

(2) if $r \in N' \leq H$ *where* $N' \leq F(a,b)$ *is a finite-index normal subgroup of* $F(a,b)$*, then* $k = j$*.*

(3) if $H \neq B_m$ *for any* m*, then with some positive probability* $k = 1$*.*

Proof. Let Γ be a graph representing H. Then $|V\Gamma| = j$. Observe first that checking for possibilities of the index $[G : K]$ amounts to checking for the number of vertices in principal quotients of Gamma (see Definition d:pq). In the proof below by an x-edge we mean an edge of the graph labeled by x.

Part (1) The index in G falling by exactly 1 would mean that we consider the principal quotient obtained by identifying precisely two vertices of Γ. Further even after folding in this case, the number of vertices should remain $j-1$. Let vertex u, v be identified to a vertex w. This gives us the Figure 8.1. Observe that since the number of vertices can not decrease on folding, we must have that for instance, the terminal vertex of the outgoing b-edge from u must coincide with the terminal vertex of the outgoing b-edge from v. However, in this case since $u \neq v$, we get that Γ had a vertex with two incoming b-edges. This can not happen as Γ is folded. Hence the $k \neq j-1$.

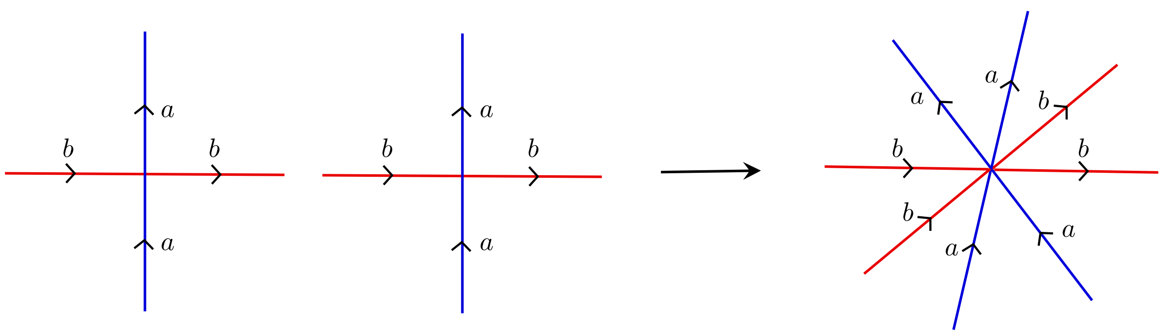

Figure 8.2: Identifying two vertices in Γ

Part (2) Suppose that H and Hx are distinct cosets. If we can show that $\beta(H)$ and $\beta(H)\beta(x)$ are distinct then we are done as in this case $[G : K] \geq j \Rightarrow [G : K] = j$. Suppose instead that $\beta(H) = \beta(H)\beta(x)$. Then there exists an $h \in H$ such that $hx^{-1} =_G 1$. That is, hx^{-1} lies in the normal closure of r in $F(a,b)$, and hence, $hx^{-1} = u_1 r_1 u_1^{-1} \dots u_m r_m u_m^{-1}$ where $m \geq 1$, and for each $1 \leq i \leq m$, we have $r_i \in \{r, r^{-1}\}$, $u_i \in F(a,b)$. Since $r \in N'$ and since N' is normal, we have that $u_i r_i u_i^{-1} \in N'$ for each i. Hence we get that $hx^{-1} \in N' \subseteq H$. This then shows that $x \in H$ which contradicts the fact that H and Hx were distinct cosets.

It merits to say that this case occurs with a finite positive probability. Observe first that since H is a finite index subgroup of $F(a, b)$, we have that $r \in H$ with probability $1/k$. Further, since H is finite index, there exists a subgroup $N' \leq H \leq F(a, b)$ where N' is a finite-index normal subgroup of $F(a, b)$. For instance, here we could take N' to be the normal core of H. Now once again since N' has finite index, with a finite probability, $r \in N'$.

Part (3) Observe that for j to equal k, we must identify all the vertices in Γ. If Γ has a loop edge labeled by a (or b) at a vertex v, we may simply identify the vertices v and v' where there is an edge labeled by b (or a) from vertex v to v'. In this case, the graph will fold to the rose, and we will be done. On the other hand, if there exist vertices $v, v' \in \Gamma$ such that there is an edge between v and v' labeled by a and another edge between v and v' labeled by b, then identifying again on identifying v and v', the graph will fold to a rose and we are done. In case neither of the these two conditions hold, then in fact $H = B_m$ for some m. Note that in particular, if k is odd, then certainly with some probability, $j = 1$.

□

We can actually say some pretty interesting things in a rather specific case. For any $n \geq 2$, let H_n be the normal subgroup represented by the graph C_n on n vertices shown in Figure 8.1.

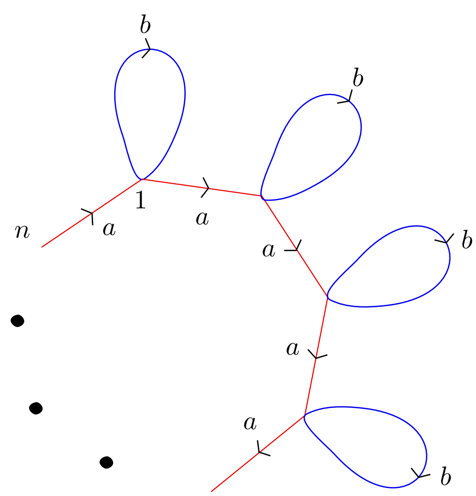

Figure 8.3: The graph C_n

In this case, the distinct cosets of $H = H_n$ in $F(a, b)$ are $H, Ha, Ha^2, \ldots, Ha^{n-1}$. Also, if $Hr = H$ i.e. if $r \in H$, then by Proposition 8.1.1 Part (2), we get that in this case $j = n$. If $Hr \neq H$, then $\beta(H) = \beta(Hr)$ since in this case for $h = 1$, we get that $hr^{-1} =_G 1$ and so $\beta(r) \in \beta(H)$. Thus once you make the identification of joining the base vertex to the vertex where r terminates, you get a normal subgroup H' where $H'r = H'$ and thus the index does not fall any further after this identification (and folding).

For any $n \geq 2$, and m a factor of n, let $p_n(m)$ denote the probability that $[G : \beta(H_n)] = m$. Let $n = p_1^{\alpha_1} p_2^{\alpha_2} \dots p_s^{\alpha_s}$ be the prime decomposition of n. Let $m = p_1^{\beta_1} p_1^{\beta_1} \dots p_t^{\beta_t}$ be a factor of n such that $\alpha_i > \beta_i$ strictly for all $1 \leq i \leq s$. Then using number theory and studying the "figure 8" diagrams we achieve after doing this identification, we get that

$$p_n(m) = \frac{\phi(n)}{nm}$$

where $\phi(n)$ is the Euler phi function. That is $\phi(n)$ counts the number of positive integers $\leq n$ that are relatively prime to n. Further, if $n = q^l$ where q is a prime, then we have: $p_n(q^l) = \frac{1}{q^l}$ and

$$p_n(q^\alpha) = \frac{(q-1)}{q^{\alpha+1}}$$

for $0 \leq \alpha \leq l-1$. Alternatively, if $n = q^l$ where q is a prime, then we have: $p_n(q^l) = \frac{1}{q^l}$ and

$$p_n(q^\alpha) = \frac{\phi(q^l)}{q^{l+\alpha}}$$

for $0 \leq \alpha \leq l-1$. If we enumerate the vertices of H_n as $\{1, \dots, n\}$, and if r terminates at the vertex j then we get that:

$$[G : \beta(H_n)] = \begin{cases} \gcd(n, j-1), & 2 \leq j \leq \lfloor \frac{n}{2} \rfloor + 1 \\ \gcd(n, n-j+1), & \lfloor \frac{n}{2} \rfloor + 2 \leq j \leq n \end{cases}$$

There are a couple of questions that are raised by this:

Q1 Fix a prime q. For $j \in \{0, 1, \dots, l-1\}$, let $p_{q^l}(\frac{j}{l}) = \frac{(q-1)}{q^{j+1}}$, and let $p_{q^l}(\frac{l}{l}) = \frac{1}{q^l}$. Then as $l \to \infty$ does this distribution approach a continuous distribution?

Q2 For $n = q^l$ where q is a prime, would a generating function like $\sum_{\alpha,l} x^{q^\alpha} y^{q^l} \frac{q-1}{q^{\alpha+1}}$ or $\sum_{\alpha,l} x^\alpha y^l \frac{q-1}{q^{\alpha+1}}$ tell us anything interesting?

We now discuss some interesting questions that remain open with regards to Chapters 1-6 in this thesis.

8.2 True Behavior and Asymptotics of Our Functions

As we have seen, while upper bounds for the functions we considered in Chapter 1 are roughly linear, lower bounds remain roughly logarithmic. This then raises the following two questions:

- Q1 What are the true asymptotics of the functions $f^{nfill}_{\Sigma,\rho}, f_{\Sigma,\rho}(L), f_{nfill}(n), f_{simp}(n), f_{prim}(n)$?
- Q2 How do we understand the true behavior of $\deg^{nfill}_{\Sigma,\rho}(\gamma_n), \deg_{\Sigma,\rho}(\gamma_n)$, $d_{nfill}(w_n), d_{simp}(w_n)$, $d_{prim}(w_n)$ for "natural sequences" of curves/elements? For instance, nothing is known even about the behavior of $d_{prim}(a^n b^n)$ in the free group $F(a, b)$.

On a different note, a recent paper of Puder [52] (see also [53, 54] for related work) introduces the notion of a *primitivity rank* $\pi(g)$ for an element $g \in F_N$. Namely, $\pi(g)$ is defined as the smallest rank of a subgroup $H \leq F_N$ such that $g \in H$ but g is not primitive in H. Puder proves in [52, Corollary 4.2] that for an element $g \in F_N$ one has either $\pi(g) = \infty$ or $0 \leq \pi(g) \leq N$, and that every integer between 0 and N does occur as a value of $\pi(g)$ for some g. He also defines and studies the primitivity rank $\pi(H)$ for a finitely generated subgroup $H \leq F_N$, where $\pi(H)$ is defined as the minimum rank of J such that $H \leq J \leq F_N$ and that H is not a proper free factor of J. These notions are related to and in some sense dual to our definitions of $d_{prim}(g)$ and $d_{simp}(g)$, but the precise connection of our results with Puder's work remains to be understood. Malestein and Putman [47] obtained a number of lower bound results (in terms of k) for the minimal self-intersection number of nontrivial elements of the k-term of the lower central series and the derived series of a surface group. It would be interesting to see if their methods can be used to obtain lower bounds for the function $f_{\Sigma,\rho}$. It would also be interesting to investigate if looking inside the lower central series and the derived series of F_N may produce new lower bounds for $f_{prim}(n)$ and $f_{simp}(n)$. Finally, Bou-Rabee raised the important point that currently for a nonabelian free group G, the best upper and lower bounds for $f_{prim}(n)$ and $\mathrm{RF}_G(n)$ have the same asymptotic behavior. It remains to be seen if $f_{prim}(n)$ and $\mathrm{RF}_G(n)$ also have the same asymptotic behavior.

References

[1] G. Arzhantseva, *On groups in which subgroups with a fixed number of generators are free*, (Russian) Fundam. Prikl. mat. **3**(3) (1997), 675683

[2] G. Arzhantseva, *A dichotomy for finitely generated subgroups of word-hyperbolic groups*, preprint, University of Geneve, 2001; http://www.unige.ch/math/biblio/preprint/ pp2001.html

[3] G. Arzhantseva and A. Ol'shanskii, *Generality of the class of groups in which subgroups with a lesser number of generators are free*, (Russian) Mat. Zametki **59** (1996), no. 4, 489–496; translation in: Math. Notes **59** (1996), no. 3-4, 350–355

[4] A. Basmajian, *Universal length bounds for non-simple closed geodesics on hyperbolic surfaces*, J. Topol. **6** (2013), no. 2, 513–524

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

[6] I. Biringer, K. Bou-Rabee, M. Kassabov, F. Matucci, *Intersection growth in groups*, preprint, 2013; arXiv:1309.7993; to appear in Transact. Amer. Math. Soc.

[7] F. Bonahon, *The geometry of Teichmüller space via geodesic currents*, Invent. Math. **92** (1988), no. 1, 139–162

[8] K. Bou-Rabee, *Quantifying residual finiteness*, J. Algebra **323** (2010), no. 3, 729–737

[9] K. Bou-Rabee and Y. Cornulier, *Systolic growth of linear groups*, Proc. Amer. Math. Soc., to appear; 1408.5837.pdf

[10] K. Bou-Rabee, M. Hagen and P. Patel, *Residual finiteness growths of virtually special groups*, Math. Z. **279** (2015), no. 1-2, 297–310

[11] K. Bou-Rabee, and T. Kaletha, *Quantifying residual finiteness of arithmetic groups*, Compos. Math. **148** (2012), no. 3, 907–920

[12] K. Bou-Rabee, and D. B. McReynolds, *Bertrand's postulate and subgroup growth*, J. Algebra **324** (2010), no. 4, 793–819

[13] K. Bou-Rabee, and D. B. McReynolds, *Asymptotic growth and least common multiples in groups*, Bull. Lond. Math. Soc. **43** (2011), no. 6, 1059–1068

[14] K. Bou-Rabee, and D. B. McReynolds, *Extremal behavior of divisibility functions*, Geom. Dedicata **175** (2015), 407–415

[15] K. Bou-Rabee and D. B. McReynolds, *Characterizing linear groups in terms of growth properties*, Michigan Math. J., to appear; arXiv:1403.0983

[16] K. Bou-Rabee and A. Myropolska, *Groups with near exponential residual finiteness growth*, preprint, 2015; arXiv:1509.01372

[17] K. Bou-Rabee and B. Seward, *Arbitrarily large residual finiteness growth*, J. Reine Angew. Math., to appear; arXiv:1304.1782

[18] K. Bou-Rabee and D. Studenmund, *Full residual finiteness growths of nilpotent groups*, Israel J. Math., to appear; arXiv:1406.3763

[19] N. V. Buskin, *Efficient separability in free groups* (Russian), Sibirsk. Mat. Zh. **50** (2009), no. 4, 765–771; translation in Sib. Math. J. **50** (2009), no. 4, 603–608

[20] D. Calegari and J. Maher, *Statistics and compression of scl*, Ergodic Theory Dynam. Systems **35** (2015), no. 1, 64–110

[21] C. Cashen and J. Manning, *Virtual geometricity is rare*, LMS J. Comput. Math. **18** (2015), no. 1, 444–455

[22] M. Culler, K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, pp 91119

[23] J. Gaster, *Lifting curves simply*, preprint, January 2015; arXiv:1501.00295; to appear in Internat. Math. Res. Notices

[24] R. Gimadeev and M. Vyalyi, *Identical relations in symmetric groups and separating words with reversible automata*, Computer science theory and applications, Lecture Notes in Comput. Sci., vol. 6072, Springer, Berlin, 2010, pp. 144–155

[25] M. Greendlinger, *Dehn's algorithm for the word problem*, Communications in Pure and Applied Mathematics, **13** (1960), pp. 6783

[26] M. Greendlinger, *On Dehn's algorithms for the conjugacy and word problems, with applications*, Communications in Pure and Applied Mathematics, **13** (1960), pp. 641677

[27] M. Greendlinger, *An analogue of a theorem of Magnus*, Archiv der Mathematik, **12** (1961), pp. 94-96

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

[29] N. Gupta and I. Kapovich, *The primitivity index function for a free group, and untangling closed curves on hyperbolic surfaces*, preprint, November 2014; http://arxiv.org/abs/1411.5523

[30] M. Hall, *Coset representations in free groups*, Trans. Amer. Math. Soc. **67** (1949), 421–432

[31] I. Kapovich, *The frequency space of a free group*, Internat. J. Alg. Comput. **15** (2005), no. 5-6, pp. 939969

[32] I. Kapovich, *Currents on free groups*, Topological and Asymptotic Aspects of Group Theory (R. Grigorchuk, M. Mihalik, M. Sapir and Z. Sunik, Editors), AMS Contemporary Mathematics Series, vol. 394, 2006, pp. 149-176

[33] I. Kapovich, *Clusters, currents and Whitehead's algorithm*, Experimental Mathematics **16** (2007), no. 1, pp. 67-76

[34] I. Kapovich and M. Lustig, *Intersection form, laminations and currents on free groups*, Geom. Funct. Anal. **19** (2010), no. 5, pp. 1426-1467

[35] I. Kapovich and A. Myasnikov, *Stallings foldings and the subgroup structure of free groups*, J. Algebra **248** (2002), no 2, pp. 608–668

[36] I. Kapovich and C. Pfaff, *A Train track Directed Random Walk on $Out(F_r)$*, Internat. J. Algebra Comput. **25** (2015), no. 5, 745–798

[37] I. Kapovich, P. Schupp, *Genericity, the Arzhantseva-Olshanskii method and the isomorphism problem for one-relator groups*, Math. Ann. **331**, 119, 2005

[38] I. Kapovich, P. Schupp, and V. Shpilrain, *Generic properties of Whitehead's algorithm and isomorphism rigidity of random one-relator groups*, Pacific J. Math. **223** (2006), no. 1, 113–140

[39] I. Kapovich, R. Weidmann, *Nielsen Equivalence in a Class of Random Groups*, to appear in Journal of Topology ; arXiv:1309.7458

[40] I. Kapovich, R. Weidmann, *Nielsen Equivalence in Small Cancellation Groups*, preprint, 2012, arXiv:1011.5862

[41] M. Kassabov and F. Matucci, *Bounding the residual finiteness of free groups*, Proc. Amer. Math. Soc. **139** (2011), no. 7, 2281–2286

[42] G. Kozma and A. Thom, *Divisibility and Laws in Finite Simple Groups*, to appear in Math. Ann. ; arXiv:1403.2324

[43] M. Lustig, *Conjugacy and Centralizers for Iwip Automorphisms of Free Groups*, Geometric group theory: Geneva and Barcelona conferences (G. Arzhantseva, L. Bartholdi, J. Burillo and E. Ventura, Editors), Trends in Mathematics, Birkhäuser Verlag Basel/Switzerland, 2007, pp. 197224

[44] A. Lubotzky, and D. Segal, *Subgroup growth*, Progress in Mathematics, 212. Birkhäuser Verlag, Basel, 2003

[45] R. Lyndon and P. Schupp, *Combinatorial group theory*, Reprint of the 1977 edition. Classics in Mathematics. Springer-Verlag, Berlin, 2001

[46] J. McCool, *Some finitely presented subgroups of the automorphism group of a free group*, J. Algebra **35** (1975), 20513

[47] J. Malestein, and A. Putman, *On the self-intersections of curves deep in the lower central series of a surface group*, Geom. Dedicata **149** (2010), 73–84

[48] A. Myasnikov, and V. Shpilrain, *Automorphic orbits in free groups*, J. Algebra **269** (2003), no. 1, 18–27

[49] J. Nielsen, *Die Isomorphismen der allgemeinen, unendlichen Gruppe mit zwei Erzeugenden* , Math. Ann. **78** (1917), no. 1, 385-397

[50] J. Nielsen, *Über die Isomorphismen unendlicher Gruppen ohne Relation*, Math. Ann. **79** (1918), no. 3, 269-272

[51] P. Patel, *On a theorem of Peter Scott*, Proc. Amer. Math. Soc. **142** (2014), no. 8, 2891–2906

[52] D. Puder, *Primitive words, free factors and measure preservation*, Israel J. Math. **201** (2014), no. 1, 25–73

[53] D. Puder, and O. Parzanchevski, *Measure preserving words are primitive*, J. Amer. Math. Soc. **28** (2015), no. 1, 63–97

[54] D. Puder, and C. Wu, *Growth of primitive elements in free groups*, J. Lond. Math. Soc. (2) **90** (2014), no. 1, 89–104

[55] E. Rips, *Subgroups of small cancellation groups*, Bull. Lond. Math. Soc. **14**, 45-47, 1982

[56] I. Rivin, *Geodesics with one self-intersection, and other stories*, Adv. Math. **231** (2012), no. 5, 2391–2412

[57] A. Roig, E. Ventura, and P. Weil, *On the complexity of the Whitehead minimization problem*, Internat. J. Algebra Comput. **17** (2007), no. 8, 1611–1634

[58] S. Schleimer, *Polynomial-time word problems*, Comment. Math. Helv. **83** (2008), no. 4, 741–765

[59] P. Scott, *Subgroups of surface groups are almost geometric*, J. London Math. Soc. (2), **17** (1978), no. 3, 555–565

[60] P. Scott, *Correction to: "Subgroups of surface groups are almost geometric"* [J. London Math. Soc. (2) 17 (1978), no. 3, 555–565], J. London Math. Soc. (2) **32** (1985), no. 2, 217–220

[61] B. Solie, *Algorithmic and statistical properties of filling elements of a free group, and quantitative residual properties of* Γ*-limit groups*, PhD Thesis, (2011)

[62] B. Solie, *Genericity of filling elements*, Internat. J. Algebra Comput. **22** (2012), no. 2

[63] J. R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), 552-565

[64] J. R. Stallings, *Whitehead graphs on handlebodies*, Geometric group theory down under (Canberra, 1996), 317–330, de Gruyter, Berlin, 1999

[65] R. Stong, *Diskbusting elements of the free group*, Math. Res. Lett. **4** (1997), no. 2-3, 201–21

[66] M.-P. Schützenberger, *Sur l'Equation* $a^{2+n} = b^{2+m}c^{2+p}$ *dans un groupe libre*, C. R. Acad. Sci. Paris **248** (1959), 2435–2436.

[67] S. Wang, and B. Zimmermann,*The maximum order of finite groups of outer automorphisms of free groups*, Math. Z. **216**(1994), no. 1, 837

[68] J. H. C. Whitehead, *On equivalent sets of elements in a free group*, Ann. of Math. (2) **37**(1936), no. 4, 782–800