

A RECOGNITION THEOREM FOR POLYNOMIAL
GROWTH OUTER AUTOMORPHISMS OF THE FREE
GROUP

By Gregory MacLean Schinke Fein

A dissertation submitted to the Graduate School-Newark

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Graduate Program in Mathematical Sciences

written under the direction of

Professor Mark E. Feighn

and approved by

Newark, New Jersey

May, 2013

©2013

Gregory MacLean Fein

ALL RIGHTS RESERVED

Abstract

A Recognition Theorem for Polynomially Growing Outer

Automorphisms of the Free Group

By Gregory MacLean Schinke Fein

Dissertation Director: Professor Mark E. Feighn

Feighn and Handel's recognition theorem for $Out(F_n)$ provides invariants that canonically determine any forward rotationless outer automorphism of the free group. We ask to what extent those invariants can be extended to outer automorphisms with some periodic behavior. Many of the same constructions do not have natural analogs, in particular because of the possible lack of principal representatives in this setting. However, by restricting our attention to polynomial growth outer automorphisms and using train track technology, we are able to find a special set of lines in the free group that encode all the dynamical information of these non-forward rotationless maps.

*To Amanda for all of her support,
and to my grandmother, who has waited a long time for this day.*

Acknowledgments:

First and foremost I must thank my advisor, Mark Feighn, for the mountains of time and thought he gave me in my effort to complete this thesis. He is a man of endless patience, and he has the uncanny ability to see a new way around seemingly any mathematical problem, thereby restoring his students' confidence in their work and ability. For these things I am eternally grateful. Thanks also must go to Lee Mosher for many illuminating and helpful conversations; he was always willing to give me time when I had questions, and his answers saved me tons of time and helped me avoid plenty of headaches. I also thank Jane Gilman and Ulrich Oertel who talked with me early in my graduate career and showed me some of the many topics that topology and geometry have to offer. I thank Sylvain Cappell for first instilling in me a love of topology, and for pointing me toward a graduate school and department full of the great people listed above who would help me to continue exploring my favorite field. I thank William Keigher, Tim McDonald, David Keys, and John Loftin for helping me become a stronger teacher of mathematics, thereby lighting up opportunities along my future career path. And I thank Ara Basmajian and Jason Behrstock who, through their seminars and courses, have helped make CUNY's Graduate Center my second mathematical home,

a place filled with friends and colleagues.

At my first mathematical home, Rutgers-Newark, my time has been made much easier thanks to the efforts of Zhengyu Mao, Susan Tel, and Jane Manay, who, as Graduate Director and administrators, have helped smooth out countless bureaucratic wrinkles. I thank them for all of the time they put in for me and for all of my fellow R-N graduate students. Among those students, I must particularly thank Andrew Silverio, Pei Wang, and Pritam Ghosh for helping me expand my mathematical horizons through the topics covered in our group theory student seminar. And I thank the other R-N students for the talks they contributed to our general student seminar as well. I also owe an enormous debt of gratitude to one former R-N student: Karan Puri. He has served as a guide for me all through graduate school, from my first day as a teaching assistant to my current journey on the job market. There is no way I can repay him all of the time and thought he has put in for me.

Finally, these acknowledgments would be woefully incomplete if I did not thank my partner, Amanda Fein Schinke. She deserves to receive a doctorate herself for all of the work she has done. She took on extra responsibilities at home when I needed time to work, and she reassured me in my times of greatest doubt. She has been a constant source of strength, and I am positive that this thesis would not have been completed without her.

Contents

1	Introduction	1
2	Organization	2
3	Definitions and Preliminaries	2
3.1	Free Groups and Automorphisms	2
3.2	Growth Rates	5
3.3	Fixed Sets	7
3.4	Marked Graphs and Topological Representatives	7
3.5	Lines in the Free Group	11
3.6	Principal Representatives	14
3.7	Linear Growth	16
3.8	The Forward Rotationless Recognition Theorem	19
3.9	Free Factor Systems and Filtrations	22
3.10	Relative Train Track Maps	28
3.11	Stallings Graphs	34
4	Guesses and Counterexamples	39
5	Annular Fix_N-Lines	44
6	The Main Theorem	61

List of Figures

1	The Cayley graph of F_2	4
2	A Linear Growth Mapping Class	18
3	The Stallings graph $\Gamma_{R_4}(\pi)$ of Counterexample 1	40
4	A homeomorphism of a surface inducing Counterexample 1 . .	41
5	A graph G for Counterexample 2	42
6	The stallings graph $\Gamma_G(\pi)$ of Counterexample 2	43
7	A sequence of components of the Stallings graph of π	53
8	A line of the form $\dots \tilde{w}_1^{-1} \tilde{w}_1^{-1} \tilde{w}_1^{-1} \tilde{g} \tilde{w}_1 \tilde{w}_1 \tilde{w}_1 \dots$ in \tilde{G}	53
9	A sequence of components of the Stallings graph of π when $p(E_{i-1}) = E_{i-1} w_{i-1}^{t_{i-1}}$ and $p(E_i) = E_i w_i^{t_i}$	55
10	A sequence of components of the Stallings graph of π when $p(E_{i-1}) = w_{i-1}^{t_{i-1}} E_{i-1}$ and $p(E_i) = w_i^{t_i} E_i$	55
11	A sequence of components of the Stallings graph of π when $p(E_{i-1}) = w_{i-1}^{t_{i-1}} E_{i-1}$ and $p(E_i) = E_i w_i^{t_i}$	56
12	A sequence of components of the Stallings graph of π when $p(E_{i-1}) = E_{i-1} w_{i-1}^{t_{i-1}}$ and $p(E_i) = w_i^{t_i} E_i$	56
13	Two components of $\Gamma_G(\pi)$ with lifts of the Nielsen path u . . .	71
14	A lift of an annular Fix_N -line to \tilde{G} in case (a)	71
15	A lift of an annular Fix_N -line to \tilde{G} in case (a)	72
16	A graph G on which to demonstrate the unwinding homotopy	73

1 Introduction

There is a strong analogy between $Out(F_n)$ and another gem of geometric group theory: the mapping class group of a surface. Mapping class groups are groups of self maps in which one ignores differences via homotopy, and much greater detail on these groups can be found in Farb and Margalit's book on the subject [FaM11]. The analogy with $Out(F_n)$ is also explored in a survey paper by Bestvina [B02]. The close relationship between these two collections of groups has resulted in many theorems in one inspiring searches for analogs in the other.

In “The Recognition Theorem for $Out(F_n)$ ” [FH11], Feighn and Handel drew on Nielsen's approach to the mapping class group [HT85] to provide a set of qualitative and quantitative invariants that, taken together, uniquely and canonically determine an outer automorphism. Their theorem applies only to *forward rotationless* outer automorphisms; these are, roughly, outer automorphisms which have no periodic behavior. Our is goal to generalize this result to all of $Out(F_n)$.

In the study of $Out(F_n)$, it is often useful to categorize elements based on their growth rates under iteration. Any ϕ in $Out(F_n)$ can be described as either *of at most polynomial growth* (PG) or *exponential growth* (EG), depending on how quickly word lengths in the free group grow under an automorphism Φ of F_n that represents the class ϕ . By focusing on PG outer automorphisms, we find a ϕ -invariant set of lines, known as *annular Fix_N -lines*, so that ϕ is canonically determined by its action on this set.

2 Organization

The organization of this work is as follows. In Section 3 we review the necessary background, including free groups, automorphisms, fixed sets, the forward rotationless recognition theorem, and train track maps. It ends by describing a special class of graphs called Stallings graphs. Much of this background can be found in [FH11], though [BFH00] and [BH92] were also used as resources. The only exception is Section 3.11, where some of the definitions are new, and so we needed to prove a few facts about these new objects. In Section 4, we suggest two fairly reasonable guesses as to what the general recognition theorem might look like, and we provide counterexamples to show that both of these guesses are incorrect. The counterexamples will be analyzed via Stallings graphs. In Section 5, we define the main construction of this thesis: the set of *annular Fix_N -lines* of a forward rotationless PG outer automorphism. As we did with the counterexamples, we will visualize these lines using Stallings graph machinery. In the final section, Section 6, we will prove our main theorem: that any PG outer automorphism is uniquely determined by its action on the set of annular Fix_N -lines of a forward rotationless power.

3 Definitions and Preliminaries

3.1 Free Groups and Automorphisms

Let F_n be the *free group on n generators*, with basis denoted $\{a, b, c, \dots\}$. An element w of the free group is a *finite reduced word* w in these generators and

their inverses, for instance a , b^2 , or $a^3b^{-1}c^6bcba$. The word *reduced* means that w contains no basis element adjacent to its own inverse, and the word *finite* refers to the finitely many appearances of the generators and their inverses. The group operation on these words is concatenation, followed by the reduction of adjacent pairs of generators and inverses. For convenience, the inverse w^{-1} of a word w will sometimes be denoted \bar{w} and sometimes W .

Let $Aut(F_n)$ be the group of automorphisms of the free group. An automorphism of the free group can be described by listing the images under the automorphism of a set of generators. For example, the following is a fairly simple element of $Aut(F_2)$.

Example 3.1.

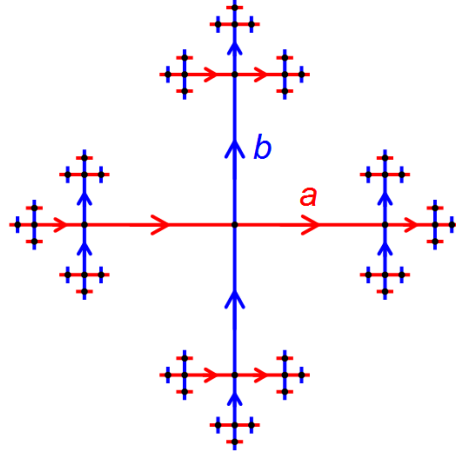
$$\Phi : a \mapsto a \quad b \mapsto ba$$

An *inner automorphism* is an automorphism i_x that conjugates every element of the free group by some fixed element x ; i.e. for all $w \in F_n$, $i_x(w) = xwx$. The set of all inner automorphisms forms a normal subgroup of $Aut(F_n)$, denoted $Inn(F_n)$. The resultant quotient group $Aut(F_n)/Inn(F_n)$ is called the *outer automorphism group of the free group* and is denoted $Out(F_n)$. Because we are modding out by conjugations to form $Out(F_n)$, individual outer automorphisms (as they are called) act not on elements of F_n but on conjugacy classes of elements of F_n .

The *boundary of F_n* , ∂F_n , is the set of reduced infinite words $P = y_1y_2y_3 \dots$ where each y_i is an element of the chosen basis or the inverse of a basis element. This is equivalent to the definition of the Gromov boundary of F_n as

a hyperbolic group, where boundary points are given by geodesic rays in the Cayley graph of F_n . Full details on the construction of the boundary can be found in [BrH99, Chapter III.H] and [BK02].

Figure 1: The Cayley graph of F_2



A boundary point is called *periodic* if it takes the form $wyyy\dots$ for some $w, y \in F_n$. The free group acts on its own boundary in the following way: for any $x \in F_n$, x sends $y_1y_2y_3\dots$ to the reduced infinite word that results from reducing $xy_1y_2y_3\dots$. The quotient of ∂F_n by this action is denoted $F_n \backslash \partial F_n$, and elements of this set are called *classes of rays*. We will usually denote a class of rays by ρ .

The boundary of the free group can be topologized in the following way: for $P = y_1y_2y_3\dots$ a reduced infinite word in ∂F_n and $t > 0$, define $V(P, t)$ to be the set of all $Q = z_1z_2z_3\dots$ so that $z_i = y_i$ for all $i \leq t$. The collection of sets $V(P, t)$ as P ranges over ∂F_n and t ranges over \mathbb{N} form a basis for the compact-open topology on ∂F_n . Any automorphism Φ of F_n induces a homeomorphism

$\partial\Phi$ of ∂F_n with this topology [F80]. The homeomorphism ∂i_x induced by an inner automorphism i_x is the one that sends $y_1 y_2 y_3 \dots$ to the reduction of $(x y_1 x^{-1})(x y_2 x^{-1})(x y_3 x^{-1}) \dots$, which is the same as the reduction of $x y_1 y_2 y_3 \dots$. In other words, $\text{Inn}(F_n)$'s action on ∂F_n is the same as F_n 's action, and so there is an induced action of $\text{Out}(F_n)$ on $F_n \backslash \partial F_n$.

Lemma 3.2. *Let $P \in \partial F_n$ be non-periodic, let $\rho \in F_n \backslash \partial F_n$ be the class of P , and let $\phi \in \text{Out}(F_n)$ so that $\phi(\rho) = \rho$. Then there is a unique representative Φ of ϕ so that $\partial\Phi(P) = P$.*

Proof. If Φ is any representative of ϕ , then by assumption, $\partial\Phi(P) = xP$ for some $x \in F_n$. Under the action of F_n on its own boundary, xP can only be equal to P if $P = xxx \dots$ or $x^{-1}x^{-1}x^{-1} \dots$, and both of these possibilities are ruled out by P being non-periodic. But $\partial i_x^{-1}\Phi(P)$ is equal to P (as $x^{-1}xP$ reduces to P), and so $i_x^{-1}\Phi$ is the only representative of ϕ whose induced homeomorphism of ∂F_n fixes P . \square

3.2 Growth Rates

One useful way to classify outer automorphisms is by their *growth rates*. For any $\Phi \in \text{Aut}(F_n)$, any $x \in F_n$, and any integer j , the *growth function of Φ with respect to x* , $g_{\Phi,x}(j)$, is defined to be the word length of $\Phi^j(x)$.

When $g_{\Phi,x}(j)$ can be bounded below by an exponential function in j , we say that x is *exponentially growing*, or *EG*, under Φ . When $g_{\Phi,x}(j)$ can be bounded above by a polynomial in j , we say that x is *polynomially growing*, or *PG*,

under Φ . If $g_{\Phi,x}(j)$ can be bounded above by a linear function in j , we say x is *linearly growing*, or *LG*, under Φ . Finally, if $g_{\Phi,x}(j)$ can be bounded above by a constant function M , then we say that x *has zero growth* under Φ , or is *ZG*. Because there are only finitely many words of length at most M , x is ZG under Φ if and only if it is periodic under Φ .

Each of these terms can be extended to apply to automorphisms of F_n rather than elements of F_n . In particular, if $g_{\Phi,x}(j)$ can be bounded by a constant for every choice of x , then we say that Φ *has zero growth*, or is *ZG*. This is of course equivalent to Φ being a finite order map by what we said above. Similarly, one can define what it means for Φ to be *linearly growing* (LG) or more generally *polynomially growing* (PG). If there exists an element of the free group that grows exponentially under Φ , then we say that Φ *has exponential growth* (EG), though in general an automorphism Φ may have mixed growth, meaning that there are elements of the free group of all different growth rates under Φ .

We say a conjugacy class α is *ZG* (respectively *LG*, *PG*, *EG*) under an outer automorphism ϕ if a cyclically reduced representative a of α is ZG (respectively LG, PG, EG) under some representative Φ of ϕ . Lastly, we say that $\phi \in \text{Out}(F_n)$ is *ZG* (respectively *LG*, *PG*, *EG*) if any (or equivalently, every) representative of ϕ is ZG (respectively LG, PG, EG).

Going forward, we will focus on PG outer automorphisms, particularly when we reach the statement of our main theorem. Much of what follows applies to all of $\text{Out}(F_n)$, but when it is convenient for simplicity, we will give definitions

and background lemmas in the PG case only.

3.3 Fixed Sets

For any $\Phi \in \text{Aut}(F_n)$, we define $\text{Fix}(\Phi)$ to be the subgroup of elements of F_n that are fixed by Φ . For instance, the fixed subgroup of the automorphism in Example 3.1 is $\langle a, baB \rangle$. But we can go even further: we define $\text{Fix}(\partial\Phi)$ to be the subset of ∂F_n fixed by the homeomorphism $\partial\Phi$ induced by Φ . We say $P \in \text{Fix}(\partial\Phi)$ is *attracting for Φ* or is an *attractor for Φ* if there is a neighborhood $V = V(P, t)$ of P such that $\partial\Phi(V) \subset V$ and such that $\bigcap_{k=1}^{\infty} \partial\Phi^k(V) = \{P\}$. A *repeller for Φ* is an attractor for Φ^{-1} . We define $\text{Fix}_N(\partial\Phi)$ to be the subset of $\text{Fix}(\partial\Phi)$ consisting of all fixed points that are not repelling, and we define $\text{Fix}_+(\partial\Phi)$ to be the set of attracting fixed points for $\partial\Phi$. Every element of $\text{Fix}_+(\partial\Phi)$ is isolated in $\text{Fix}(\partial\Phi)$ [BFH04, Lemma 2.5].

3.4 Marked Graphs and Topological Representatives

While these fixed boundary points and subgroups are useful tools, throughout much of the history of the study of $\text{Out}(F_n)$, individual outer automorphisms have been analyzed not directly by their action on the free group, but by thinking of them in a more topological light. In particular, we replace F_n by a rank n graph, and we represent the outer automorphism by a homotopy equivalence of this graph. Of course, there are many ways to identify the fundamental group of a graph with F_n , and this choice will impact the outer automorphism represented.

A *graph* is a one-dimensional CW complex, composed of 0-cells called *vertices* which are connected by 1-cells called *edges*. Any connected graph has free fundamental group, and for any finite, connected graph G , the *rank of G* , $rk(G)$, is the rank of $\pi_1(G)$. If $rk(G) = 0$ (meaning that G is contractible), then we say that G is a *tree*. For example, the universal cover of any connected graph is a tree. A *forest* is a disjoint union of trees.

Each edge E in G is oriented from one vertex to another. Its initial vertex is denoted $\iota(E)$, and its terminal vertex is denoted $\tau(E)$. The same edge with the opposite orientation is denoted \bar{E} , though sometimes to avoid conflicting with other notation, we will denote it by E^{-1} . For any vertex v of G , a *direction at v* is the germ of an initial segment of an oriented edge E with $\iota(E) = v$. The set of directions at v is denoted $T_v(G)$.

Let R_n be the *rose with n petals*, meaning a graph with a single vertex v and n edges. Assume that F_n has been identified with $\pi_1(R_n, v)$. A *marked graph* is a pairing (G, m) , where G is a rank n graph with every vertex of valence at least two, and $m : R_n \rightarrow G$ is a homotopy equivalence. If we let $b = m(v)$, then m gives an identification between F_n and $\pi_1(G, b)$. We will often suppress the marking map m , and say that G is a marked graph. The universal cover \tilde{G} of G is often called a *marked tree*, and we can equivalently think of the marking as a continuous function $\tilde{m} : \tilde{R}_n \rightarrow \tilde{G}$ that forms a commutative diagram with m and the two covering maps. For any marked tree T and any finitely generated subgroup A of F_n , there is a subtree T_A of T that represents A .

Once a particular marked graph G is chosen, any homotopy equivalence $f : G \rightarrow G$ induces an outer automorphism ϕ of the free group. We will say that f *represents* ϕ . Choosing a basepoint $*$ in G and a path from $*$ to $f(*)$ induces an automorphism Φ in the class ϕ . For every $\phi \in \text{Out}(F_n)$, one can find an f representing it that maps vertices to vertices and restricts to an immersion on any edge [BH92]. Such a homotopy equivalence is called a *topological representative*. If v is a vertex of G and $f(v)$ is a vertex w , then f maps $T_v(G)$ to $T_w(G)$.

Let \tilde{G} be the universal cover of G . A *path in \tilde{G}* is a proper map $\tilde{\sigma} : I \rightarrow \tilde{G}$ with domain a (possibly infinite) closed interval I such that either: $\tilde{\sigma}$ is an embedding, or I is finite and the image of $\tilde{\sigma}$ is a single point. In the latter case we say that $\tilde{\sigma}$ is a *trivial path*. If I is finite and $\tilde{\sigma} : I \rightarrow \tilde{G}$ is any continuous map, then $\tilde{\sigma}$ is homotopic rel endpoints to a unique (possibly trivial) path $[\tilde{\sigma}]$. We say that $[\tilde{\sigma}]$ *is obtained from $\tilde{\sigma}$ by tightening*. A map $\tilde{\sigma}$ that satisfies $[\tilde{\sigma}] = \tilde{\sigma}$ is called *tight*. If $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ is a lift of a homotopy equivalence $f : G \rightarrow G$, we denote $[\tilde{f}(\tilde{\sigma})]$ by $\tilde{f}_\#(\tilde{\sigma})$. When talking of paths, we will often suppress the map from an interval and think of $\tilde{\sigma}$ as the image of this map.

Any path with finite domain can be expressed as a concatenation $\tilde{E}_1\tilde{E}_2\ldots\tilde{E}_j$, where $\tilde{E}_2, \dots, \tilde{E}_{j-1}$ are edges of \tilde{G} , and \tilde{E}_1 and \tilde{E}_j may be edges or segments of edges. We call $\tilde{E}_1\tilde{E}_2\ldots\tilde{E}_j$ the *edge path associated to $\tilde{\sigma}$* . We can extend this notion to paths with infinite domains by allowing for infinite concatenations of edges.

A *path in G* is the image of a path in \tilde{G} under the covering map from \tilde{G} to G .

Equivalently, a path in G is a map σ from a closed interval to G which is an immersion (or possibly has image a point if the interval is finite.) The notions of tightening and edge paths associated to paths can likewise be transferred from \tilde{G} to G by considering images under the covering map. A *circuit in G* is an immersion $\sigma : S^1 \rightarrow G$. For any finite graph G , the *core* of G is the subgraph of G consisting of edges that are crossed by some circuit in G .

If $f : G \rightarrow G$ is a topological representative, then a path σ in G is called a *periodic Nielsen path for f* if $f_{\#}^k(\sigma) = \sigma$ for some $k \geq 1$. If we can take $k = 1$ (meaning that $f_{\#}(\sigma) = \sigma$), then we call σ a *Nielsen path*. A (periodic) Nielsen path is *indivisible* if it does not decompose as a concatenation of non-trivial (periodic) Nielsen subpaths.

The set of periodic points of f is denoted $Per(f) \subset G$; the set of fixed points of f is denoted $Fix(f)$. Two points x and y in $Fix(f)$ are called *Nielsen equivalent* if they are the endpoints of the Nielsen path for f , and the resulting equivalence classes are called *Nielsen classes*. These fixed paths and classes will be exceedingly important later on, as they carry information about the fixed subgroups of the outer automorphism that f represents.

When considering such fixed paths, we may also ask about the fixed paths of a lift $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$. In the universal cover, however, these fixed paths may be infinite in length, limiting toward points in the boundary of the tree \tilde{G} , denoted $\partial\tilde{G}$. As is true of automorphisms representing ϕ , any lift \tilde{f} of a topological representative induces a homeomorphism $\partial\tilde{f} : \partial\tilde{G} \rightarrow \partial\tilde{G}$ [F80]. We define $Fix(\partial\tilde{f})$ to be the subset of $\partial\tilde{G}$ fixed by this homeomorphism, and

define $\text{Fix}_N(\partial\tilde{f})$ to be the set of non-repelling fixed points of \tilde{f} .

Given a choice of basepoint, the boundary of \tilde{G} is naturally identified with ∂F_n , where a based, tight infinite ray in \tilde{G} is identified with the reduced infinite word in ∂F_n that it carries under the marking map from \tilde{R}_n . There is, in addition, an identification between representatives of ϕ and lifts of a topological representative f of ϕ such that, if Φ corresponds to \tilde{f} , then $\partial\Phi$ and $\partial\tilde{f}$ are the same homeomorphism of $\partial\tilde{G} = \partial F_n$ with the same fixed sets [BFH04, Lemma 2.1][F80]. Under this identification between automorphisms and maps of \tilde{G} , an inner automorphism i_c is identified with a covering translation T_c , and $\text{Fix}_N(\partial i_c) = \text{Fix}_N(\partial T_c)$ consists of two points T_c^+ and T_c^- , and the bi-infinite edge path that connects these two points, denoted \tilde{A}_c , is called the *axis of c in \tilde{G}* .

3.5 Lines in the Free Group

The axis of an element of the free group is one example of a line in the free group, meaning an object represented by a bi-infinite path in either G or \tilde{G} . More formally, we define the *space of directed lines in F_n* , to be $\tilde{\mathcal{B}} = \tilde{\mathcal{B}}(F_n) = \partial F_n \times \partial F_n$, where an element (P, Q) is *the line from P to Q* , and (P, P) is the *trivial line at $P \in \partial F_n$* . (All other pairings are referred to as *nontrivial lines*.) The point P is called the *initial endpoint* of the directed line (P, Q) , and Q is its *terminal endpoint*. For any two directed lines of the form (P, Q) and (Q, R) , their product is defined as follows: $(P, Q) || (Q, R) = (P, R)$. This operation endows $\tilde{\mathcal{B}}$ with the structure of a groupoid. Under the operation, (Q, P) is the

inverse of (P, Q) . With the exception of Lemma 3.3, the definitions and facts given in this section can be found in [BFH00, Section 2.2], though there the lines utilized are unoriented.

For any marked graph G , define $\tilde{\mathcal{B}}(\tilde{G})$ to be the set of bi-infinite directed paths in $\tilde{G} \cup \partial\tilde{G}$, called the *space of directed lines in \tilde{G}* . There is an identification between $\tilde{\mathcal{B}}(F_n)$ and $\tilde{\mathcal{B}}(\tilde{G})$ induced by the identification between ∂F_n and $\partial\tilde{G}$, and so $\tilde{\mathcal{B}}(\tilde{G})$ carries the same groupoid structure. On $\tilde{\mathcal{B}}(\tilde{G})$, this structure takes the following form: for any two directed lines λ and λ' in $\tilde{\mathcal{B}}$ such that the terminal end of λ is the initial end of λ' , their product $\lambda||\lambda'$ is obtained by concatenating the two lines via their common end and reducing, thereby yielding the unique directed line which runs from the initial end of λ to the terminal end of λ' . As before, there is a *trivial line* at every point of $\partial\tilde{G}$, and the inverse of a directed line is the same line with the opposite direction.

There is a natural action of F_n on $\tilde{\mathcal{B}}(F_n) = \tilde{\mathcal{B}}(\tilde{G})$ induced by its action on its own boundary. Namely, for any $x \in F_n$ and (P, Q) in $\tilde{\mathcal{B}}$, $x(P, Q) = (xP, xQ)$. If we take the quotient by this action, then we obtain the sets $\mathcal{B}(F_n)$ and $\mathcal{B}(G)$, with an induced identification between them. We call any element of $\tilde{\mathcal{B}}(\tilde{G})$ the *realization in \tilde{G}* of the corresponding element of $\tilde{\mathcal{B}}(F_n)$; likewise, any element of $\mathcal{B}(G)$ is called the *realization in G* of the corresponding element of $\mathcal{B}(F_n)$.

Under the quotient from $\tilde{\mathcal{B}}(\tilde{G})$ to $\mathcal{B}(G)$, every nontrivial line in $\tilde{\mathcal{B}}(\tilde{G})$ has image a bi-infinite directed path in G . By the marking on G , every such path yields an element of $\mathcal{B}(F_n)$ represented by a bi-infinite string of generators and their inverses. Any automorphism acts on such a word, and the action of an

inner automorphism is trivial. And so, we have a natural action of $Out(F_n)$ on $\mathcal{B}(F_n)$ and (by extension) on $\mathcal{B}(G)$. Equivalently, this action of $Out(F_n)$ on \mathcal{B} can be induced by the action of $Aut(F_n)$ on $\tilde{\mathcal{B}}$ by $\Phi(P, Q) = (\partial\Phi(P), \partial\Phi(Q))$.

In this notation, we can now view the axis \tilde{A}_c of $c \in F_n$ from Section 3.4 as an element of $\tilde{\mathcal{B}}(\tilde{G}) = \tilde{\mathcal{B}}(F_n)$. We denote by A_c the image of \tilde{A}_c in $\mathcal{B}(G) = \mathcal{B}(F_n)$, and we call it the *axis of the conjugacy class of c* . The axis of a conjugacy class is also sometimes referred to as a *periodic line*, and it is represented by the bi-infinite word $\dots ccc \dots$.

We will now prove a line analog of Lemma 3.2. The argument will be similar.

Lemma 3.3. *Let $\tilde{\ell} = (P, Q) \in \tilde{\mathcal{B}}(F_n)$ be a lift of a non-periodic directed line $\ell \in \mathcal{B}(F_n)$, and let $\phi \in Out(F_n)$ so that $\phi(\ell) = \ell$. Then there is a unique representative Φ of ϕ so that $\Phi(\tilde{\ell}) = \tilde{\ell}$.*

Proof. If Φ is any representative of ϕ , then Φ must send (P, Q) to another ordered pair of boundary points in the same equivalence class ℓ . In other words, $\Phi(P, Q) = (xP, xQ)$ for some $x \in F_n$. The line (xP, xQ) cannot be equal to (P, Q) unless it is a lift of the axis of the conjugacy class of x . Because ℓ is non-periodic, it cannot be this axis, and so $i_x^{-1}\Phi$ is the only representative of ϕ whose induced map of $\tilde{\mathcal{B}}(F_n)$ will fix $\tilde{\ell}$. \square

The space $\tilde{\mathcal{B}}(\tilde{G})$ can be topologized in the following way: for any finite path $\tilde{\alpha}$ in \tilde{G} , define the *neighborhood of $\tilde{\alpha}$* , $N(\tilde{\alpha})$, to be the subset of $\tilde{\mathcal{B}}(\tilde{G})$ consisting of all lines that contain $\tilde{\alpha}$. The collection of neighborhoods $N(\tilde{\alpha})$, with $\tilde{\alpha}$ ranging

over all finite paths in \tilde{G} , forms a basis for a topology. We use the natural quotient map from $\tilde{\mathcal{B}}(\tilde{G})$ to $\mathcal{B}(G)$ to define the quotient topology on $\mathcal{B}(G)$, and we extend the topologies to $\tilde{\mathcal{B}}(F_n)$ and $\mathcal{B}(F_n)$ via their identifications with $\tilde{\mathcal{B}}(\tilde{G})$ and $\mathcal{B}(G)$. The induced actions of $Aut(F_n)$ on $\tilde{\mathcal{B}}(F_n)$ and $Out(F_n)$ on $\mathcal{B}(F_n)$ are by homeomorphisms. Any closed subset Λ of $\mathcal{B}(F_n) = \mathcal{B}(G)$ is called a *lamination*. This name is also applied to any closed F_n -invariant subset of $\tilde{\mathcal{B}}(\tilde{F}_n) = \tilde{\mathcal{B}}(\tilde{G})$.

For any subset Y of ∂F_n , define *the set of lines carried by Y* , $\tilde{\mathcal{B}}(Y)$, to be the subset of $\tilde{\mathcal{B}}(F_n)$ consisting of all directed lines both of whose endpoints are in Y . This subset is equal to the product $Y \times Y$. We say a subgroup A of F_n is *root closed* if, whenever $a \in A$ and $b = a^j$ for some $j \in \mathbb{Z}$, b is in A as well. For any root closed subgroup A of F_n , define *the set of lines carried by A* , $\tilde{\mathcal{B}}(A)$, to be $\tilde{\mathcal{B}}(\partial A)$, where $\partial A \subset \partial F_n$ is the boundary of A .

Finally, For $[A]$ the conjugacy class of a root closed subgroup of F_n , define *the set of lifted lines carried by $[A]$* , $\tilde{\mathcal{B}}([A])$, to be the subset of $\tilde{\mathcal{B}}(F_n)$ given by taking the union of the sets $\tilde{\mathcal{B}}(A')$ as A' ranges over representatives of $[A]$, and let $\mathcal{B}([A])$ (called *the set of lines carried by $[A]$*) be their image under the quotient from $\tilde{\mathcal{B}}$ to \mathcal{B} .

3.6 Principal Representatives

It is clear from the definition of $Out(F_n)$ that every outer automorphism has infinitely many choices of representative automorphism, just as any topological representative has infinitely many possible lifts. We will restrict to a special set

of representatives that will be distinguished by the size of their fixed sets. We say that Φ representing ϕ is a *principal representative* if $\text{Fix}_N(\partial\Phi)$ contains at least three points, or if it contains two points that are not the endpoints of the axis of some element of F_n .¹ The set of all principal representatives of ϕ is denoted $P(\phi)$. Any lift of a topological representative that corresponds to an element of $P(\phi)$ is called a *principal lift*.

Intuitively, in choosing principal representatives, we select the automorphisms with the largest fixed sets, so as to rule out those that do not carry much information about the outer automorphism. To see why this is necessary, consider the following two automorphisms of F_2 .

$$\Phi_1 : \quad a \mapsto a \quad b \mapsto b \quad \quad \Phi_2 : \quad a \mapsto baB \quad b \mapsto babAB$$

The automorphism Φ_2 is the conjugation i_{ba} , and so both Φ_1 and Φ_2 represent the identity in $\text{Out}(F_2)$. However, while Φ_1 fixes every element of F_2 , Φ_2 fixes only the subgroup generated by ba . Which of these is a more accurate description of the identity outer automorphism: that it fixes a small subgroup, or that it fixes everything? The map Φ_1 carries more information about the identity, and so we say Φ_1 is principal, while Φ_2 is not.

Given two representatives Φ and Φ' of an outer automorphism ϕ , we say they are *isogredient* if there is some $x \in F_n$ so that $\Phi = i_x \Phi' i_x^{-1}$. Isogredience

¹In [FH11], the definition of *principal* is slightly different; it also rules out possibility that the two points are the endpoints of a certain line associated to an EG outer automorphism. As we will be restricting our attention to PG maps for the main results of this paper, I have left out this case for simplicity.

defines an equivalence relation on representatives of ϕ . In the language of marked graphs, if $f : G \rightarrow G$ is a topological representative of ϕ , we say that lifts \tilde{f} and \tilde{f}' are *isogredient* if $\tilde{f} = T_x \tilde{f}' T_x^{-1}$ for some x . Isogredient lifts correspond to isogredient automorphisms, and the equivalence classes of either automorphisms or lifts are called *isogredience classes*. It is easy to see that, for any Φ and Φ' in the same isogredience class, $Fix(\Phi') = i_x Fix(\Phi)$ and $Fix_N(\partial\Phi') = \partial i_x Fix_N(\partial\Phi)$, where i_x is the inner automorphism that conjugates Φ' to Φ , and likewise for isogredient lifts of f . This also shows that the elements of an isogredience class are either all principal or all non-principal.

When f is a particular kind of topological representative of ϕ called a relative train track map (defined in Section 3.10), the isogredience classes of principal lifts of f (and by extension principal representatives of ϕ) are in one-to-one correspondence with the Nielsen classes of f . And, as there can only be finitely many such Nielsen classes, it follows that there are only finitely many isogredience classes of principal representatives of any $\phi \in Out(F_n)$. This fact is shown in Lemma 3.8, Remark 3.9, and Corollary 3.17 of [FH11], and it implies that, for any $\phi \in Out(F_n)$, the set of all Fix_N 's of its principal representatives is finite when considered up to conjugacy.

3.7 Linear Growth

At this point, it might be helpful to discuss how some of the definitions given above can be simplified in the case that ϕ is of at most linear growth (LG). For example, LG maps have no isolated fixed points [CL99, Corollary 7.7], and so,

by the remark at the end of Section 3.3, if ϕ is LG, then for any Φ representing ϕ , $Fix_N(\partial\Phi) = \partial Fix(\Phi)$, where $\partial Fix(\Phi)$ is the boundary of $Fix(\Phi)$ as a subset of ∂F_n . This means that, in defining principal representatives of LG outer automorphisms, we can ignore the case that $|Fix_N(\partial\Phi)| = 2$ and those two points are not the endpoints of an axis. Every point in $\partial Fix(\Phi)$ is the endpoint of some axis, and so, if $|\partial Fix(\Phi)| = 2$, then those two points are necessarily the endpoints of the axis of the element generating this rank one subgroup. Therefore, in the case that ϕ is LG, we may define a principal representative of ϕ to be any Φ so that the rank of $Fix(\Phi)$ is at least two.

We are nearing the point where we can state Feighn and Handel's recognition theorem for $Out(F_n)$. First, we need to describe one more aspect of linear growth: twisting. We say that a root-free conjugacy class α in F_n is an *axis* for $\phi \in Out(F_n)$ if there is a $\Phi \in P(\phi)$ and an element a in α such that $a \in Fix(\Phi)$ and $i_a^j \Phi \in P(\phi)$ for some $j \in \mathbb{Z}$. The number j is referred to as the *twist coefficient* of the ordered pair $(\Phi, i_a^j \Phi)$ with respect to a . Intuitively, a is the element of F_n by which Φ performs a “twist” that gives it linear growth. This can be seen clearly by returning to Example 3.1.

Example 3.4.

$$\Phi : a \mapsto a \quad b \mapsto ba \qquad i_a \Phi : a \mapsto a \quad b \mapsto ab$$

Both Φ and $i_a \Phi$ are principal for the outer automorphism ϕ that they represent. The conjugacy class of the element a is an axis for ϕ , and we can see that the element b grows linearly under either map by “twisting” in a . If instead we

looked at the outer automorphism ϕ^2 represented by Φ^2 , then $i_a\Phi^2$ is not principal for ϕ^2 ; instead, $i_a^2\Phi^2$ is; the twist coefficient with respect to a is two, as can be seen in Example 3.5.

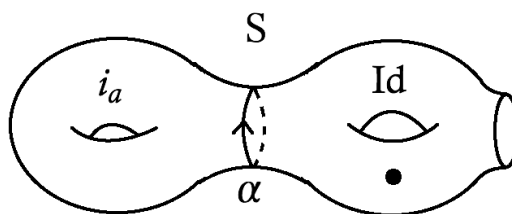
Example 3.5.

$$\Phi^2 : a \mapsto a \quad b \mapsto ba^2 \quad i_a\Phi^2 : a \mapsto a \quad b \mapsto aba \quad i_a^2\Phi^2 : a \mapsto a \quad b \mapsto a^2b$$

The use of the word “twisting” here has its origins in the world of the mapping class group of a surface S , where the homeomorphisms that induce linear growth on word lengths in the fundamental group are roots of *Dehn multitwists* on S . This analogy can be further visualized and understood through the following surface example. Note that, because the surface depicted has nontrivial boundary, the homeomorphisms described induce elements of $\text{Out}(F_n)$.

Example 3.6.

Figure 2: A Linear Growth Mapping Class



Let D_α be the homeomorphism of S given by performing a Dehn Twist in the curve α . (This means to cut an annular neighborhood of α out of S and glue it back with a 360 degree twist at one end.) If we choose as basepoint the point marked in the right half of S in Figure 2, D_α determines an automorphism Δ_α

of $\pi_1(S) = F_4$ that is the identity on the subgroup of F_4 represented by the right half of S and is a conjugation by a on the subgroup of F_4 represented by left half of S (where a is the element of F_4 represented by the curve based at the chosen basepoint in the homotopy class of α .) This automorphism is principal for the outer automorphism it represents, as it fixes the rank two fundamental group of the right side of S . A representative given by placing the basepoint on the left side of S is also principal, and the automorphism of $\pi_1(S)$ it induces is equal to $i_a^{-1}\Delta_\alpha$. The curve α that defines the twisting (or equivalently the conjugacy class in F_n that it represents) is an axis for the element δ_α of $Out(F_4)$ represented by both Δ_α and $i_a^{-1}\Delta_\alpha$.

A fuller exploration of LG outer automorphisms from the perspective of graphs of groups is available in [CL95], [CL99], and [KLV01]. In those papers, every axis of ϕ is an edge stabilizer in a graph of groups decomposition of F_n corresponding to ϕ . As we'll see shortly, these axes go on a short list of very important invariants that uniquely determine a PG outer automorphism of the free group.

3.8 The Forward Rotationless Recognition Theorem

The recognition theorem of Feighn and Handel applies only to *forward rotationless* outer automorphisms, those with no periodic behavior. More precisely, we say $\phi \in Out(F_n)$ is forward rotationless if, for all $k > 0$ and for all $\Phi \in P(\phi)$, the map that sends Φ to Φ^k defines a bijection b_k between $P(\phi)$ and $P(\phi^k)$ such that $Fix_N(\partial b_k(\Phi)) = Fix_N(\partial \Phi)$. This requirement on Fix_N shows that all periodic elements of Φ are in fact fixed. We are now able to

state the forward rotationless version of recognition theorem in the PG case.

Theorem. [FH11] Let $\phi, \psi \in \text{Out}(F_n)$ be forward rotationless and PG. If there exists a bijection $B : P(\phi) \rightarrow P(\psi)$ such that:

$$(a) \text{ Fix}_N(\partial\Phi) = \text{Fix}_N(\partial B(\Phi))$$

$$(b) \text{ If } w \in \text{Fix}(\Phi) \text{ and } \Phi, i_w\Phi \in P(\phi), \text{ then } B(i_w\Phi) = i_w B(\Phi).$$

Then $\phi = \psi$.

Note that the colloquial version of this theorem might be, “If ϕ and ψ have the same fixed sets, axes, and twist coefficients, then they are equal,” as those are the rough statements contained in hypotheses (a) and (b), given in terms of principal automorphisms representing ϕ and ψ .

Instead of working with individual representatives of ϕ and ψ , one might try to speak in more global terms. As was explained in Section 3.6, for any $\Phi_1, \Phi_2 \in P(\phi)$ that are isogredient, $\text{Fix}(\Phi_1)$ and $\text{Fix}(\Phi_2)$ are conjugate subgroups of F_n . And so, if we consider set of the conjugacy classes of the fixed subgroups of principal representatives of ϕ , this is a finite set. We call the disjoint union of this set $\text{Fix}(\phi)$. However, as the recognition theorem shows, conjugacy classes of fixed subgroups are not enough. When higher-than-linear polynomial growth occurs, we must include attracting fixed points on the boundary.

Let X be a subset of ∂F_n which is invariant under the action of some subgroup H of F_n ; X is called an H -subset. The pairing (H, X) is considered *conjugate* to another such pairing (H', X') if there is an element $a \in F_n$ so that $H' = H^a$

and $X' = aX$. The conjugacy class of (H, X) is defined to be $\bigcup_{a \in F_n} (H^a, aX)$ and is denoted $[H, X]$. The following lemma allows us to think of the fixed set of an automorphism of the free group in this notation.

Lemma 3.7. *For any $\Phi \in \text{Aut}(F_n)$, $\text{Fix}_N(\partial\Phi)$ and $\text{Fix}_+(\partial\Phi)$ are both $\text{Fix}(\Phi)$ -subsets of ∂F_n .*

Proof. Let $P \in \partial F_n$ be an element of $\text{Fix}_N(\partial\Phi)$, and let $w \in F_n$ be in $\text{Fix}(\Phi)$. Then $\partial\Phi(wP) = \Phi(w)\partial\Phi(P) = wP$, meaning that wP is in $\text{Fix}(\partial\Phi)$.

Assume that P is an attractor for Φ . By the definition in Section 3.3, this means there is a neighborhood V of P such that $\partial\Phi(V) \subset V$ and $\bigcap_{k=1}^{\infty} \partial\Phi^k(V) = \{P\}$. And so

$$\partial\Phi(wV) = \Phi(w)\partial\Phi(V) = w\partial\Phi(V) \subset wV,$$

and

$$\bigcap_{k=1}^{\infty} \partial\Phi^k(wV) = \bigcap_{k=1}^{\infty} \Phi^k(w)\partial\Phi^k(V) = \bigcap_{k=1}^{\infty} w\partial\Phi^k(V) = w \bigcap_{k=1}^{\infty} \partial\Phi^k(V) = \{wP\}.$$

In other words, wP is an attractor for Φ with attracting neighborhood wV . If we instead assume that P is a repeller for Φ , then the same argument above with Φ^{-1} in place of Φ shows that wP is a repeller for Φ . This shows the invariance under $\text{Fix}(\Phi)$ of both $\text{Fix}_N(\partial\Phi)$ and $\text{Fix}_+(\partial\Phi)$. \square

With this in mind, for any $\Phi \in \text{Aut}(F_n)$, consider the conjugacy class $[\text{Fix}(\Phi), \text{Fix}_N(\partial\Phi)]$ of the pair $(\text{Fix}(\Phi), \text{Fix}_N(\Phi))$, and for any forward rotationless $\phi \in \text{Out}(F_n)$, define $\text{Fix}_N(\phi)$ to be the disjoint union of the pairings

$[Fix(\Phi), Fix_N(\Phi)]$ as Φ ranges over all principal representatives of ϕ . Because we are only considering pairings up to conjugacy, this is a finite union. Similarly, define the set of *eigenrays for Φ* , $\tilde{\mathcal{R}}(\Phi)$, to be the conjugacy class $[Fix(\Phi), Fix_+(\partial\Phi)]$ of the pair $(Fix(\Phi), Fix_+(\Phi))$, and define the set of *eigenrays for ϕ* to be the disjoint union of these sets over $P(\phi)$.

To prove their forward rotationless recognition theorem, Feighn and Handel made use of certain topological representatives of outer automorphisms with added structure. These representatives, called relative train tracks, have played a crucial role in the proofs of many important results in the study of $Out(F_n)$. We will also require them to prove of the main results in this thesis.

3.9 Free Factor Systems and Filtrations

We say that a subgroup $F < F_n$ is a *free factor* if there is another subgroup F' so that F_n is equal to the free product of F and F' . A *subgroup system* is a collection $\mathcal{S} = \{[A_1], \dots, [A_j]\}$ of the conjugacy classes of finitely many nontrivial, finite rank, root-free subgroups of F_n . If there are subgroups A'_i representing $[A_i]$ so that every A'_i and their free product $A'_1 * \dots * A'_j$ are free factors, then we call \mathcal{S} a *free factor system*. We will often denote a free factor system by \mathcal{F} . A good reference for free factor systems is [BFH00, Section 2.6], and many of the definitions we quote below can be found in that section.

For $[F]$ the conjugacy class of a free factor F , the conjugacy class $[x]$ of $x \in F_n$ is *carried by $[F]$* if F contains a representative of $[x]$. We will sometimes abuse notation and say that x is carried by F . We say that an abstract directed line

$\ell \in \mathcal{B}(F_n)$ is carried by $[F]$ if it is a limit of axes of conjugacy classes carried by $[F]$.

If G is a marked graph and G_F is a connected subgraph of G so that $[\pi_1(G_F)] = [F]$, then a conjugacy class $[x]$ is carried by $[F]$ if and only if the circuit representing $[x]$ in G is contained in G_F . We can equivalently say ℓ is carried by $[F]$ if its realization in G is contained in such a G_F . When such a G_F is chosen for every free factor in \mathcal{F} in the same graph G , we call their disjoint union $G_{\mathcal{F}}$. We say $G_{\mathcal{F}}$ is a realization of \mathcal{F} in G . Finally, if $\mathcal{F} = \{[F^1], \dots, [F^j]\}$ is a free factor system and W is a collection of abstract directed lines and conjugacy classes of elements, then we say W is carried by \mathcal{F} if every element of W is carried by some $[F^i]$.

Following the notation of Section 3.5, for \mathcal{S} a subgroup system in F_n , define the set of lifted lines carried by \mathcal{S} , $\tilde{\mathcal{B}}(\mathcal{S})$, to be the disjoint union of the sets $\tilde{\mathcal{B}}([A])$ as $[A]$ ranges over the conjugacy classes of subgroups in \mathcal{S} , and define the set of lines carried by \mathcal{S} , $\mathcal{B}(\mathcal{S})$, to be the disjoint union of the $\mathcal{B}([A])$'s. Note that the subgroup system in question may be a free factor system.

For free factors F^1 and F^2 , we say $[F^1] \sqsubset [F^2]$ if F^1 is conjugate to a free factor of F^2 . And for free factor systems \mathcal{F}_1 and \mathcal{F}_2 , we say $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ if for every $[F^i] \in \mathcal{F}_1$ there exists $[F^j] \in \mathcal{F}_2$ so that $[F^i] \sqsubset [F^j]$. This defines a partial order on the set of all free factor systems in F_n . The *complexity* of a free factor system $\mathcal{F} = \{[F^1], \dots, [F^j]\}$, is defined to be the non-increasing sequence of numbers given by $cx(\mathcal{F}) = rank(F^1), \dots, rank(F^j)$ after possibly reordering the F^i 's. By Grushko's Theorem [G40], the sum of the numbers in

the complexity of any free factor system is less than or equal to n , the rank of the ambient free group, and so there are only finitely many such possible complexities for a fixed n . We can order these complexities lexicographically, meaning that, for example, $7, 6 > 3, 3, 3, 2 > 2 > 1, 1$. Each number in the complexity of a free factor system is called a *digit*.

It is clear from the definitions that $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ implies $cx(\mathcal{F}_1) \leq cx(\mathcal{F}_2)$. We call \mathcal{F}_2 a *one edge extension* of \mathcal{F}_1 if one of the following two situations occurs:

- \mathcal{F}_1 and \mathcal{F}_2 contain the same number of free factors, one digit of $cx(\mathcal{F}_2)$ is one more than the corresponding digit of $cx(\mathcal{F}_1)$, and all other digits are equal.
- \mathcal{F}_2 has one fewer free factor than \mathcal{F}_1 , one digit of $cx(\mathcal{F}_2)$ is equal to the sum of two digits of $cx(\mathcal{F}_1)$, and all other digits of $cx(\mathcal{F}_2)$ are the same as the other digits of $cx(\mathcal{F}_1)$.

Equivalently, \mathcal{F}_2 is a one edge extension of \mathcal{F}_1 if G_1 and G_2 are core subgraphs of a marked graph G realizing \mathcal{F}_1 and \mathcal{F}_2 and if \mathcal{F}_2 is equal to the union of \mathcal{F}_1 with a single edge or the concatenation of edges in which every vertex of $G_2 - G_1$ is valence two in G_2 . We also sometimes say that G_2 is a one edge extension of G_1 .

Lemma 3.8. *If \mathcal{S} is a subgroup system and Λ is a lamination that contains every periodic line in $\mathcal{B}(\mathcal{S})$, then Λ contains $\mathcal{B}(\mathcal{S})$.*

Proof. Let $\ell \in \mathcal{B}(\mathcal{S})$. By definition, $\mathcal{B}(\mathcal{S})$ is equal to the disjoint union of the sets $\mathcal{B}([A])$ for $[A]$ the conjugacy class of a subgroup in \mathcal{S} , and so ℓ is in $\mathcal{B}([A])$

for some $[A] \in \mathcal{S}$. Let A' be a subgroup representing $[A]$, and choose a basis for A' denoted by the letters a, b, c and so on up to the rank of A' . Then ℓ can be represented by a reduced bi-infinite word $\dots a_{-3}a_{-2}a_{-1}a_0a_1a_2a_3\dots$, where each a_i is a basis element or the inverse of a basis element. Consider the sequence of bi-infinite words $\ell_j = \dots (a_{-j}\dots a_0\dots a_j)(a_{-j}\dots a_0\dots a_j)(a_{-j}\dots a_0\dots a_j)\dots$. If each of these bi-infinite words is reduced (meaning that a_{-j} is never equal to a_j^{-1}), then ℓ_j represents a sequence of periodic elements of $\mathcal{B}([A]) \subset \mathcal{B}(\mathcal{S})$ whose limit is ℓ , and as Λ is closed by definition, it must therefore contain ℓ , and we are done.

If there is at least one j so that $a_{-j} = a_j^{-1}$, but if there exists a subsequence ℓ_{j_i} of ℓ_j so that a_{-j_i} never equals a_{j_i} , then again we are done because ℓ is the limit of this subsequence of periodic lines. If there does not exist such a subsequence, then there is some $J > 0$ such that for all $j \geq J$, $a_{-j} = a_j^{-1}$. In this case, we define ℓ'_j to be the sequence of bi-infinite words $\dots a'_ja'_ja'_j(a_{-j}\dots a_0\dots a_j)a'_ja'_ja'_j\dots$, where $a'_j = b$ if $a_j = a$ or a^{-1} , and $a'_j = a$ otherwise. The lines ℓ'_j are reduced by construction; they therefore represent elements of $\mathcal{B}(\mathcal{S})$, and their limit as j goes to infinity is ℓ .

To finish the proof, we need only check that each ℓ'_j can be approximated by a periodic lines in $\mathcal{B}([A])$. The sequence of periodic lines represented by the reduced bi-infinite words

$$\ell_{j,k} = \dots ((a'_j)^k a_{-j}\dots a_0\dots a_j (a'_j)^k) ((a'_j)^k a_{-j}\dots a_0\dots a_j (a'_j)^k) ((a'_j)^k a_{-j}\dots a_0\dots a_j (a'_j)^k) \dots$$

gives this approximation. Their limit as k goes to infinity is ℓ'_j . \square

Lemma 3.9. *If \mathcal{S} is a subgroup system, then $\mathcal{B}(\mathcal{S})$ is a lamination.*

Proof. Let $[A]$ be one of the conjugacy classes of nontrivial, finite rank, root free subgroups in \mathcal{S} . Let ℓ be a line in the closure $\overline{\mathcal{B}([A])}$ of $\mathcal{B}([A])$, the space of lines carried by $[A]$. Then there is a sequence of lines $\ell_k \in \mathcal{B}([A])$ so that the limit as i goes to infinity of ℓ_i is ℓ . Let G be a marked graph, and think of ℓ and ℓ_k as elements of $\mathcal{B}(G)$.

Let $\tilde{\ell}$ be a lift of ℓ to \tilde{G} . Then we may choose lifts $\tilde{\ell}_k$ of the ℓ_k that converge to $\tilde{\ell}$ and that are contained in the subtree T_A of \tilde{G} determined by some subgroup A' in $[A]$. Assume for the purposes of this proof that \tilde{G} has been given a metric where each edge is assigned length one. Then by the definition of convergence in $\tilde{\mathcal{B}}(\tilde{G})$, for every M there exists a K so that for all $k \geq K$, $\tilde{\ell}_k$ and $\tilde{\ell}$ share a common segment of length at least M . Therefore, arbitrarily long segments of $\tilde{\ell}$ are contained in T_A , and so $\tilde{\ell}$ is contained in T_A as well. This means that ℓ is contained in $\mathcal{B}([A])$, and so $\mathcal{B}([A])$ is closed as a subset of $\mathcal{B}(G) = \mathcal{B}(F_n)$.

The set $\mathcal{B}(\mathcal{S})$ is the finite union of the sets $\mathcal{B}([A])$ for $[A] \in \mathcal{S}$ and so is closed itself. Therefore, by definition, $\mathcal{B}(\mathcal{S})$ is a lamination. \square

Suppose that G is a marked graph and that $f : G \rightarrow G$ is a homotopy equivalence representing $\phi \in \text{Out}(F_n)$. A *filtration* of G is an increasing sequence $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ of subgraphs, each of whose components contains at least one edge. If $f(G_i) \subset G_i$ for all i then we say that $f : G \rightarrow G$

respects the filtration or that the *filtration is f -invariant*. A path or circuit has *height r* if it is contained in G_r but not G_{r-1} . A lamination has height r if each directed line in its realization in G has height at most r and some directed line has height r .

The r^{th} *stratum* H_r is defined to be the closure of $G_r - G_{r-1}$. To each stratum H_r there is an associated square matrix M_r , called the *transition matrix for H_r* , whose ij^{th} entry is the number of times that the i^{th} edge (in some ordering of the edges of H_r) crosses the j^{th} edge in either direction. By enlarging the filtration, we may assume that each M_r is either irreducible or the zero matrix. A filtration with this property is called *maximal*. We say that H_r is an *irreducible stratum* if M_r is irreducible and is a *zero stratum* if M_r is the zero matrix.

There is a *transition matrix* associated to any topological representative of a graph $f : G \rightarrow G$ defined similarly to the transition matrix of a stratum in a filtration of G . Namely, put some ordering on all of the edges of G , and set the ij^{th} entry m_{ij} of the matrix M_f to be the number of times that the i^{th} edge of G crosses the j^{th} edge in either direction. This matrix induces a unique *maximal filtration of G associated to f* by reordering the edges of G so that an edge e precedes an edge e' whenever there is a sequence of edges $e = e_0, e_1, \dots, e_K = e'$ such that $m_{i,i+1}$ is non-zero. This filtration is f -invariant, and full details of its construction can be found in [BH92, Section 5].

In the case that $\phi \in Out(F_n)$ represented by f is PG, by subdividing edges or splitting filtration elements in two, we may assume that each stratum H_r in this

maximal filtration is a cycle of edges E_1, \dots, E_q satisfying that $f(E_i) = E_{i+1}u_i$ for all i , where $u_i \subset G_{r-1}$ and where indices are taken mod q [FH11, Section 2.6].

Recall that every connected subgraph C of a marked graph G represents the conjugacy class of the free factor that is C 's fundamental group, and that a disconnected subgraph of G represents the free factor system composed of the fundamental groups of each of its connected components. In particular, if $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ is a filtration of G , then each G_r represents a free factor system $\mathcal{F}(G_r)$, and if $f : G \rightarrow G$ is a topological representative of ϕ that respects the filtration, then ϕ leaves each $\mathcal{F}(G_r)$ invariant. We say the nested sequence of free factor systems $\mathcal{C} = \mathcal{F}(G_0) \sqsubset \mathcal{F}(G_1) \sqsubset \dots \sqsubset \mathcal{F}(G_N)$ is *realized by* the filtration of G .

3.10 Relative Train Track Maps

In [BH92], Bestvina and Handel defined special topological representatives called *train track maps* for *irreducible* outer automorphisms, a special class of maps that are either finite order or EG. They also developed *relative train track maps* to handle outer automorphisms that leave invariant certain free factors, possibly mixing some EG and lower growth behavior. However, as we are currently only working with PG outer automorphisms, the definition of a relative train track simplifies considerably, and this is the definition I'll give here. For $\phi \in \text{Out}(F_n)$ of at most polynomial growth, a *relative train track map* is a topological representative $f : G \rightarrow G$ of a marked graph where

the induced maximal filtration corresponding to f satisfies that the transition matrix of every stratum is an irreducible matrix. Every definition that follows in this section will assume that $\phi \in \text{Out}(F_n)$ is PG, and many of them are significantly more complicated in their most general form.

In [FH11], Feighn and Handel prove the following theorem that gives relative train tracks even more structure.

Theorem 3.10. *[Theorem 2.19 of [FH11]] For every $\phi \in \text{Out}(F_n)$ there is a relative train track map $f : G \rightarrow G$ and filtration that represents ϕ and satisfies the following properties.*

- (V) *The endpoints of all indivisible periodic Nielsen paths are vertices.*
- (P) *If a stratum $H_m \subset \text{Per}(f)$ is a forest then there exists a filtration element G_j such $\mathcal{F}(G_j) \neq \mathcal{F}(G_l \cup H_m)$ for any G_l .*
- (NEG) *The terminal endpoint of an edge in a non-periodic stratum H_i is periodic and is contained in a filtration element of height less than i that is its own core.*
- (F) *The core of each filtration element is a filtration element.*
- (FFS) *If \mathcal{C} is a nested sequence of non-trivial ϕ -invariant free factor systems then we may choose $f : G \rightarrow G$ to realize \mathcal{C} .*

When working with forward rotationless outer automorphisms, one can modify these homotopy equivalences further still. We say that $x \in \text{Per}(f)$ is *principal* if x is not contained in a component C of $\text{Per}(f)$ that is topologically a circle

where each point in C has exactly two periodic directions. Lifts to \tilde{G} of principal periodic points in G are also said to be *principal*. If each principal vertex and each periodic direction at a principal vertex has period one then we say that $f : G \rightarrow G$ is *rotationless*.

Recall from the end of Section 3.9 that we may assume every stratum H_r takes the form $\{E_1, \dots, E_q\}$, where $f(E_i) = E_{i+1}u_i$ for some $u_i \subset G_{r-1}$. The initial endpoint of each E_i is a principal periodic point for f , and so, if f is rotationless, then we may assume each stratum H_r is a single edge E with $f(E) = Eu$ for $u \subset G_{r-1}$ [FH11, Section 3.3].

The following statement from [FH11] gives the correspondence between forward rotationless outer automorphisms and rotationless relative train track maps.

Proposition 3.11. *[Proposition 3.29 of [FH11]] Suppose that $f : G \rightarrow G$ represents ϕ and satisfies the conclusions of Theorem 3.10. Then $f : G \rightarrow G$ is rotationless if and only ϕ is forward rotationless.*

But we can modify f even further to our advantage. We need only a few more definitions to describe a special kind of train track known as a CT.

If E_i and E_j are linear edges and if there are $m_i, m_j > 0$ and a closed root-free Nielsen path w such that $f(E_i) = E_i w^{m_i}$ and $f(E_j) = E_j w^{m_j}$, then a path of the form $E_i w^p \bar{E}_j$ with $p \in \mathbb{Z}$ is called an *exceptional path*.

A filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ is said to be *reduced* (with respect to ϕ) if, whenever a free factor system \mathcal{F}' is ϕ^k -invariant for some $k > 0$ and

$\mathcal{F}(G_{r-1}) \sqsubset \mathcal{F}' \sqsubset \mathcal{F}(G_r)$, it follows that either $\mathcal{F}' = \mathcal{F}(G_{r-1})$ or $\mathcal{F}' = \mathcal{F}(G_r)$. Roughly, this tells us that the filtration encodes every ϕ^k -invariant free factor system for all k .

A non-trivial path or circuit σ in G is *completely split* if it has a splitting, called a *complete splitting*, into subpaths, each of which is a single edge in an irreducible stratum, an indivisible Nielsen path, or an exceptional path. A relative train track map is *completely split* if $f(E)$ is completely split for each edge E in each irreducible stratum.

Definition 3.12. A relative train track map $f : G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ is said to be a *CT* (for completely split improved relative train track map) if it satisfies the following properties.

1. **(Rotationless)** $f : G \rightarrow G$ is rotationless.
2. **(Completely Split)** $f : G \rightarrow G$ is completely split.
3. **(Filtration)** The filtration is reduced. The core of each filtration element is a filtration element.
4. **(Vertices)** The endpoints of all indivisible periodic (necessarily fixed) Nielsen paths are (necessarily principal) vertices. The terminal endpoint of each non-fixed edge is principal (and hence fixed).
5. **(Periodic Edges)** Each periodic edge is fixed and each endpoint of a fixed edge is principal. If the unique edge E_r in a fixed stratum H_r is not a loop then G_{r-1} is a core graph and both ends of E_r are contained in G_{r-1} .

6. **(Zero Strata)** There are no zero strata.
7. **(Linear Edges)** For each linear E_i there is a closed root-free Nielsen path w_i such that $f(E_i) = E_i w_i^{d_i}$ for some $d_i \neq 0$. If E_i and E_j are distinct linear edges with the same axes then $w_i = w_j$ and $d_i \neq d_j$.
8. **(Indivisible Nielsen Paths)** For every indivisible Nielsen path σ , there is a linear edge E_i with w_i as in (Linear Edges) and there exists $k \neq 0$ such that $\sigma = E_i w_i^k \bar{E}_i$.

For an edge E_i as in (Linear Edges), we call the conjugacy class $[w_i]$ represented by the Nielsen path w_i the *axis of E_i* .

Theorem 3.13. *[Theorem 4.29 of [FH11]] Suppose that $\phi \in \text{Out}(F_n)$ is forward rotationless and that $\mathcal{F} = \{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_t\}$ is a nested sequence of ϕ -invariant free factor systems. Then ϕ is represented by a CT $f : G \rightarrow G$ with a filtration that realizes \mathcal{F} .*

Some important properties follow directly from the definition of a CT, for example the following result of Feighn and Handel.

Lemma 3.14. *[Lemma 4.37 of [FH11]] Assume that $f : G \rightarrow G$ is a CT. The following properties hold for every principal lift $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$.*

1. *If $\tilde{v} \in \text{Fix}(\tilde{f})$ and a non-fixed edge \tilde{E} determines a fixed direction at \tilde{v} , then $\tilde{E} \subset \tilde{f}_\#(\tilde{E}) \subset \tilde{f}_\#^2(\tilde{E}) \subset \dots$ is an increasing sequence of paths whose union is a ray \tilde{R} that converges to some $P \in \text{Fix}_N(\partial \tilde{f})$ and whose interior is fixed point free.*

2. For every isolated $P \in \text{Fix}_N(\partial\tilde{f})$ there exists \tilde{E} and \tilde{R} as in (1) that converges to P . The edge E is non-linear.

For $f : G \rightarrow G$ a CT representing ϕ , set $K_0 = K_0(f)$ to be the subgraph of G consisting of all fixed and linear growth edges. We say an edge E of G is a *growing edge* if it is not contained in K_0 . Let \mathcal{E}_f be the set of all growing edges of G . Then Lemma 3.14 sets up a correspondence between growing edges for f and eigenrays for ϕ .

We can define a partial order on \mathcal{E}_f by $E < F$ if E is a term of the complete splitting of $f^N(E)$ for some $N \geq 0$. This in turn induces a partial order on $\mathcal{R}(\phi)$. Extend these partial orders to total orders so that $\mathcal{R}(\phi) = \{\rho_1, \dots, \rho_t\}$ and $\mathcal{E}_f = \{E_1, \dots, E_t\}$, where each set is listed in increasing order and ρ_i corresponds to E_i .

Set $K_j = K_j(f) = K_0(f) \cup \bigcup_{i=1}^j E_i$. Then the (possibly disconnected) subgraphs K_j determine free factor systems $\mathcal{F}_j = \mathcal{F}(K_j)$. It is clear from the definition that $\mathcal{F}_j \sqsubset \mathcal{F}_{j'}$ if $j \leq j'$. We call $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \dots \sqsubset \mathcal{F}_t = \mathcal{F}_n$ the *sequence of free factor systems determined by ϕ, G , and the total order on \mathcal{E}_f* , and sometimes we write the sequence as $\mathcal{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_t\}$. The choice of how to obtain a total order from the partial order on \mathcal{E}_f is a finite one. If we take the union of the sets \mathcal{F} over all possible choices of finite order, we get the partially ordered *hierarchy of free factor systems associated to ϕ and G* .

Note 1. If we let $\mathcal{F}_0(\phi)$ equal the smallest free factor system that carries every conjugacy class that grows at most linearly under ϕ , then for any CT $f : G \rightarrow G$ representing ϕ , $\mathcal{F}(K_0(f)) = \mathcal{F}_0(\phi)$, as by definition every edge of

K_0 grows at most linearly under f (and so does every circuit in K_0), and, in the other direction, every conjugacy class carried by $\mathcal{F}_0(\phi)$ can be represented by a circuit in $K_0(f)$.

3.11 Stallings Graphs

In [S83], Stallings gave an algorithm that, for any subgroup H of F_n , provides a core graph $\Gamma(H)$ carrying a marking by H . This graph is called the *Stallings graph* of H . The Stallings graph represents the conjugacy class of H , and choosing a basepoint gives a representative subgroup in that conjugacy class. Given any automorphism Φ of F_n , the Stallings graph of $Fix(\Phi)$ could be a useful topological stand-in for the conjugacy class of this fixed subgroup, an algebraic invariant. However, $\Gamma(Fix(\Phi))$ does not keep track of the isolated fixed points of $\partial\Phi$ in ∂F_n , points that were essential in the forward rotationless recognition theorem. And so, we will need a more general notion of Stallings graphs.

Definition 3.15. Following the notation of Section 3.8, (H, X) is a pairing of a subgroup H of F_n and a subset X of ∂F_n that is H -invariant, and $[H, X]$ is the conjugacy class of (H, X) . Let G be any marked graph and define the *Stallings graph of $[H, X]$ with respect to G* to be the graph $\Gamma_G([H, X])$ obtained as follows: let \bar{G} be \tilde{G} compactified by adding its set of ends. The graph $\Gamma_G([H, X])$ is equal to the quotient by the action of H of the convex hull of X in \bar{G} .

As the convex hull of X is the union of X with a subtree of \tilde{G} , it follows that

$\Gamma_G([H, X])$ is the union of a subgraph of the cover \hat{G} of G corresponding to H with a set of boundary points of \hat{G} . In particular, the core of the Stallings graph $\text{core}(\Gamma_G([H, X]))$ has fundamental group H , and we can think of $\Gamma_G([H, X])$ as the core of \hat{G} union any infinite embedded rays that limit toward isolated points of $X \bmod H$. These isolated points make up what we call the *boundary of $\Gamma_G([H, X])$* .

Let Φ in $\text{Aut}(F_n)$ be a representative of a forward rotationless ϕ , and let $f : G \rightarrow G$ be any CT representing ϕ . Then we define the *Stallings graph of Φ with respect to G* to be $\Gamma_G(\Phi) = \Gamma_G([\text{Fix}(\Phi), \text{Fix}_N(\partial\Phi)])$.

Lemma 3.16. *If Φ_1 and Φ_2 are isogredient representatives of a forward rotationless map ϕ , then $\Gamma_G(\Phi_1)$ and $\Gamma_G(\Phi_2)$ are isomorphic as marked graphs.*

Proof. Suppose $\Phi_2 = i_x \Phi_1 i_x^{-1}$ for $x \in F_n$. Then $\text{Fix}(\Phi_2) = i_x(\text{Fix}(\Phi_1))$ and $\text{Fix}_N(\partial\Phi_2) = \partial i_x(\text{Fix}_N(\partial\Phi_1))$. Therefore $[\text{Fix}(\Phi_2), \text{Fix}_N(\partial\Phi_2)] = [\text{Fix}(\Phi_1), \text{Fix}_N(\partial\Phi_1)]$, and so $\Gamma_G(\Phi_1) = \Gamma_G(\Phi_2)$ □

As is explained in Section 3.6, any outer automorphism ϕ has only finitely many isogredience classes of principal representatives, and so there are only finitely many such Stallings graphs associated to $P(\phi)$.

Definition 3.17. The *Stallings graph of ϕ with respect to G* , $\Gamma_G(\phi)$ is the disjoint union of $\Gamma_G(\Phi)$ as Φ ranges over $P(\phi)$. The *boundary of the Stallings graph*, $\partial\Gamma_G(\phi)$, is the disjoint union of the boundaries of the $\Gamma_G(\Phi)$.

The following lemmas describe what the Stallings graph of a forward rotation-less PG outer automorphism looks like. We know from the definition that each component \mathcal{C} of $\Gamma_G(\phi)$ corresponds to the fixed sets of one isogredience class of principal representatives of ϕ (or equivalently, to Nielsen classes of fixed points of f). We will now discuss the different pieces that can occur in one of those components by examining the roles played by the growing and non-growing edges of G .

Lemma 3.18. *1. Every point of the boundary of $\Gamma_G(\phi)$ corresponds to an eigenray for ϕ .*

2. For every growing edge $E \in \mathcal{E}_f$, the component \mathcal{C}_E of $\Gamma_G(\phi)$ corresponding to the Nielsen class $\iota(E)$ contains lift \hat{E} that is the initial endpoint of a ray \hat{R}_E toward the boundary of \mathcal{C}_E in which every vertex but $\iota(\hat{E})$ is valence two in $\Gamma_G(\phi)$. There are lifts \tilde{E} of E and \tilde{R}_E of \hat{R}_E and a principal lift \tilde{f} of f to \tilde{G} so that \tilde{R}_E is obtained as the union of the increasing sequence of paths $\tilde{E} \subset \tilde{f}_\#(\tilde{E}) \subset \tilde{f}_\#^2(\tilde{E}) \subset \dots$

3. For every point \hat{P} in the boundary of $\Gamma_G(\phi)$, there is a growing edge E as in (2) so that \hat{P} is the endpoint of the corresponding ray \hat{R}_E .

Proof. To prove (1), recall that any component C of $\Gamma_G(\phi)$ corresponds to an isogredience class of a principal representative Φ of ϕ by definition. The boundary points of C are the images under the quotient by the action of $Fix(\Phi)$ of points in $Fix_N(\partial\Phi)$ that are not in $\partial Fix(\Phi)$. Because $Fix_N(\partial\Phi)$ contains only elements of $\partial Fix(\Phi)$ and isolated attractors, it follows that the bound-

any points of C must be the images of isolated attractors, which naturally correspond to eigenrays for ϕ .

Statements (2) and (3) are just the translations to Stallings graph language of the statements in Lemma 3.14. \square

Lemma 3.19. *Every loop in $\text{core}(\Gamma_G(\phi))$ is a lift of a closed Nielsen path for p . Moreover, if γ is a circuit in G composed of fixed edges, then the component \mathcal{C}_γ of $\Gamma_G(\phi)$ corresponding to the Nielsen class of γ contains a circuit $\hat{\gamma}$ that is a lift of γ .*

Proof. By definition, the core of each component of a Stallings graph has fundamental group the conjugacy class $[Fix(\Phi)]$ for some principal representative Φ of ϕ . Every conjugacy class $[x]$ carried by $[Fix(\Phi)]$ is fixed by ϕ , and so is realized in G by a closed Nielsen path of f .

If γ is composed of fixed edges, then it represents a conjugacy class carried by the subgroup system $Fix(\phi)$, and that conjugacy class is represented by a circuit in $\Gamma_G(\phi)$. \square

Lemma 3.20. *Every edge in $\text{core}(\Gamma_G(\phi))$ is a lift of an edge that grows at most linearly under p .*

Proof. Let E be an edge of G that does not grow at most linearly under f . In other words, $E \in \mathcal{E}_f$. Item (Indivisible Nielsen Paths) of the definition of a CT states that any indivisible Nielsen path of f takes the form $\sigma = E_i w_i^k \bar{E}_i$,

where E_i is a linear edge, and $[w_i]$ is the axis of E_i . Thus, E cannot be an edge in an indivisible Nielsen path. And as E is not fixed pointwise, it cannot be an edge in a divisible Nielsen path. Therefore, by Lemma 3.19, no lift of E can be contained $\text{core}(\Gamma_G(\phi))$. The lemma is the contrapositive of this statement. \square

Recall item (Linear Edges) of the definition of a CT: for every linear edge E_i there is a closed root-free Nielsen path w_i so that $f(E_i) = E_i w_i^{d_i}$ for some $d_i \neq 0$. Note that, because of this property, the initial vertex of E_i is always fixed by f .

Lemma 3.21. *For E_i a linear edge of G , the component C_i corresponding to the Nielsen class of $\iota(E_i)$ contains a lift \hat{E}_i of E_i with a closed lift \hat{w}_i of w_i at its terminal endpoint.*

Proof. If $f(E_i) = E_i w_i^{d_i}$, then the edge path $E_i w_i \bar{E}_i$ is an indivisible Nielsen path for f , and any automorphism Φ in the isogredience class corresponding to the Nielsen class of $\iota(E)$ will fix the element of F_n represented by $E_i w_i \bar{E}_i$ up to conjugacy. Therefore, there must be a path $\hat{E}_i \hat{w}_i \hat{E}_i^{-1}$ in C_i based at a lift $\iota(\hat{E}_i)$ of $\iota(E_i)$, exactly as described. \square

The union of \hat{E}_i and \hat{w}_i is referred to as a *lollipop*.

Because of the above lemmas, given any CT G , we can construct the Stallings graph with respect to G by gluing together closed lifts of circuits composed of fixed edges, lollipops coming from linear edges, and infinite rays corresponding to growing edges. We will see this in action in the examples of Section 4.

4 Guesses and Counterexamples

In trying to find a recognition theorem for general PG outer automorphisms, the following fact is a good starting point: for any n , there is a universal K such that, for every element ϕ of $Out(F_n)$, ϕ^K is forward rotationless [FH11, Lemma 4.43]. Because of this, if we are given any two non-forward rotationless outer automorphisms ϕ and ψ , we may first ask whether ϕ^K and ψ^K are equal via the forward rotationless recognition theorem. If ϕ^K and ψ^K are different, then ϕ and ψ must be as well. And so, if we assume $\phi^k = \psi^k = \pi$ is forward rotationless, then we can use invariants of π to try to recognize ϕ and ψ .

As a first attempt, one might hope that, because ϕ and ψ have a common power, if they agree on their shared periodic subgroups and boundary points (namely, the fixed subgroups and boundary points of π), then ϕ and ψ will be the same outer automorphism. This is the intuition that inspires the following guess at a recognition theorem.

Guess 1. Let $\phi, \psi \in Out(F_n)$ be such that $\phi^k = \psi^k = \pi$ is forward rotationless. If $\phi|_{Fix_N(\pi)} = \psi|_{Fix_N(\pi)}$, then $\phi = \psi$.

However, things aren't quite so simple, as the following counterexample to this statement shows.

Counterexample 1. Let ϕ and ψ be elements of $Out(F_4)$ represented by Φ and Ψ as follows:

$$\begin{array}{ll} \Phi : & \begin{array}{ll} a \mapsto c & c \mapsto a \\ b \mapsto dc & d \mapsto ba \end{array} \end{array} \qquad \begin{array}{ll} \Psi : & \begin{array}{ll} a \mapsto c & c \mapsto a \\ b \mapsto d & d \mapsto ba^2 \end{array} \end{array}$$

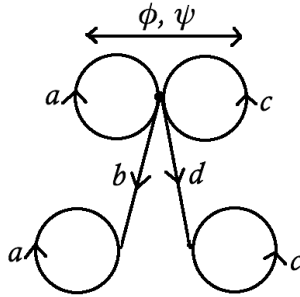
Then $\phi^2 = \psi^2 = \pi$ can be represented by the rotationless map Π as follows:

$$\Pi : \begin{array}{ll} a \mapsto a & c \mapsto c \\ b \mapsto ba^2 & d \mapsto dc^2 \end{array}$$

In this case, because all the automorphisms are linear growth, $Fix_N(\pi) = Fix(\pi)$, which is the conjugacy class of the fixed subgroup $\langle a, baB, c, dcD \rangle$ of Π . And indeed, on this set, Φ and Ψ agree. (Both switch a with c and baB with dcD .) However, it is clear that ϕ is not equal to ψ : for example, ψ sends the conjugacy class $[b]$ to the conjugacy class $[d]$, while $\phi([b]) = [dc]$, and dc is not conjugate to d by any element of the free group.

If we think of Π as a CT acting on the rose R_4 with the edges marked by a, b, c , and d , the Stallings graph $\Gamma_G(\pi)$ has a single component in the form of a two-petaled rose marked by a and c , and two lollipops representing baB and dcD . Both ϕ and ψ act by switching the two petals of the rose and the two corresponding lollipops.

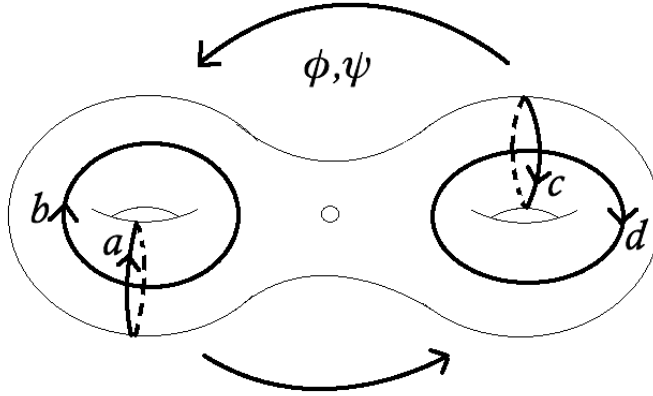
Figure 3: The Stallings graph $\Gamma_{R_4}(\pi)$ of Counterexample 1



The reason this counterexample is possible is that $Fix_N(\pi)$ does not keep

track of the twisting numbers associated to the edges b and d , and so we can choose to distribute the twisting in two different ways over these edges. This twisting issue becomes even more apparent if we realize all three automorphisms as homeomorphisms of a surface S of genus two with one puncture. This example and the associated surface picture are inspired by an example in [KLV01] that was used to illustrate conjugacy between two roots of an LG outer automorphism.

Figure 4: A homeomorphism of a surface inducing Counterexample 1



The automorphisms Φ and Ψ of $\pi_1(S)$ are induced by homeomorphisms that switch the two sides of S and perform the appropriate number of Dehn Twists in the simple closed curves on S representing $[a]$ and $[c]$. Meanwhile, Π is induced by the double Dehn twist in both curves.

However, this twisting issue is a linear phenomenon, and one might hope that, at least in the PG case, if you force ϕ and ψ to agree on their LG conjugacy classes, then any twisting will have to agree. This inspires our next guess.

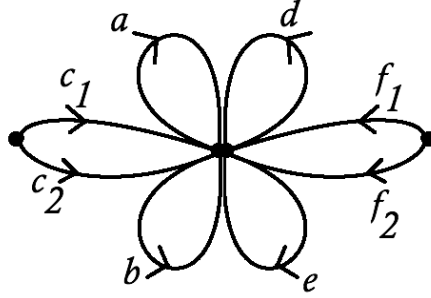
Guess 2. Let $\phi, \psi \in \text{Out}(F_n)$ be PG, with $\phi^k = \psi^k = \pi$ forward rotationless.

If $\phi|_{\text{Fix}_N(\pi)} = \phi|_{\text{Fix}_N(\pi)}$ and $\phi|_{\mathcal{F}_0(\pi)} = \phi|_{\mathcal{F}_0(\pi)}$, then $\phi = \psi$.

Unfortunately, this guess is also false. What follows is a counterexample to this statement.

Counterexample 2. Let G be the rose R_6 with a looped marked by each of the generators of $F_6 = \langle a, b, c, d, e, f \rangle$, but with the edges marked c and f each subdivided into two edges by a vertex, so that $c = \bar{c}_1 c_2$ and $f = \bar{f}_1 f_2$ as depicted.

Figure 5: A graph G for Counterexample 2



Let h and s as follows be homotopy equivalences representing ϕ and ψ respectively.

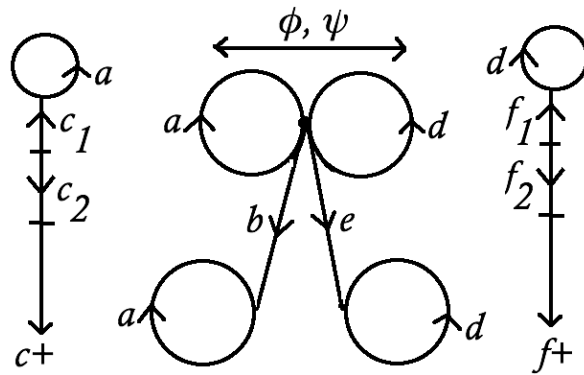
	$a \mapsto d$	$d \mapsto a$	
	$b \mapsto ed$	$e \mapsto ba$	
$h :$	$c_1 \mapsto f_1 d^2$	$f_1 \mapsto c_1 a^2$	
	$c_2 \mapsto f_2 e$	$f_2 \mapsto c_2 b$	
			$s :$
			$a \mapsto d$
			$d \mapsto a$
			$b \mapsto ed$
			$e \mapsto ba$
			$c_1 \mapsto f_1$
			$f_1 \mapsto c_1 a^4$
			$c_2 \mapsto f_2 e$
			$f_2 \mapsto c_2 b$

Then $h^2 = s^2 = p$ is a CT representing ϕ 's and ψ 's common power as follows:

$$\begin{array}{lcl}
& a \mapsto a & d \mapsto d \\
& b \mapsto ba^2 & e \mapsto ed^2 \\
p : & c_1 \mapsto c_1 a^4 & f_1 \mapsto f_1 d^4 \\
& c_2 \mapsto c_2 bba & f_2 \mapsto f_2 eed
\end{array}$$

This CT has three distinct Nielsen classes of fixed points. The Nielsen class of the central vertex includes the loops a and c , and the corresponding component of $\Gamma_G(\pi)$ is similar to the Stallings graph of Counterexample 1, with loops representing a and d and lollipops representing baB and edE . The component of $\Gamma_G(\pi)$ corresponding to the Nielsen class of the common initial vertex of c_1 and c_2 has a lollipop representing the loop $c_1 a C_1$ because of the linear growth of c_1 and an infinite ray because of the quadratic growth of c_2 . The same is true for the Nielsen class of the common initial vertex of f_1 and f_2 .

Figure 6: The stallings graph $\Gamma_G(\pi)$ of Counterexample 2



Here, $\mathcal{F}_0(\pi)$ is the conjugacy class of the free factor $\langle a, b, d, e \rangle$, and ϕ and

ψ agree on this free factor. They also agree on $Fix_N(\pi)$, both switching the two isolated points c^+ and f^+ . And yet, once again, we can see that they are not the same outer automorphism. The key here is that the edges c_1 and f_1 exhibit linear growth, and so we can distribute twisting along edges a and d over c_1 and f_1 , just as in Counterexample 1 we distributed twisting by a and c over the edges b and d . However, c_1 and f_1 are not closed paths, and they are not contained in any closed paths in G that represent linearly growing conjugacy classes. This is why the difference between ϕ and ψ is not picked up by $\mathcal{F}_0(\pi)$.

However, this example also gives a hint at how we might find a recognition theorem that works. Consider the directed line ℓ in $\mathcal{B}(G)$ represented by $\dots aaaB\bar{c}_1c_2bbaba^2ba^3\dots$. You can visualize this line in the Stallings graph above as a path that begins by winding around the lower lift of a in the center component, runs up the edge marked b , jumps to the left component, and runs down the infinite ray toward c^+ . Note that $h(\ell) = \dots dddED\bar{f}_1f_2eeded^2ed^3\dots$, while $s(\ell) = \dots dddE\bar{f}_1f_2eeded^2ed^3\dots$. The fact that these are not the same image line in $\mathcal{B}(G)$ gives us a way to differentiate ϕ from ψ . This is the idea that we will describe in more generality in the next section.

5 Annular Fix_N -Lines

Lemma 5.1. *Let $\pi, \chi \in Out(F_n)$ so that $[\pi, \chi] = 1$. Then $\chi(Fix_N(\pi)) = Fix_N(\pi)$.*

Proof. Because χ commutes with π , χ^{-1} does as well. Therefore, it is enough to

show that $\chi(Fix_N(\pi)) \subset Fix_N(\pi)$, because if this is true, then $\chi^{-1}(Fix_N(\pi)) \subset Fix_N(\pi)$, and applying χ to this containment shows that $Fix_N(\pi) = \chi\chi^{-1}(Fix_N(\pi)) \subset \chi(Fix_N(\pi))$.

If $P(\pi)$ is trivial, then $Fix_N(\pi)$ is trivial, and the statement is vacuously true. Therefore, we may assume that π has principal representatives and call one such representative Π . If X is any representative of χ , then because π and χ commute, Π and X will commute up to conjugacy, meaning that $X\Pi = i_c\Pi X$ for some $c \in F_n$. Meanwhile, X and Π induce homeomorphisms $\partial\Pi$ and ∂X of ∂F_n with the property that $\partial X(Fix_N(\partial\Pi)) = Fix_N(\partial X\Pi X^{-1})$. It follows that $X\Pi X^{-1}$ is principal for the outer automorphism $\chi\pi\chi^{-1}$, because Π is principal for π , and because $Fix_N(\partial X\Pi X^{-1})$ is a homeomorphic image of $Fix_N(\partial\Pi)$. But by assumption, π and χ commute, and so $X\Pi X^{-1}$ is a principal representative of π . Therefore, the conjugacy class of $\partial X(Fix_N(\partial\Pi))$ is contained in $Fix_N(\pi)$. As $Fix_N(\pi)$ is the finite union of the conjugacy classes of these $Fix_N(\partial\Pi)$'s, we have that $\chi(Fix_N(\pi)) \subset Fix_N(\pi)$ as desired.

□

We are now able to define the main object of study for this section, which is also the invariant in use in the main theorem of this thesis.

Definition 5.2. Let $\pi \in Out(F_n)$ be forward rotationless and PG. We wish to associate to π a particular set of lines in F_n which will carry much of π 's dynamic information. To do this, let Π be any principal representative of π , and consider the set $\tilde{\mathcal{B}}(Fix_N(\partial\Pi)) \subset \tilde{\mathcal{B}}(F_n)$ of lines carried by $Fix_N(\partial\Pi) \subset \partial F_n$.

Note that if Π' is any other principal representative of π which is isogredient to Π , then $\tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi'))$ and $\tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi))$ have the same image in \mathcal{B} . We wish to consider all possible principal representatives and their fixed sets, so we define $\tilde{\mathcal{B}}(\text{Fix}_N(\pi))$ to be disjoint union of the sets $\tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi))$ as Π ranges over $P(\pi)$, and define $\mathcal{B}(\text{Fix}_N(\pi))$ to be the image of $\tilde{\mathcal{B}}(\text{Fix}_N(\pi))$ under the quotient to \mathcal{B} .

However, as we saw in Counterexample 2 in Section 4, $\text{Fix}_N(\pi)$ is not a strong enough invariant to uniquely determine π 's roots; we need a slightly larger set of lines. Define $\tilde{\mathcal{L}}_N(\pi)$ to be the subgroupoid of $\tilde{\mathcal{B}}$ generated by the union of the sets $\tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi))$ as Π ranges over $P(\pi)$. Finally, let $\mathcal{L}_N(\pi)$ be the image of $\tilde{\mathcal{L}}_N(\pi)$ under the natural quotient to \mathcal{B} . Each element ℓ of $\mathcal{L}_N(\pi)$ will be called an *annular Fix_N -line for π* . Those directed lines that are contained in $\mathcal{B}(\text{Fix}_N(\pi)) \subset \mathcal{L}_N(\pi)$ are simply called *Fix_N -lines for π* .

Note that $\mathcal{B}(\text{Fix}(\pi))$, the set of lines carried by the subgroup system $\text{Fix}(\pi)$, is contained in $\mathcal{L}_N(\pi)$ by definition. We will also find it helpful to use $\mathcal{B}(\text{Fix}(\pi))$ to define one further set of lines in \mathcal{B} . Following the notation of Section 3.5, for any $\Pi \in P(\pi)$, $\tilde{\mathcal{B}}(\text{Fix}(\Pi)) = \partial\text{Fix}(\Pi) \times \partial\text{Fix}(\Pi)$. Let $\tilde{\mathcal{L}}_0(\pi)$ be the subgroupoid of $\tilde{\mathcal{B}}$ generated by the sets $\tilde{\mathcal{B}}(\text{Fix}(\Pi))$ as Π ranges over $P(\pi)$, and define $\mathcal{L}_0(\pi)$ to be the image of $\tilde{\mathcal{L}}_0(\pi)$ under the quotient to \mathcal{B} . We call any $\ell \in \mathcal{L}_0(\pi)$ an *annular Fix -line for π* . Any line in $\mathcal{B}(\text{Fix}(\pi))$ is a *Fix -line for π* .

While the sets of annular Fix_N - and Fix -lines may seem like unwieldy objects, the concatenations that yield them are actually quite simple in nature, as the

following lemma shows.

Lemma 5.3. *If $(P, Q) \in \tilde{\mathcal{B}}(Fix_N(\partial\Pi_1))$ (respectively $\tilde{\mathcal{B}}(Fix(\Pi_1))$) and $(Q, R) \in \tilde{\mathcal{B}}(Fix_N(\partial\Pi_2))$ (respectively $\tilde{\mathcal{B}}(Fix(\Pi_2))$) for $\Pi_1 \neq \Pi_2 \in P(\pi)$, then $Q \in \partial F_n$ is periodic and represented by an infinite word of the form $xxx\dots$, where $x \in F_n$ is a representative of an axis for π .*

Proof. By assumption, $Q \in \partial F_n$ is fixed by both $\partial\Pi_1$ and $\partial\Pi_2$. Because Π_1 and Π_2 both represent π , $\Pi_2 = i_x^j \Pi_1$ for some root-free $x \in F_n$ and some $j \in \mathbb{Z}$, meaning that $Q = \partial\Pi_2(Q) = \partial i_x^j \partial\Pi_1(Q) = x^j Q$. This cannot happen unless Q is equal to $xxx\dots$ or $x^{-1}x^{-1}x^{-1}\dots$. This means that x is in both $Fix(\Pi_1)$ and $Fix(\Pi_2)$, and an axis is by definition a conjugacy class with a representative that is fixed by two different principal representatives. Therefore, x represents an axis for π . \square

And so, to construct $\tilde{\mathcal{L}}_N(\pi)$ from $\tilde{\mathcal{B}}(Fix_N(\pi))$, we need only consider concatenations that occur at endpoints of lifts of the finitely many axes of π , as these are the only concatenations that result in a reduced line not already carried by $Fix_N(\pi)$.

If $\mathcal{L}_N(\pi)$ is to be a useful object, however, we will need to know that it is invariant under not only π but under any other map which commutes with π (and in particular under π 's roots.) The following lemma checks this invariance.

Lemma 5.4. *Let $\pi \in Out(F_n)$ be forward rotationless and PG, and let $\chi \in Out(F_n)$ so that $[\pi, \chi] = 1$. Then $\chi(\mathcal{L}_N(\pi)) = \mathcal{L}_N(\pi)$.*

Proof. By Lemma 5.1, $\chi(Fix_N(\pi)) = Fix_N(\pi)$. Or, equivalently, in the language of directed lines, we can say that $\chi(\mathcal{B}(Fix_N(\pi))) = \mathcal{B}(Fix_N(\pi))$. Any $\ell \in \mathcal{L}_N(\pi)$ is obtained as the quotient of a concatenation of the form $\tilde{\ell}_1 || \tilde{\ell}_2 || \dots || \tilde{\ell}_j$, where each $\tilde{\ell}_i$ is a lift of a line $\ell_i \in \mathcal{B}(Fix_N(\pi))$, and each adjacent pair of lifts has an endpoint in common. Therefore, because χ acts on \mathcal{B} by a homeomorphism, (or equivalently, because any representative X of χ acts on $\tilde{\mathcal{B}}$ by a homeomorphism) it follows that $\chi(\ell)$ may also be realized as such a concatenation in $\tilde{\mathcal{B}}$, and so is also contained in $\mathcal{L}_N(\pi)$. □

The set of annular Fix_N -lines is by its nature an algebraic object. It requires some topology on the space of abstract lines to define it, but the lines themselves can be thought of as pairs of boundary points or even as bi-infinite strings of letters in the free group. However, in proving our main results, we will find it useful to represent elements of $\mathcal{L}_N(\pi)$ in a more topological manner: namely as bi-infinite paths in a CT for π .

Definition 5.5. Let $p : G \rightarrow G$ be any CT representing π . Consider a component \mathcal{C} of $\Gamma_G(\pi)$ and define $\mathcal{B}(\mathcal{C})$ to be the set of directed lines in $\mathcal{B}(G)$ that lift to \mathcal{C} . Define $\tilde{\mathcal{B}}(\mathcal{C})$ to be the set of all lifts of elements of $\mathcal{B}(\mathcal{C})$ to $\tilde{\mathcal{B}}(\tilde{G})$, let $\tilde{\mathcal{B}}(Fix_N(p))$ be the union of the sets $\tilde{\mathcal{B}}(\mathcal{C})$ as \mathcal{C} ranges over the finitely many components of $\Gamma_G(\pi)$, and set $\tilde{\mathcal{L}}_N(p)$ to be the subgroupoid of $\tilde{\mathcal{B}}(\tilde{G})$ generated by this union. Finally, define the set of *annular Fix_N -lines for p* , $\mathcal{L}_N(p)$, to be the image of $\tilde{\mathcal{L}}_N(p)$ under the quotient back to $\mathcal{B}(G)$.

As we did for $\mathcal{L}_N(\pi)$, we can define a subset of $\mathcal{L}_N(p)$ that ignores isolated fixed rays. To do this, we define $\mathcal{B}_0(\mathcal{C})$ to be the set of directed lines in $\mathcal{B}(G)$ that lift to the core of \mathcal{C} . Take the set of all lifts of these lines to $\tilde{\mathcal{B}}(\tilde{G})$ to be $\tilde{\mathcal{B}}_0(\mathcal{C})$. The union of all the $\tilde{\mathcal{B}}_0(\mathcal{C})$'s is denoted $\tilde{\mathcal{B}}_0(\text{Fix}(p))$, and the subgroupoid generated by this union is $\tilde{\mathcal{L}}_0(p)$. Finally, $\mathcal{L}_0(p)$, the set of *annular Fix-lines* for p is the image of $\tilde{\mathcal{L}}_0(p)$ under the quotient to $\mathcal{B}(G)$.

Lemma 5.6. *Let $\pi \in \text{Out}(F_n)$ be forward rotationless and PG. Then for any choice of CT $p : G \rightarrow G$ representing π , $\mathcal{L}_N(p) = \mathcal{L}_N(\pi)$ and $\mathcal{L}_0(p) = \mathcal{L}_0(\pi)$.*

Proof. By the definitions of the Stallings graphs of π and any representative Π (Definitions 3.15 and 3.17) and by Lemma 3.16, there is a one-to-one correspondence between connected components of $\Gamma_G(\pi)$ and isogredience classes of principal representatives of π . And for any such component \mathcal{C} , the set $\tilde{\mathcal{B}}(\mathcal{C}) \subset \tilde{\mathcal{B}}(\tilde{G})$ is equal to the union of the sets $\tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi))$, as Π ranges over the principal representatives of π in the isogredience class corresponding to \mathcal{C} . Therefore, we have the following equality of unions:

$$\bigcup_{\substack{\mathcal{C} \text{ a component} \\ \text{of } \Gamma_G(\pi)}} \tilde{\mathcal{B}}(\mathcal{C}) = \bigcup_{\Pi \in P(\pi)} \tilde{\mathcal{B}}(\text{Fix}_N(\partial\Pi)) \subset \tilde{\mathcal{B}}(\tilde{G})$$

Thus, the subgroupoids of $\tilde{\mathcal{B}}(\tilde{G})$ generated by these two unions are equal as well. Or, in other words, $\tilde{\mathcal{L}}_N(p) = \tilde{\mathcal{L}}_N(\pi)$, and so $\mathcal{L}_N(p) = \mathcal{L}_N(\pi)$.

Similarly, the core of each \mathcal{C} represents the conjugacy class of the fixed subgroup of an isogredience class of principal representatives of π , and the set

$\tilde{\mathcal{B}}_0(\mathcal{C})$ is equal to the union of the sets $\tilde{\mathcal{B}}(\partial \text{Fix}(\Pi))$ for Π in that isogredience class. Therefore, we once again have equality of the unions of these sets, and that equality carries through to the subgroupoid generated, and to the quotients $\mathcal{L}_0(p)$ and $\mathcal{L}_0(\pi)$.

□

Because of this lemma, we can think of $\mathcal{L}_N(\pi)$ as either an abstract algebraic object or as a more concrete topological one. We will use this to our advantage in later proofs, often suppressing the difference in order to switch between abstract annular Fix_N -lines and those associated to a particular CT representative.

We will now prove a few results that more precisely describe which lines $\mathcal{L}_N(\pi)$ and $\mathcal{L}_0(\pi)$ do and do not include. For example, $\mathcal{L}_0(p)$ does not contain the axis of every linearly growing conjugacy class. However, as the following lemmas show, it contains sequences of lines that approximate those axes. Recall that, for π forward rotationless and PG, $\mathcal{F}_0(\pi)$ is the smallest free factor system that carries every conjugacy class that grows at most linearly under π .

Lemma 5.7. *Let $\pi \in \text{Out}(F_n)$ be forward rotationless and PG. Then $\mathcal{L}_0(\pi) \subset \mathcal{B}(\mathcal{F}_0(\pi))$.*

Proof. Let $p : G \rightarrow G$ be a CT representing π . Then, by Lemma 3.20, if $\hat{\ell}$ is a bi-infinite edge path in $\text{core}(\Gamma_G(\pi))$, then every edge of $\hat{\ell}$ is a lift of an edge of G that is either fixed or linearly growing under p . If we take a lift $\tilde{\ell}$ of $\hat{\ell}$ to \tilde{G} , then this is still true. And if we concatenate finitely many such $\tilde{\ell}$, then

the line in \tilde{G} that we get after reducing still has this property. And so, if we consider the annular *Fix*-line in $\mathcal{B}(G)$ corresponding to this concatenation, it will be composed of edges of at most linear growth, meaning it represents an abstract directed lined that is contained in $\mathcal{B}(\mathcal{F}_0(\pi))$. \square

Lemma 5.8. *For \mathcal{F} a free factor system, the set $\tilde{\mathcal{B}}(\mathcal{F})$ of lifted lines carried by \mathcal{F} is closed under the groupoid operation on $\tilde{\mathcal{B}}$, meaning that is a subgroupoid.*

Proof. Let G be a marked graph so that there is a unique, possibly disjoint core subgraph $G_{\mathcal{F}}$ of G with fundamental group \mathcal{F} . Let ℓ be a directed line in $\mathcal{B}(\mathcal{F})$, let $\tilde{\ell}$ be a lift of ℓ to $\tilde{\mathcal{B}}(\mathcal{F})$, and let ℓ_G and $\tilde{\ell}_G$ be the corresponding tight directed lines in $\mathcal{B}(\tilde{G})$ and $\tilde{\mathcal{B}}(\tilde{G})$. Then ℓ_G is contained in $G_{\mathcal{F}}$, and $\tilde{\ell}_G$ is contained in the preimage $\tilde{G}_{\mathcal{F}}$ of $G_{\mathcal{F}}$ in \tilde{G} .

If ℓ' is any other element of $\mathcal{B}(\mathcal{F})$ with representative ℓ'_G and lifts $\tilde{\ell}'$ and $\tilde{\ell}'_G$ so that $\tau(\tilde{\ell}_G) = \iota(\tilde{\ell}'_G)$, then the product $\tilde{\ell}''_G = \tilde{\ell}_G || \tilde{\ell}'_G$ is still contained in $\tilde{G}_{\mathcal{F}}$. And so the abstract directed line $\tilde{\ell}'' \in \tilde{\mathcal{B}}(F_n)$ represented by $\tilde{\ell}''_G$ is in $\tilde{\mathcal{B}}(\mathcal{F})$, and $\tilde{\mathcal{B}}(\mathcal{F})$ is closed under the groupoid operation. \square

In particular, this means that, for $\pi \in \text{Out}(F_n)$ forward rotationless, the set $\tilde{\mathcal{B}}(\mathcal{F}_0(\pi))$ of lifted linear lines for π is a subgroupoid of $\tilde{\mathcal{B}}$. The following lemma gives another description of $\mathcal{B}(\mathcal{F}_0(\pi))$ in terms of annular *Fix*-lines.

Lemma 5.9. *Let $\pi \in \text{Out}(F_n)$ be forward rotationless and PG. Then the closure $\overline{\mathcal{L}_0(\pi)}$ of $\mathcal{L}_0(\pi)$ in \mathcal{B} is equal to $\mathcal{B}(\mathcal{F}_0(\pi))$.*

Proof. Because $\mathcal{F}_0(\pi)$ is a subgroup system, $\mathcal{B}(\mathcal{F}_0(\pi))$ is closed by Lemma 3.9. By Lemma 5.7, $\mathcal{L}_0(\pi) \subset \mathcal{B}(\mathcal{F}_0(\pi))$, and so the closure of $\mathcal{L}_0(\pi)$ is also contained in $\mathcal{B}(\mathcal{F}_0(\pi))$.

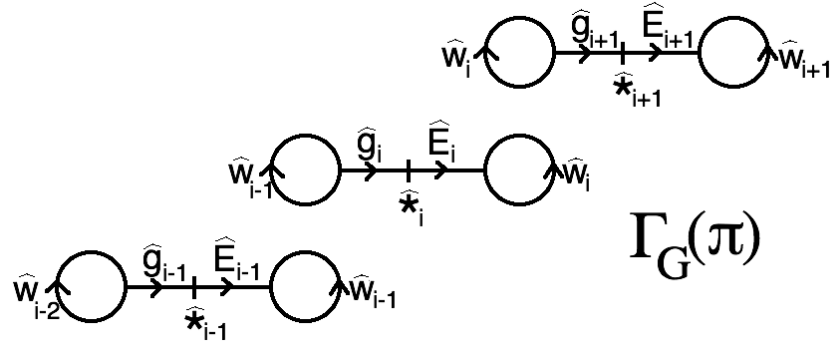
To prove the opposite containment, by Lemma 3.8, $\bar{\ell}$ contains $\mathcal{B}(\mathcal{F}_0)$ if it contains every periodic line in $\mathcal{B}(\mathcal{F}_0)$, and so, it suffices to show that, for any conjugacy class γ that is carried by \mathcal{F}_0 , the axis ℓ_γ of γ is in $\overline{\mathcal{L}_0(\pi)}$. Let γ be such a conjugacy class and let $p : G \rightarrow G$ be a CT representing π chosen so that there is a core subgraph G^0 of G with fundamental group \mathcal{F}_0 .

Then there is a loop g in G that represents γ and is entirely contained in G^0 . Choosing a basepoint for G in g , we may write g as a concatenation $g = g_1 E_1 g_2 E_2 \dots g_m E_m$, where each g_i is a concatenation of fixed edges for p , each E_i is a non-fixed linear edge for p , and some of the g_i or E_i may be trivial. By item (Linear Edges) of the definition of a CT (Definition 3.12), for each i , there is a root-free closed Nielsen path w_i so that either $p(E_i) = E_i w_i^{t_i}$ or $p(E_i) = w_i^{t_i} E_i$. Assume for now that $p(E_i) = E_i w_i^{t_i}$ for all i .

Let $*_i$ be the initial vertex of E_i . Then $*_i$ is fixed by p , and we may consider the component \mathcal{C}_i of $\Gamma_G(\pi)$ corresponding to the Nielsen class of $*_i$. By Lemma 3.21, \mathcal{C}_i must contain a lift $\hat{*}_i$ of $*_i$ with an adjacent lollipop consisting of a lift \hat{E}_i of E_i and a closed lift \hat{w}_i of w_i . In addition, because g_i is a Nielsen path with terminal vertex $*_i$, \mathcal{C}_i must also contain a lift \hat{g}_i of g_i with terminal endpoint $\hat{*}_i$. But the initial vertex of g_i is also the terminal vertex of E_{i-1} (where indices are taken mod m), and so there is also a closed lift \hat{w}_{i-1} of w_{i-1} based at the initial endpoint of \hat{g}_i (which may in fact be $\hat{*}_i$ if g_i is closed or

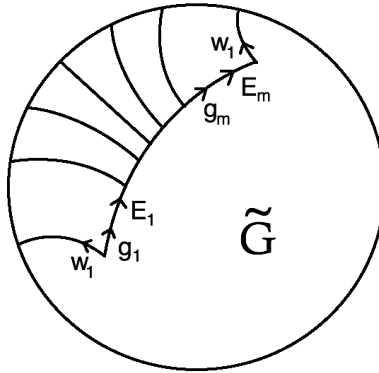
trivial.)

Figure 7: A sequence of components of the Stallings graph of π



This enables us to find a *Fix*-line in $\mathcal{B}(\mathcal{C}_i)$ represented by the bi-infinite path $\dots \hat{w}_{i-1}^{-1} \hat{w}_{i-1}^{-1} \hat{w}_{i-1}^{-1} \hat{g}_i \hat{E}_i \hat{w}_i \hat{w}_i \hat{w}_i \dots$. Similarly, if we consider the fixed point $*_{i+1}$ at the initial endpoint of E_{i+1} and the corresponding component \mathcal{C}_{i+1} of $\Gamma_G(\pi)$, we see a *Fix*-line in $\mathcal{B}(\mathcal{C}_{i+1})$ represented by $\dots \hat{w}_i^{-1} \hat{w}_i^{-1} \hat{w}_i^{-1} \hat{g}_{i+1} \hat{E}_{i+1} \hat{w}_{i+1} \hat{w}_{i+1} \hat{w}_{i+1} \dots$. (See Figure 7.)

Figure 8: A line of the form $\dots \tilde{w}_1^{-1} \tilde{w}_1^{-1} \tilde{w}_1^{-1} \tilde{g} \tilde{w}_1 \tilde{w}_1 \tilde{w}_1 \dots$ in \tilde{G}



There are lifts of these lines to $\tilde{B}(\tilde{G})$ that can be concatenated and reduced via

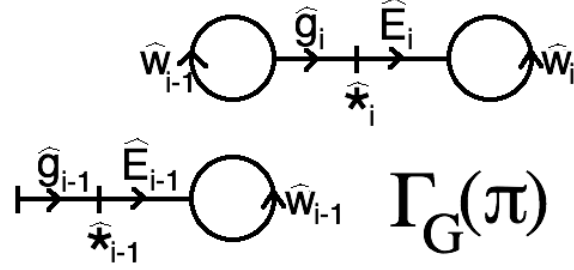
the groupoid operation to yield an annular *Fix*-line for π represented by a line $\dots \tilde{w}_{i-1}^{-1} \tilde{w}_{i-1}^{-1} \tilde{w}_{i-1}^{-1} \tilde{g}_i \tilde{E}_i \tilde{g}_{i+1} \tilde{E}_{i+1} \tilde{w}_{i+1} \tilde{w}_{i+1} \tilde{w}_{i+1} \dots$ in \tilde{G} . Continuing this operation through each C_i , we can concatenate and reduce $m - 1$ times to obtain a lift of an annular *Fix*-line given by $\dots \tilde{w}_m^{-1} \tilde{w}_m^{-1} \tilde{w}_m^{-1} \tilde{g}_1 \tilde{E}_1 \tilde{g}_2 \tilde{E}_2 \dots \tilde{g}_m \tilde{E}_m \tilde{w}_m \tilde{w}_m \tilde{w}_m \dots$. This line has the form $\dots \tilde{w}_m^{-1} \tilde{w}_m^{-1} \tilde{w}_m^{-1} \tilde{g} \tilde{w}_m \tilde{w}_m \tilde{w}_m \dots$, where \tilde{g} is a lift of the loop g . And so, the line $\dots w_m^{-1} w_m^{-1} w_m^{-1} g w_m w_m w_m \dots$ is in $\mathcal{L}_0(p) = \mathcal{L}_0(\pi)$. If we concatenate j appropriate lifts of this same line, then we find that $\dots w_1^{-1} w_m^{-1} w_m^{-1} g^j w_m w_m w_m \dots$ is an annular *Fix*-line for p as well. The limit of this sequence as j goes to infinity is the line $\dots ggg \dots$, a representative of the axis of the conjugacy class γ . Therefore, the axis of γ is in $\overline{\mathcal{L}_0(p)}$.

We must now extend the above argument to the general situation, where for each i , either $p(E_i) = E_i w_i^{t_i}$ or $p(E_i) = w_i^{t_i} E_i$. Note that there are precisely 2^m possibilities. Following the construction above, it suffices to show that in any such arrangement, for each i , we can find components of $\Gamma_G(\pi)$ yielding an $\ell_i \in \mathcal{L}_0(\pi)$ that contains $g_{i-1} E_{i-1} g_i E_i$ as a subpath, as we may then combine those ℓ_i via the groupoid operation to obtain the sequence of annular *Fix*-lines $\dots w_1^{-1} w_m^{-1} w_m^{-1} g^j w_m w_m w_m \dots$ whose limit is the axis of g .

The proof of the existence of ℓ_i requires considering a few cases. In each case, Lemma 3.21 assigns a lollipop in $\Gamma_G(\pi)$ to each E_i , with a lift \hat{w}_i of w_i at either the initial or terminal vertex of a lift \hat{E}_i of E_i . The Stallings graph pictures below are obtained by gluing together these lollipops using the p -fixed g_i . For simplicity, we choose not to depict any pieces of $\Gamma_G(\pi)$ that are not lifts of g_i , E_i , w_i , g_{i-1} , E_{i-1} , or w_{i-1} .

Case A: $p(E_{i-1}) = E_{i-1}w_{i-1}^{t_{i-1}}$ and $p(E_i) = E_iw_i^{t_i}$

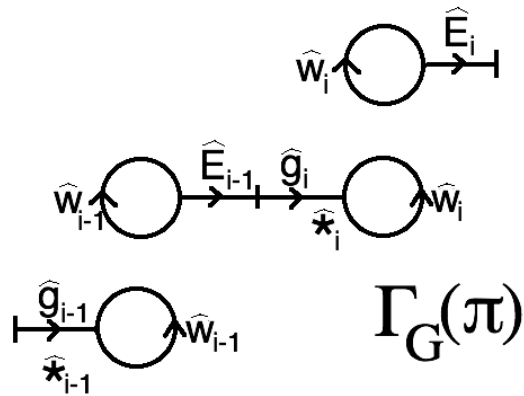
Figure 9: A sequence of components of the Stallings graph of π when $p(E_{i-1}) = E_{i-1}w_{i-1}^{t_{i-1}}$ and $p(E_i) = E_iw_i^{t_i}$



In this case, the Nielsen class of $*_i$ contains g_i and w_{i-1} , while the Nielsen class of $*_{i-1}$ contains g_{i-1} . (See Figure 9.) An annular *Fix*-line ℓ_i with $g_{i-1}E_{i-1}g_iE_i$ as a subpath can be obtained by concatenating and reducing two *Fix*-lines whose common endpoint is an endpoint of an axis for w_{i-1} .

Case B: $p(E_{i-1}) = w_{i-1}^{t_{i-1}}E_{i-1}$ and $p(E_i) = w_i^{t_i}E_i$

Figure 10: A sequence of components of the Stallings graph of π when $p(E_{i-1}) = w_{i-1}^{t_{i-1}}E_{i-1}$ and $p(E_i) = w_i^{t_i}E_i$

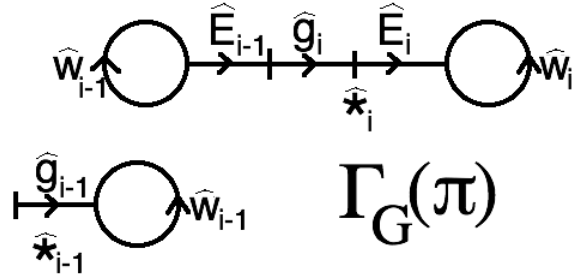


Here, the Nielsen class of $*_i$ contains all of g_i and w_i , while the Nielsen class of

$*_{i-1}$ contains g_{i-1} and w_{i-1} . (See Figure 10.) Two concatenations at endpoints of axes of w_{i-1} and w_i yield the desired ℓ_i .

Case C: $p(E_{i-1}) = w_{i-1}^{t_{i-1}} E_{i-1}$ and $p(E_i) = E_i w_i^{t_i}$

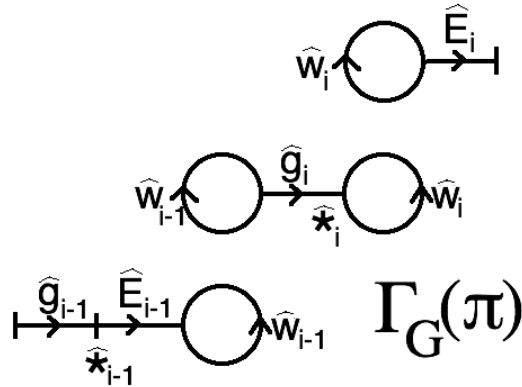
Figure 11: A sequence of components of the Stallings graph of π when $p(E_{i-1}) = w_{i-1}^{t_{i-1}} E_{i-1}$ and $p(E_i) = E_i w_i^{t_i}$



As described in Figure 11, the concatenation of two *Fix*-lines at an endpoint of an axis of w_{i-1} gives the desired annular *Fix*-line ℓ_i .

Case D: $p(E_{i-1}) = E_{i-1} w_{i-1}^{t_{i-1}}$ and $p(E_i) = w_i^{t_i} E_i$

Figure 12: A sequence of components of the Stallings graph of π when $p(E_{i-1}) = E_{i-1} w_{i-1}^{t_{i-1}}$ and $p(E_i) = w_i^{t_i} E_i$



Here again two concatenations are necessary, as the Nielsen class of $*_i$ contains g_i , w_i and w_{i-1} , but $*_{i-1} = \iota(E_{i-1})$ and $\tau(E_i)$ are each in another Nielsen class. (See Figure 12.)

Thus, in any arrangement of the E_i , we can find the sequence of annular *Fix*-lines $\dots w_1^{-1}w_m^{-1}w_m^{-1}g^jw_mw_mw_m\dots$ that shows $\overline{\mathcal{L}_0(\pi)} = \mathcal{B}(\mathcal{F}_0(\pi))$.

□

Now that we know this fact about $\mathcal{L}_0(\pi)$, we can use it to deduce some facts about how roots of π act on $\mathcal{L}_N(\pi)$. If ϕ and ψ are two roots of π , then by Lemma 5.4, they both leave $\mathcal{L}_N(\pi)$ invariant. If ϕ and ψ induce the same action on $\mathcal{L}_N(\pi)$, then it turns out they have the same action on some other invariants of π as well. The next lemma make this precise.

Lemma 5.10. *Let $\phi, \psi \in \text{Out}(F_n)$, with $\phi^k = \psi^k = \pi$ forward rotationless and PG. If $\phi|_{\mathcal{L}_N(\pi)} = \psi|_{\mathcal{L}_N(\pi)}$, then $\phi|_{\mathcal{F}_0(\pi)} = \psi|_{\mathcal{F}_0(\pi)}$ and $\phi|_{\mathcal{R}(\pi)} = \psi|_{\mathcal{R}(\pi)}$.*

Proof. By Lemma 5.9, $\mathcal{B}(\mathcal{F}_0(\pi))$ is contained in the closure of $\mathcal{L}_0(\pi)$. As $\mathcal{L}_0(\pi)$ is contained in $\mathcal{L}_N(\pi)$, it follows that $\mathcal{B}(\mathcal{F}_0(\pi)) \subset \overline{\mathcal{L}_N(\pi)}$ as well. Therefore, by the continuity of the action of ϕ and ψ on \mathcal{B} , we have that ϕ and ψ agree on $\mathcal{B}(\mathcal{F}_0(\pi))$, and so on $\mathcal{F}_0(\pi)$ itself.

The agreement on eigenrays is also next to immediate. Set $\phi^{-1}\psi = \delta$ and suppose there is an element ρ of $\mathcal{R}(\pi)$ such that $\delta(\rho) \neq \rho$. Then if ℓ_ρ is any *Fix*_N-line that lifts to a directed line in $\tilde{\mathcal{B}}$ with terminal endpoint a lift of the endpoint of ρ , then δ cannot fix ℓ_ρ , a contradiction.

□

For ease of reference, the relationships between the various sets of directed lines we've defined are summarized below. As usual, let π be a forward rotationless PG outer automorphism and p a CT representing it. The sets in the middle column are obtained by taking the subgroupoids generated by lifts of the sets on the left. The set on the right is obtained by taking the closure in \mathcal{B} of the set in the middle.

$$\begin{array}{ccc}
 \mathcal{B}(\text{Fix}_N(\pi)) & \subset & \mathcal{L}_N(\pi) = \mathcal{L}_N(p) \\
 \cup & & \cup \\
 \mathcal{B}(\text{Fix}(\pi)) & \subset & \mathcal{L}_0(\pi) = \mathcal{L}_0(p) \subset \mathcal{B}(\mathcal{F}_0(\pi))
 \end{array}$$

Before we can get to the main theorem, we need two more lemmas. One is a result on free factor systems that will allow us to organize an inductive proof of the main theorem, and the other is a corollary from work of Bestvina, Feighn, and Handel [BFH00] that will restrict what can happen in each step of the induction.

Lemma 5.11. *Let $\phi, \psi \in \text{Out}(F_n)$, with $\phi^k = \psi^k = \pi$ forward rotationless and PG, let $p : G \rightarrow G$ be a CT representing π , and let $\mathcal{F} = \{\mathcal{F}_0, \dots, \mathcal{F}_t\}$ be a choice of totally ordered sequence extracted from the hierarchy of free factor systems associated to π and G so that for each i , \mathcal{F}_{i+1} is a one edge extension of \mathcal{F}_i . If $\phi|_{\mathcal{L}_N(\pi)} = \psi|_{\mathcal{L}_N(\pi)}$, then $\delta = \phi^{-1}\psi$ leaves invariant the conjugacy class of every free factor in each of the \mathcal{F}_i 's.*

Proof. By the definition of the hierarchy of free factor systems in Section 3.10, there is a sequence of core subgraphs $G^0 \subset G^1 \subset \dots \subset G^t$ of G such that G^i has a marking by \mathcal{F}_i , and each G^i is a filtration element G_r for some r . By the correspondence between filtration elements and free factor systems outlined in that same section, each G^{i+1} is a one edge extension of G^i , and $G^t = G$.

Set $\delta = \phi^{-1}\psi$. Because $\phi|_{\mathcal{F}_0(\pi)} = \psi|_{\mathcal{F}_0(\pi)}$ by Lemma 5.10, it follows that $\delta([F^i]) = [F^i]$ for each of the conjugacy classes of subgroups $[F^i]$ which make up $\mathcal{F}_0(\pi)$, and moreover, that $\delta|_{[F^i]} = Id_{[F^i]}$. And so any topological representative $d : G \rightarrow G$ of δ is homotopic to a map which leaves G^0 invariant and is the identity when restricted to this subgraph. This will form the base case of an inductive argument.

Assume that δ leaves invariant the conjugacy class of every free factor in each of $\mathcal{F}_0, \dots, \mathcal{F}_J$ for some J between 0 and t . We will show that δ leaves invariant the conjugacy class of every free factor in \mathcal{F}_{J+1} .

Because we are assuming that the first $J+1$ free factor systems are δ -invariant, any topological representative $d : G \rightarrow G$ of δ must leave each of G^0, \dots, G^J invariant up to homotopy. And because \mathcal{F} is a sequence of one edge extensions, it follows that $G^{J+1} - G^J$ is a single edge ϵ with $\iota(\epsilon)$ and $\tau(\epsilon)$ contained in G^J , and with $p(\epsilon) = u\epsilon v$ for circuits u and v in G^J . When u and v are both nontrivial, one subdivides ϵ into two edges E and F so that $\epsilon = \bar{E}F$, $p(E) = E\bar{u}$, and $p(F) = Fv$. In the case that ϵ is subdivided, there is one filtration element G_r in between G^J and G^{J+1} given by the union of G^J with either E or F . However, in this case, note that $\mathcal{F}(G_r) = \mathcal{F}(G^J)$. If ϵ is not

subdivided, then G^J and G^{J+1} are adjacent in the filtration of G .

Let L_ϵ be the set of all annular Fix_N -lines ℓ for p of the form $\bar{R}_1 \in R_2$, where R_1 and R_2 are infinite rays entirely contained in G^J . It is possible to find such lines because $\iota(\epsilon), \tau(\epsilon) \subset G^J$. In particular, if an endpoint of ϵ is contained G^0 , then we may choose the ray based at that endpoint to be of the form $\sigma\sigma\sigma\dots$ for some closed Nielsen path σ . And if an endpoint of ϵ is not contained in G^0 , then we can choose the ray to represent some eigenray carried by G^J .

Let \mathcal{F}_ϵ be the smallest free factor system so that \mathcal{F}_ϵ carries L_ϵ and so that $\mathcal{F}_J \sqsubset \mathcal{F}_\epsilon$. Note that \mathcal{F}_{J+1} is a free factor system that satisfies these last two properties, and so we have $\mathcal{F}_\epsilon \sqsubset \mathcal{F}_{J+1}$.

Because $\delta(\ell) = \ell$ for every $\ell \in L_\epsilon$ and $\delta(\mathcal{F}_J) = \mathcal{F}_J$ by assumption, it follows that $\delta(\mathcal{F}_\epsilon) = \mathcal{F}_\epsilon$ as well. Similarly, because G^J is a filtration element for p , we know that $\pi(\mathcal{F}_J) = \mathcal{F}_J$. This fact paired with the fact that $p(\epsilon) = u\epsilon v$ with $u, v \subset G^J$ shows that L_ϵ is π -invariant also. Therefore, $\pi(\mathcal{F}_\epsilon) = \mathcal{F}_\epsilon$.

However, by item (Filtration) of the definition of a CT (Definition 3.12), the filtration of G must be reduced with respect to π , meaning that, because \mathcal{F}_ϵ is π -invariant, it must be equal to either \mathcal{F}_J or \mathcal{F}_{J+1} . Because \mathcal{F}_J does not carry every (or in fact, any) element of L_ϵ , it must be that $\mathcal{F}_\epsilon = \mathcal{F}_{J+1}$, and so $\delta(\mathcal{F}_{J+1}) = \mathcal{F}_{J+1}$ as desired.

□

Lemma 5.12. *(Corollary 3.2.2 of [BFH00]) If $f : G \rightarrow G$ is a topological*

representative and H_i is a stratum that consists of a single edge ϵ , then $f(\epsilon)$ crosses ϵ , in either direction, at most once.

6 The Main Theorem

Theorem 6.1. *Let $\phi, \psi \in \text{Out}(F_n)$ be PG, with $\phi^k = \psi^k = \pi$ forward rotationless. If $\phi|_{\mathcal{L}_N(\pi)} = \psi|_{\mathcal{L}_N(\pi)}$, then $\phi = \psi$.*

Proof. Let $p : G \rightarrow G$ be a CT representing π , and from the hierarchy of free factor systems associated to π and G , extract a totally ordered sequence $\mathcal{F} = \{\mathcal{F}_0, \dots, \mathcal{F}_t\}$ so that \mathcal{F}_{i+1} is a one edge extension of \mathcal{F}_i . By the definition of this hierarchy in Section 3.10, there is a sequence of core subgraphs $G^0 \subset G^1 \subset \dots \subset G^t$ of G such that G^i has a marking by \mathcal{F}_i , and each G^i is a filtration element G_r for some r . Note that each G^{i+1} is a one edge extension of G^i , and that $G^t = G$.

Set $\delta = \phi^{-1}\psi$. Because $\phi|_{\mathcal{F}_0(\pi)} = \psi|_{\mathcal{F}_0(\pi)}$ by Lemma 5.10, it follows that $\delta([F^i]) = [F^i]$ for each of the conjugacy classes of subgroups $[F^i]$ that make up $\mathcal{F}_0(\pi)$, and moreover, that $\delta|_{[F^i]} = \text{Id}_{[F^i]}$. And so any homotopy equivalence $d : G \rightarrow G$ of δ must be homotopic to a map which leaves G^0 invariant and is the identity when restricted to this subgraph. We will construct a homotopy equivalence d , beginning with the requirement that no homotopy is necessary, namely that $d(G^0) = G^0$ and $d|_{G^0} = \text{Id}_{G^0}$. This will form the base case of an inductive argument.

Up until this point, the proof has mimicked the opening steps of the proof

of Lemma 5.11. Here is where it diverges. Let G^i be any of our sequence of core subgraphs such that $d(G^i) = G^i$ and $d|_{G^i} = Id_{G^i}$. We will show that we can homotope d so that d is a topological representative of δ whose induced maximal filtration contains G^0, \dots, G^t as a subfiltration. In particular, we will show that the resulting map d satisfies $d(G^j) = G^j$ for all $j > i$. Afterwards, we will show that $d|_{G^{i+1}} = Id_{G^{i+1}}$, thus completing the induction. Because we are now trying to show we can find such a d that leaves G^{i+1} not only invariant up to homotopy as we did in Lemma 5.11, but fixed, the argument is longer and more difficult, and it requires being broken into several cases below.

Because we chose \mathcal{F} to be a sequence of one edge extensions, it follows that $G^{i+1} - G^i$ consists of a single edge ϵ with $\iota(\epsilon)$ and $\tau(\epsilon)$ contained in G^i , and with $p(\epsilon) = u\epsilon v$ for circuits u and v contained in G^i . When u and v are both nontrivial, one subdivides ϵ into two edges E and F so that $\epsilon = \bar{E}F$, $p(E) = E\bar{u}$, and $p(F) = Fv$.

Let C^{i+1} be the connected component of G^{i+1} that contains ϵ , and let $[F^{i+1}]$ be the conjugacy class of a free factor in \mathcal{F}^{i+1} represented by C^{i+1} . By Lemma 5.11, $[F^{i+1}]$ is δ -invariant, and so d leaves the graph C^{i+1} invariant up to homotopy. Let C^i be the connected component or union of connected components of G^i contained in C^{i+1} ; then $C^{i+1} = C^i \cup \epsilon$. By inductive assumption $d(C^i) = C^i$, but, a priori, $d(\epsilon)$ may include edges outside C^{i+1} . However, by the invariance of C^{i+1} up to homotopy, it is possible to deform d through continuous maps, each of which is the identity when restricted to C^i , until the resulting homotopy equivalence (which we shall still call d) leaves C^{i+1} invari-

ant. We can consider the restriction $d|_{C^{i+1}}$ of d to C^{i+1} , which is a homotopy equivalence representing $\delta|_{[F^{i+1}]} \in \text{Out}(F^{i+1})$.

We will now find additional homotopies that will transform $d|_{C^{i+1}}$ into a topological representative of $\delta|_{[F^{i+1}]}$. We must first check that $d|_{C^{i+1}}$ sends vertices to vertices. This is trivially true for every vertex in C^i , where $d|_{C^{i+1}}$ is the identity. Every vertex of C^{i+1} is contained in C^i except for the vertex that appears in the interior of ϵ after subdivision in the case that both u and v are nontrivial. By choosing a homotopy of d which is the identity everywhere except for on the interior of ϵ , we can assure that this vertex is sent to a vertex as well.

By the definition of a topological representative in Section 3.4 we must now check that $d|_{C^{i+1}}$ can be made locally injective on every edge. Again, we need only worry about ϵ . The current $d|_{C^{i+1}}$ may not be locally injective on ϵ but we can make it so by once again choosing a homotopy through continuous maps that differ only on ϵ and remain the identity on C^i . The resulting homotopy equivalence (which we shall still call $d|_{C^{i+1}}$) is a topological representative of $\delta|_{[F^{i+1}]}$.

The transition matrix of this topological representative is the identity on an m by m block, where m is the number of edges in C^i , and the entries in its final column corresponding to the image of ϵ . Therefore, in the maximal filtration of C^{i+1} induced by this topological representative (as defined near the end of Section 3.9), each edge of C^{i+1} is a single stratum. In particular, both C^i and C^{i+1} are filtration elements.

We are now in the situation where we can apply Lemma 5.12 to show that $d|_{C^{i+1}}(\epsilon)$ crosses ϵ at most once. However, if $d|_{C^{i+1}}(\epsilon)$ were to cross ϵ zero times, then $\delta|_{[F^{i+1}]}$ would not be an onto map, an impossibility. Therefore, $d(\epsilon)$ must cross ϵ exactly once. However, by assumption, the endpoints of ϵ are fixed by d (as they are contained in C^i), and so either $d|_{C^{i+1}}(\epsilon) = x\epsilon y$ or $d|_{C^{i+1}}(\epsilon) = x\bar{\epsilon}y$ for loops x and y in C^i . We can now extend all the homotopies we chose on $d|_{C^{i+1}}$ to homotopies of G by leaving d unchanged on $G - C^{i+1}$, so that we have d the identity on G^i , with $d(\epsilon) = x\epsilon y$ or $x\bar{\epsilon}y$.

We seek to show that the topological representative d that we've constructed can be modified so that both x and y are trivial and so that ϵ is not flipped to $\bar{\epsilon}$, meaning $d(\epsilon) = \epsilon$. We'll begin with a couple of claims. The first will provide us the with machinery to make that modification, and the second will greatly reduce the number of cases we need to consider. Recall from Section 3.4 that, for any continuous function $\sigma : I \rightarrow G$ from an interval into a marked graph, $[\sigma]$ is the path obtained by tightening σ .

Claim 1. Assume there is an annular Fix_N -line ℓ_ϵ represented by a directed line $\bar{R}_1 \epsilon R_2$ in G , where R_1 and R_2 are infinite paths entirely contained in G^i . Then $d(\epsilon) = x\epsilon y$, and not $x\bar{\epsilon}y$.

Further, if ℓ_ϵ can be chosen so that R_1 is not equal to either $[xxx\dots]$ or $[\bar{x}\bar{x}\bar{x}\dots]$, then x may be assumed trivial. If there is such an ℓ_ϵ so that R_2 is not equal to either $[yyy\dots]$ or $[\bar{y}\bar{y}\bar{y}\dots]$, then y may be assumed trivial.

Note: If x and y are circuits in G , meaning that $[x] = x$ and $[y] = y$, then

the brackets can be dropped from $xxx \dots, \bar{x}\bar{x}\bar{x} \dots, yyy \dots$, and $\bar{y}\bar{y}\bar{y} \dots$ in the statement of Claim 1. In the case that x and y are not circuits, the addition of the brackets ensures that ℓ_ϵ is a path, and hence is an element of $\mathcal{B}(G)$.

We will prove this claim by considering a lift of ℓ_ϵ to the universal cover \tilde{G} of G . Let $\tilde{\ell}_\epsilon = \tilde{R}_1^{-1}\tilde{\epsilon}\tilde{R}_2$ be such a lift, and let \tilde{r}_1 and \tilde{r}_2 be the initial and terminal endpoints of $\tilde{\ell}_\epsilon$. Because R_1 and R_2 are contained in G^i , they never cross the edge ϵ , and so ℓ is a non-periodic line. Therefore, by Lemma 3.3, there is a unique representative Δ of δ that fixes $(\tilde{r}_1, \tilde{r}_2) \in \tilde{\mathcal{B}}(F_n)$. If we let \tilde{d} be the lift of d to \tilde{G} corresponding to Δ , then $\partial\tilde{d}(\tilde{r}_1) = \tilde{r}_1$ and $\partial\tilde{d}(\tilde{r}_2) = \tilde{r}_2$ by assumption. And so $\tilde{d}(\tilde{\ell}_\epsilon)$ must tighten to $\tilde{\ell}_\epsilon$. But even further, because R_1 and R_2 are contained in G^i (where d is the identity), these two rays will be fixed pointwise by d .

If either of the R_i is a non-periodic ray, then \tilde{d} will fix the corresponding \tilde{R}_i pointwise. If R_i is a periodic ray $www\dots$ for some root-free loop w in G^i , then $\tilde{d}(\tilde{R}_i) = \tilde{w}_1 \dots \tilde{w}_j \tilde{R}_i$, where $\tilde{w}_1, \dots, \tilde{w}_j$ are lifts of w or w^{-1} , or else d would not fix R_i pointwise. Therefore, in any event, $\tilde{d}(\tilde{R}_i)$ never crosses $\tilde{\epsilon}$ or $\tilde{\epsilon}^{-1}$.

Assume that $d(\epsilon) = x\bar{\epsilon}y$. Then $\tilde{d}(\tilde{\ell}_\epsilon) = \tilde{d}(\tilde{R}_1^{-1})\tilde{d}(\tilde{\epsilon})\tilde{d}(\tilde{R}_2) = \tilde{d}(\tilde{R}_1^{-1})\tilde{x}\tilde{\epsilon}^{-1}\tilde{y}\tilde{d}(\tilde{R}_2)$, where \tilde{x} and \tilde{y} are lifts of x and y . As this line crosses $\tilde{\epsilon}^{-1}$ once and never crosses $\tilde{\epsilon}$, it cannot tighten to $\tilde{\ell}_\epsilon$, a contradiction. Therefore, $d(\epsilon) = x\epsilon y$, and not $x\bar{\epsilon}y$.

We now have that $\tilde{d}(\tilde{\ell}_\epsilon) = \tilde{d}(\tilde{R}_1^{-1})\tilde{x}\tilde{\epsilon}\tilde{y}\tilde{d}(\tilde{R}_2)$. This line can only tighten to $\tilde{\ell}_\epsilon$ if $\tilde{x}^{-1}\tilde{d}(\tilde{R}_1)$ tightens to \tilde{R}_1 and $\tilde{y}\tilde{d}(\tilde{R}_2)$ tightens to \tilde{R}_2 . Consider the tightening

of $\tilde{y}\tilde{d}(\tilde{R}_2)$ to \tilde{R}_2 : it can occur in one of three ways. The first case is that \tilde{R}_2 is a lift of the ray $[\bar{y}\bar{y}\bar{y}\dots]$ in G , and the action of \tilde{d} shifts this ray by one \bar{y} along this lift of the axis of the conjugacy class represented by y so that \tilde{y} may cancel with this extra lift. The second case is that \tilde{R}_2 is a lift of the ray $[yyy]\dots$ and \tilde{d} 's action shifts \tilde{R}_2 by one y along this lift of an axis toward \tilde{r}_2 , making room for the lift \tilde{y} . Therefore, if we rule out both of those possibilities for R_2 , the third and only remaining case is that \tilde{y} tightens to the trivial path based at $\tau(\tilde{\epsilon})$. This means that y tightens to the trivial path at $\tau(\epsilon)$, and we may choose this tightening homotopy so that it does not change the fact that d is the identity when restricted to G^i .

Similarly, we may force x to be trivial by ruling out the possibility that R_1 is contained in an axis for the conjugacy class represented by x . The only remaining possibility is that x tightens to the trivial path at $\iota(\epsilon)$. This completes the proof of Claim 1.

Claim 2.

- If u is a non-Nielsen path for p or if $\iota(\epsilon)$ is not contained in G^0 , then x may be assumed trivial.
- If v is a non-Nielsen path for p or if $\tau(\epsilon)$ is not contained in G^0 , then y may be assumed trivial.

Note that, if u or v is trivial, then it is a Nielsen path, as p fixes both endpoints of ϵ . Therefore, we do not need to specify that the non-Nielsen path is nontrivial. The same will be true in Cases 2 and 4 below.

We need only prove the first item of the claim, as the second item is the same fact with a change of notation. Our strategy will be to find an annular Fix_N -line of the type described in Claim 1 by which we can tighten away x .

To begin with, there is at least one point of ϵ which is fixed by p . If u and v are nontrivial, it is the common initial endpoint of E and F after subdivision. If one of them is trivial, it is an endpoint of ϵ , and if both are trivial, all points of ϵ are fixed. In any event, call such a fixed point $*$, consider the component \mathcal{C}_* of $\Gamma_G(\pi)$ determined by the Nielsen class of $*$, and let \hat{G} the covering space of G in which \mathcal{C}_* embeds. This component contains a lift $\hat{*}$ of $*$ which is contained in a lift $\hat{\epsilon}$ of ϵ .

Assume that u is a non-Nielsen path. Then the initial direction of ϵ determines an eigenray $\rho \in \mathcal{R}(\pi)$, represented by a point $\hat{\rho} \in \partial\mathcal{C}_*$ that is the endpoint of a ray \hat{R} in \mathcal{C}_* beginning at $\iota(\hat{\epsilon})$ and following the direction of \hat{u} . Let R be the image of \hat{R} under the covering projection to G . Because u is contained in G^i and because R is formed from u and its images under p , it follows that R is contained in G^i as well.

If $\tau(\epsilon)$ is contained in G^0 , then because G^0 is a core graph by definition, $\tau(\epsilon)$ is contained in at least one closed path σ in G^0 . If σ is composed entirely of p -fixed edges, then it is a closed Nielsen path for p . If there is no such σ , then every closed path in G^0 containing $\tau(\epsilon)$ contains at least one non-fixed linear edge E_σ , with $p(E_\sigma) = E_\sigma u_\sigma$ for some closed Nielsen path u_σ . Moreover, there is such an E_σ and a path σ' of fixed edges so that $\sigma' E_\sigma$ is contained in σ , and so that $\iota(\sigma') = \tau(\epsilon)$. In this case, the edge path $\sigma' E_\sigma u_\sigma \bar{E}_\sigma \bar{\sigma}'$ is a closed Nielsen

path for p . Therefore, in any event, there is a closed Nielsen path based at $\tau(\epsilon)$; call this Nielsen path γ .

This path γ is contained in G^0 by assumption, and so is contained in G^i . It follows that the line in G given by $\bar{R}\epsilon\gamma\gamma\gamma\dots$ satisfies all the requirements of Claim 1 to show that x may be assumed trivial. If instead $\tau(\epsilon)$ is contained not in G^0 , then there will be no such Nielsen path γ . However, if v is also a non-Nielsen path, then the terminal direction of ϵ determines a second eigenray ρ_v under iteration by p , this one represented by a ray R_v contained in G^i . In this case, we can use the line $\bar{R}\epsilon R_v$ and Claim 1 to show that d does not flip the edge ϵ and that we may assume x to be trivial.

We may assume that v is not a nontrivial Nielsen path, or else it could have played the role of γ in the previous paragraph. Therefore, we need only worry about the case that $\tau(\epsilon) \not\subseteq G^0$ and v is trivial. This means that each of the outward directions at $\tau(\epsilon)$ determines a growing edge for p . If none of these edges were to be contained in G^i , it would follow that $\tau(\epsilon)$ is not contained in G^i , a contradiction. Therefore, there is some other growing edge ϵ' that is contained in G^i and has initial vertex $\tau(\epsilon)$. Under iteration by p , ϵ' determines an infinite ray R' in G that represents an eigenray ρ' for π . Therefore, in \mathcal{C}_* we see an embedded line of the form $\hat{R}^{-1}\hat{\epsilon}'\hat{R}'$, where $\hat{\epsilon}'$ and \hat{R}' are lifts of ϵ' and R' . Once again, this line can be used with Claim 1 to show that d does not flip the edge ϵ and that x may be assumed trivial by tightening. This finishes the proof in the case that u is a non-Nielsen path.

Now assume instead that $\iota(\epsilon) \not\subseteq G^0$. This means there can be no nontrivial,

closed Nielsen paths for p based at $\iota(\epsilon)$, and so each of the outgoing directions at $\iota(\epsilon)$ determines a growing edge for p . If none of these edges is contained in G^i , then $\iota(\epsilon)$ is not contained in G^i , a contradiction. Therefore, one of these edges (and its resulting eigenray representative) is contained in G^i . Meanwhile, if $\tau(\epsilon)$ is contained in G^0 , then we find a ray of the form $\gamma\gamma\gamma\dots$ contained in G^i , where γ is a closed Nielsen path for p . If $\tau(\epsilon)$ is not contained in G^0 , then we find a growing and associated ray contained in G^i , or else face the contradiction that $\tau(\epsilon) \not\subseteq G^i$. In any event, we find rays based at both $\iota(\epsilon)$ and $\tau(\epsilon)$ that enable us to form a line satisfying the requirements of Claim 1. Therefore, once again, d does not flip the edge ϵ , and x may be assumed trivial. This completes the proof of Claim 2.

From here, the argument breaks into several cases based on the nontriviality and growth rates of u and v . Because of Claim 2, we need only consider cases in which at least one of u and v is trivial or fixed by p , ruling out the case that both u and v are nontrivial non-Nielsen paths. As we go through each case, it will be important to recall the argument that we used to prove Claim 2 above, as we will use similar tools below, often seeking to find an annular Fix_N -line to which we can apply Claim 1.

Case 1: The circuits u and v are both nontrivial Nielsen paths.

If so, then both $\iota(\epsilon)$ and $\tau(\epsilon)$ must be contained in G^0 . Therefore ϵ is a linear edge and is contained in G^0 as well. This means ϵ is fixed by d by assumption.

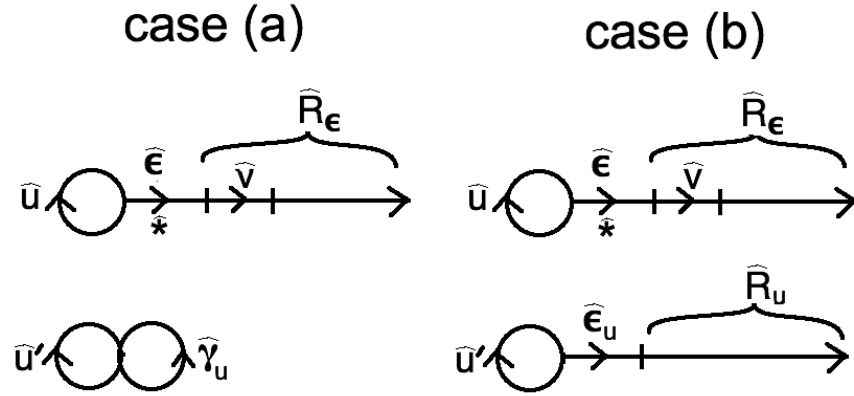
Case 2: One of u and v is a nontrivial Nielsen path, and the other is a non-Nielsen path.

Without loss of generality, we may assume that u is the Nielsen path, as the argument in the other case is symmetric. By Claim 2, y is trivial, and so we only need worry about x . As in the proof of Claim 2, consider the p -fixed point $*$ and the corresponding component \mathcal{C}_* of $\Gamma_G(\pi)$ containing lifts $\hat{*}, \hat{E}, \hat{F}, \hat{u}$, and \hat{v} . There is also an infinite embedded ray \hat{R}_ϵ based at $\tau(\hat{\epsilon})$, following the direction of \hat{v} , and representing an eigenray for π . In addition, because u is a Nielsen path, we know that \hat{u} is a loop.

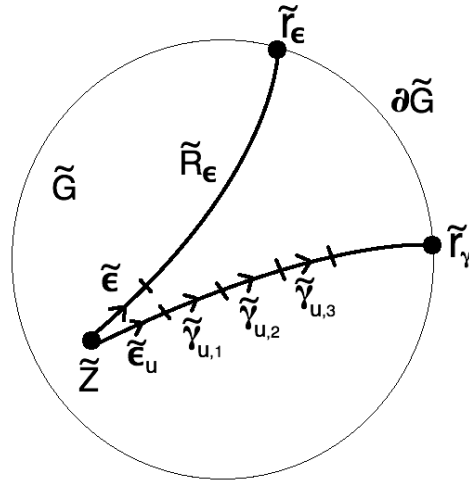
Now assume that x is nontrivial. Because u is a Nielsen path, we can consider the component \mathcal{C}_u of $\Gamma_G(\pi)$ corresponding to the Nielsen class of u . This component contains another lift \hat{u}' of u . And, because any lift \tilde{p}_u of p to \tilde{G} that corresponds to \mathcal{C}_u must be principal, \mathcal{C}_u must also contain one of the following:

- (a) a closed lift $\hat{\gamma}_u$ of a loop γ_u in G that is not homotopic to any root, power, or power of a root of x .
- (b) a lift $\hat{\epsilon}_u$ of a growing edge ϵ_u with corresponding infinite ray \hat{R}_u at its terminal vertex.

In case (a), consider the directed line in $\mathcal{B}(\mathcal{C}_*)$ represented by the infinite path $\hat{R}_\epsilon^{-1}\hat{F}^{-1}\hat{E}\hat{u}\hat{u}\hat{u}\dots$ and the directed line in $\mathcal{B}(\mathcal{C}_u)$ represented by $\dots\hat{u}'\hat{u}'\hat{u}'\hat{\gamma}_u\hat{\gamma}_u\hat{\gamma}_u\dots$. Choose lifts of these lines to \tilde{G} of the form $\tilde{R}_\epsilon^{-1}\tilde{F}^{-1}\tilde{E}\tilde{u}_1\tilde{u}_2\tilde{u}_3\dots$ and $\dots\tilde{u}_3^{-1}\tilde{u}_2^{-1}\tilde{u}_1^{-1}\tilde{\gamma}_{u,1}\tilde{\gamma}_{u,2}\tilde{\gamma}_{u,3}\dots$, where the \tilde{u}_i are all lifts of u , and the $\tilde{\gamma}_{u,j}$ are

Figure 13: Two components of $\Gamma_G(\pi)$ with lifts of the Nielsen path u 

all lifts of γ_u . These particular lifts can be concatenated and reduced to obtain a directed line of the form $\tilde{\ell}_a = \tilde{R}_\epsilon^{-1} \tilde{F}^{-1} \tilde{E} \tilde{\gamma}_{u,1} \tilde{\gamma}_{u,2} \tilde{\gamma}_{u,3} \dots$ which is a lift of an annular Fix_N -line ℓ_a for π .

Figure 14: A lift of an annular Fix_N -line to \tilde{G} in case (a)

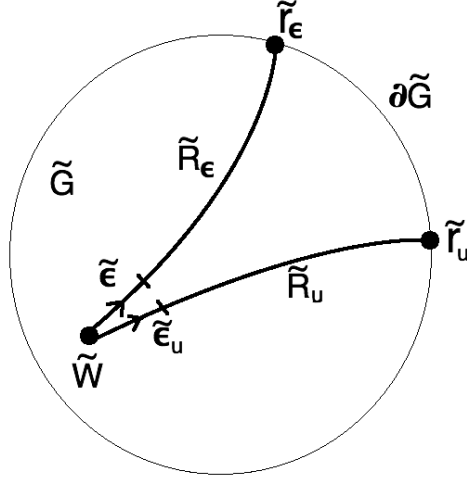
Let R_ϵ and γ_u be the images of \hat{R}_ϵ and $\hat{\gamma}_u$ under the covering map to G . The loop γ_u is contained in G^0 as it is a Nielsen path, and the ray R_ϵ is contained

in G^i because u is. Therefore ℓ_a satisfies the requirements of Claim 1, and so d does not flip the edge ϵ and x may be assumed trivial.

Similarly, in case (b), we will concatenate lifts of the lines $\hat{R}_\epsilon^{-1}\hat{F}^{-1}\hat{E}\hat{u}\hat{u}\hat{u}\dots$ and $\dots\hat{u}'\hat{u}'\hat{u}'\hat{\epsilon}_u\hat{R}_u$ in $\Gamma_G(\pi)$ to obtain the Fix_N -line ℓ_b represented by $\hat{R}_\epsilon^{-1}\hat{F}^{-1}\hat{E}\hat{\epsilon}_u\hat{R}_u$ and find a contradiction by considering its image under δ . In this case, choose a lift $\tilde{\ell}_b$ of ℓ_b of the form $\tilde{R}_\epsilon^{-1}\tilde{F}^{-1}\tilde{E}\tilde{\epsilon}_u\tilde{R}_u$, and let R_ϵ , ϵ_u , and R_u be the corresponding edge and rays in G .

Before moving on, note that the line represented by $\hat{R}_\epsilon^{-1}\hat{F}^{-1}\hat{E}\hat{u}\hat{u}\hat{u}\dots$ satisfies the requirements of Claim 1 to show that ϵ is not flipped by d . We will rely on this fact to complete the proof in this case.

Figure 15: A lift of an annular Fix_N -line to \tilde{G} in case (a)



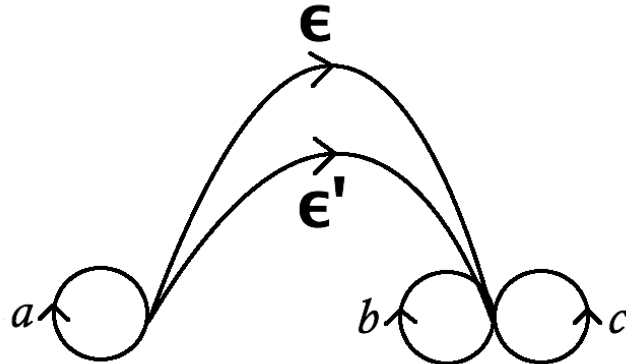
By assumption, R_ϵ is contained in G^i . If ϵ_u is contained in G^i , then R_u is as well, and we can apply Claim 1 to ℓ_b to show that d does not flip the edge ϵ and x may be assumed trivial. If $\epsilon_u \not\subseteq G^i$, then we may choose a different

such edge and ray in \mathcal{C}_u to serve as ϵ_u and R_u . If none of the edges leaving $\iota(\epsilon)$ is contained in G^i , then this means that in the component C of G^{i+1} that contains ϵ and u , no other edges emanate from $\iota(\epsilon) = \iota(u) = \iota(\epsilon)$, and so, because ϵ is not flipped by d , x is a power or root of u . The fundamental group of C is a free factor \mathbb{F}_C , and d 's restriction to C represents δ 's restriction to the conjugacy class of \mathbb{F}_C . But d is homotopic to the identity map on G^{i+1} , a homotopy given by unwinding the twisting of ϵ around x . Choose a homotopy of G that performs this unwinding while leaving all other components of G^{i+1} fixed, so that we may assume x to be trivial.

This completes the proof that $d(\epsilon) = \epsilon$ in Case 2.

Note: When performing the unwinding homotopy at the end of Case 2, twisting may be added to edges that are higher up in the filtration of G . That twisting will in turn be taken care of later in the induction. This point is illustrated by the following example.

Figure 16: A graph G on which to demonstrate the unwinding homotopy



Example 6.2.

Let G be the graph above, and let $p : G \rightarrow G$ be the CT described below, representing a forward rotationless PG $\pi \in \text{Out}(F_n)$. Suppose d below is a representative of δ . In its current state, d is the identity when restricted to $G^0(p) = a \cup b \cup c$, but we assume that d has not been shown or homotoped to be the identity on the edges ϵ or ϵ' . However, for simplicity, we will assume that d does not append any suffix to ϵ or ϵ' . In other words, assume that y and y' (in the language of this proof) have already been shown trivial.

$$\begin{array}{ll}
 a \mapsto a & a \mapsto a \\
 b \mapsto b & b \mapsto b \\
 p : \quad c \mapsto cb & d : \quad c \mapsto c \\
 \epsilon \mapsto a \epsilon c & \epsilon \mapsto a^3 \epsilon \\
 \epsilon' \mapsto a^2 \epsilon' c^2 & \epsilon \mapsto a^{-5} \epsilon'
 \end{array}$$

We must make a choice whether to homotope d to be the identity on ϵ first or on ϵ' . Choose to work with ϵ first and perform the unwinding homotopy on the graph $G^1 = a \cup b \cup c \cup \epsilon$ to force d to be the identity on this subgraph. However, this homotopy affects d 's action on ϵ' , resulting in the modified version of d below.

$$\begin{array}{ll}
 a \mapsto a & \epsilon \mapsto \epsilon \\
 d : \quad b \mapsto b & \epsilon \mapsto a^{-8} \epsilon' \\
 c \mapsto c &
 \end{array}$$

In this example, when we move up the induction to ϵ' , we do not need to use such an unwinding: we can show δ trivial using the annular Fix_N -line represented by $\bar{R}'\epsilon' \in R$, where R and R' are the eigenray representatives given by iterating the edges ϵ and ϵ' under p .

Case 3: One of u and v is trivial, and the other is a nontrivial Nielsen path.

Assume that u is the trivial one; again, the other possibility is symmetric. If $\iota(\epsilon)$ is contained in G^0 , then we are done, as ϵ is then also contained in G^0 and so is fixed by assumption. If $\iota(\epsilon)$ is not contained in G^0 , then by Claim 2, we know that x is trivial, and we may focus on y .

Because u is trivial, ϵ is not subdivided for the CT structure of p , and we will consider the p -fixed vertex $* = \iota(\epsilon)$, and the component \mathcal{C}_* of $\Gamma_G(\pi)$ corresponding to its Nielsen class. Because ϵ is a linear edge of G , by Lemma 3.21, \mathcal{C}_* must contain a lollipop consisting of a lift $\hat{\epsilon}$ of ϵ with a closed lift \hat{v} of v at its terminal endpoint and a lift $\hat{*}$ of $*$ as its initial endpoint. Because $*$ is not contained in G^0 , there cannot be any tight loop based at $\hat{*}$ other than roots and powers of $\hat{\epsilon}\hat{v}\hat{\epsilon}^{-1}$. However, because the lifts of π corresponding to this component of $\Gamma_G(\pi)$ must be principal, it follows that there is at least one growing edge ϵ_* and corresponding eigenray representative R_* based at $*$. This edge and ray lift to an edge $\hat{\epsilon}_*$ and \hat{R}_* based at $\hat{*}$ in \mathcal{C}_* . If none of the growing edges based at $*$ is contained in G^i , then $* = \iota(\epsilon)$ is not contained in G^i , a contradiction. So we may assume we have such an edge ϵ_* so that it and

R_* are contained in G^i .

From here, the argument is much the same as in Case 2. Consider the component \mathcal{C}_v of $\Gamma_G(\pi)$ corresponding to the Nielsen class of v . If \mathcal{C}_v contains any loop $\hat{\gamma}_v$ that is not homotopic to a lift of any root or power of y , then we may use the annular Fix_N -line represented by $R_*^{-1} \epsilon \gamma_v \gamma_v \gamma_v \dots$ and Claim 1 to show that ϵ is not flipped by d and that y may be assumed trivial. If there is no such $\hat{\gamma}_v$, then there must be a collection of growing edges iterating out to representatives. If any of these growing edges is contained in G^i , then we have an annular Fix_N -line with which we can apply Claim 1. If none of the edges is contained in G^i , then d 's restriction to this G^{i+1} is homotopically trivial. In any case, we may assume d does not flip ϵ and that y is trivial, and so we have that $d(\epsilon) = \epsilon$.

Case 4: One of u and v is trivial, and the other is a non-Nielsen path.

Here again, we assume u is trivial without loss of generality. Then by Claim 2, y is trivial. To show x is trivial, once more consider the component \mathcal{C}_* of $\Gamma_G(\pi)$ corresponding to the Nielsen class of the fixed initial vertex $*$ of ϵ . In this case, \mathcal{C}_* must contain an eigenray representative \hat{R}_ϵ based at the terminal endpoint of a lift $\hat{\epsilon}$ of ϵ .

If there is based at $\hat{*} = \iota(\hat{\epsilon})$ a loop $\hat{\gamma}$ that is not homotopic to a lift of any root or power of x , then the Fix_N -line represented by $\dots \hat{\gamma} \hat{\gamma} \hat{\gamma} \hat{\epsilon} \hat{R}_\epsilon$ can be used to show that ϵ is not flipped by d and that x may be assumed trivial by Claim 1.

If no such $\hat{\gamma}$ exists, then either there is a growing edge in G^i with which we can apply Claim 1, or there is no such edge and we can unwind the twisting as in other cases.

Case 5: Both u and v are trivial.

In this case, we will need to take note of the locations of the initial and terminal endpoints of ϵ . For instance, if both $\iota(\epsilon)$ and $\tau(\epsilon)$ are contained in G^0 , then ϵ is in G^0 as well, and our job is done. Similarly, if neither $\iota(\epsilon)$ nor $\tau(\epsilon)$ is contained in G^0 , then both are trivial by Claim 2. Therefore, the only remaining possibility is that one is in G^0 and the other is not.

Assume that $\iota(\epsilon) \subseteq G^0$ and $\tau(\epsilon) \not\subseteq G^0$, as once more, the other case is symmetric. By Claim 2, we know that y is trivial, so we must only check x . Once again, consider the component \mathcal{C}_* corresponding to the Nielsen class of any point $*$ of ϵ . (Every point of ϵ is fixed by p .) In this case, all we know a priori is that \mathcal{C}_* contains a lift $\hat{\epsilon}$ of ϵ .

Because of the positions of the two endpoints, based at $\iota(\hat{\epsilon})$ we have only loops, while based at $\tau(\hat{\epsilon})$ there are only infinite rays. As in the proof of Claim 2 and in Case 3, there must be a growing edge ϵ_* and corresponding eigenray representative R_* based at $\tau(\epsilon)$ that are contained in G^i , or else $\tau(\epsilon)$ would not be contained in G^i . Let $\hat{\epsilon}_*$ and \hat{R}_* be the lifts of ϵ_* and R_* to \mathcal{C}_* .

If there is any loop $\hat{\sigma}$ based at $\iota(\hat{\epsilon})$ that is not homotopic to a lift of any root or power of x , then Claim 1 applied to the Fix_N -line represented by $\dots \hat{\sigma} \hat{\sigma} \hat{\sigma} \hat{\epsilon}_* \hat{R}_*$

shows that we may assume ϵ is not flipped and x is trivial. If there is any other component \mathcal{C}_x of $\Gamma_G(\pi)$ which contains a closed lift of x , then we can follow the argument of Case 2 and check if \mathcal{C}_x also contains any loops or rays which will be carried by G^i and so enable us to apply Claim 1. If there exist no such loops or rays in either \mathcal{C}_* or in another component \mathcal{C}_x , then the lollipop composed of the edge ϵ and the loop x are isolated in G^{i+1} , and we can use our unwinding trick once more. Therefore, we may assume that both x and y are trivial in any case.

This completes the proof that we can construct d representing δ so that $d(\epsilon) = \epsilon$ in all possible cases, and so we have that d is the identity when restricted to the subgraph G^{i+1} . Applying induction shows that d as constructed is the identity on each of our subgraphs in turn, and in particular on $G^t = G$. Therefore, δ is the identity outer automorphism.

□

References

- [BK02] Nadia Benakli and Ilya Kapovich. Boundaries of Hyperbolic Groups. *Contemporary Mathematics*, 296:39-94, 2002.
- [B02] Mladen Bestvina. The topology of $Out(F_n)$. *Proc. of ICM (Beijing, 2002)*, 2:373-384, 2002.
- [BFH00] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $Out(F_n)$ I: Dynamics of exponentially-growing automorphisms. *Ann. of Math.*, 151(2):517-623, 2000.
- [BFH04] Mladen Bestvina, Mark Feighn, and Michael Handel. Solvable subgroups of $Out(F_n)$ are virtually Abelian. *Geom. Dedicata*, 104:7196, 2004.
- [BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math.*, 135(1):1-51, 1992.
- [BrH99] Martin Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, Berlin, 1999.
- [CL95] Marshall Cohen and Martin Lustig. Very small group actions on \mathbb{R} -trees and Dehn twist automorphisms. *Topology*, 34(3): 575-617, 1995.
- [CL99] Marshall Cohen and Martin Lustig. The conjugacy problem for Dehn twist automorphisms of free groups. *Comment. Math. Helv.*, 74: 179-200, 1999.

- [FaM11] Benson Farb and Dan Margalit. A Primer on Mapping Class Groups. Princeton University Press. Princeton, NJ, 2011.
- [FH11] Mark Feighn and Michael Handel. The recognition theorem for $Out(F_n)$. *Groups, Geometry, and Dynamics*, 5(1):39-103, 2011.
- [F80] Floyd, W. J. Group completions and limit sets of Kleinian groups. *Invent. Math.*, 57(3):205-218, 1980.
- [G40] I. A. Grushko. On the bases of a free product of groups. *Matematicheskii Sbornik*, 8:169-182, 1940.
- [HT85] Michael Handel and William P. Thurston. New proofs of some results of Nielsen. *Adv. in Math.*, 56(2):173-191, 1985.
- [KLV01] Sava Krstic, Martin Lustig, and Karen Vogtmann. An equivariant Whitehead algorithm and conjugacy for roots of Dehn twist automorphisms. *Proc. of the Edinburgh Mathematical Society*, 44: 117-141, 2001.
- [S83] John Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71(3):551-565, 1983.

Curriculum Vitae

1984	Born November 26th in Middle Island, New York, U.S.A.
1999-2002	Attended and graduated from Longwood High School, Middle Island.
2002-2006	Attended the College of Arts and Sciences of New York University in New York, New York, U.S.A.
2006	B.A.(Hons.) Mathematics and Journalism, New York University.
2006-2013	Graduate work in Mathematics, Rutgers University, Newark, New Jersey.
2006-2012	Teaching Assistantship, Department of Mathematics and Computer Science at Rutgers-Newark.
2012-2013	Dissertation Fellowship, Rutgers University-Newark.
2013	Ph.D. in Mathematical Sciences.