On automorphisms of free groups and free products and their fixed points.

Armando Martino

Queen Mary and Westfield College

Thesis submitted to the University of London

for the degree of Doctor of Philosophy

January 1998

Abstract

Free group outer automorphisms were shown by Bestvina and Handell to have fixed subgroups whose rank is bounded in terms of the rank of the underlying group. We consider the case where this upper bound is achieved and obtain combinatorial results about such outer automorphisms thus extending the work of Collins and Turner. We go on to show that such automorphisms can be represented by certain graph of group isomorphisms called Dehn Twists and also solve the conjuagacy problem in a restricted case, thus reproducing the work of Cohen and Lustig, but with different methods.

We rely heavily on the relative train tracks of Bestvina and Handell and in fact go on to use an analogue of these for free product automorphisms developed by Collins and Turner. We prove an index theorem for such automorphisms which counts not only the group elements which are fixed but also the points which are fixed at infinity - the infinite reduced words.

Two applications of this theorem are considered, first to irreducible free group automorphisms and then to the action of an automorphism on the boundary of a hyperbolic group. We reduce the problem of counting the number of points fixed on the boundary to the case where the hyperbolic group is indecomposable and provide an easy application to virtually free groups.

Acknowledgements

I would like to thank EPSRC for their support and Don Collins, my supervisor for introducing me to such an interesting area of mathematics. I would also like to thank Clare and the students in room 203 for their support and Ian Chiswell for interesting conversations.

Contents

Introduction **5**

1 Preliminaries **7**

1.1 Relative Train Track Maps 7

1.2 Stable Train Track Maps 13

2 Maximal Rank Outer Automorphisms **17**

2.1 Bases and Trees 17

2.2 Refining the structure 36

3 Dehn Twists **45**

3.1 Graphs of Groups and Dehn Twists 45

3.2 Relation to Maximal Rank 47

4 Conjugacy Problems **52**

4.1 A Combinatorial Approach 52

5 Index Theorem **67**

5.1 The Attractivity of Infinite Fixed Words 67

5.2 Main Theorem 81

6 Irreducible Automorphisms **103**

6.1 Train Tracks and Maximal Index 103

6.2 Rank Two Automorphisms 113

6.3 Normal Forms 116

7 Hyperbolic Groups **125**

Bibliography **135**

Introduction

The theory of $\mathbb{R}$-trees has been particularly rich especially in the context of free group automorphisms. In [CV86] the outer automorphism group of the free group of rank n, OutF_n, was shown to act on a finite dimensional contractible space. The closure of this space also turns out to be contractible (in [Sko]), compact (by [CM87]) and finite dimensional consisting of the space of *very small* actions (shown in [BF]) which were studied in [CL95]. What may be striking then is that $\mathbb{R}$-trees do not appear in what follows. Our main tool is the machinery of relative train tracks developed in [BH92] and extended in [CT94] to the case of free product automorphisms. However, $\mathbb{R}$-trees have influenced some of this work, for instance we have a generalisation of the main theorem in [GJLL] which is based solidly in the theory of $\mathbb{R}$-trees using the analysis of [Lev95]. The advantage of using relative train tracks is that it is quite straightforward to generalise proofs to the case of free products while $\mathbb{R}$-tree methods do not have (as yet) free product counterparts. The disadvantage is that with arbitary free products, one usually cannot exploit any of the rich geometry available for free groups, except by analogy and the proofs are thus less intuitive.

After introducing relative train tracks we consider in chapter 2 free group outer automorphisms whose fixed subgroups are in a certain sense maximal. We find good bases for these as well as assigning posets to them which describe how to build up maximal rank outer automorphisms from those of smaller rank.

In chapter 3, we use the ideas of the previous chapter to prove by induction that maximal rank outer automorphisms are in fact represented by dehn twist automorphisms. This is the converse of a result of [CL]. In that paper it was also shown that dehn twist automorphisms have a solvable conjugacy problem. This solution of the conjugacy problem applies to the maximal rank automorphisms (not outer) studied in [CT96]. In chapter 4 we provide an alternative solution which is combinatorial in nature and which we hope will generalise to free products.

We then move on to consider fixed points at infinity of automorphisms of free products. In chapter 5 we prove a close generalisation of the index theorem proved in [GJLL]. The main difficulty here is that the factor groups of the free product are not assumed to be geometric - that is they do not neccessarily satisfy a hypothesis such as hyperbolicity or automaticity. Thus a classification of infinite fixed words (Proposition 5.1.14) is cumbersome and ad hoc.

In chapter 6 we use the method of proof of the index theorem, which is significantly different to that in [GJLL], to describe irreducible automorphisms of free groups of maximal index and to apply a solution of the conjugacy problem given in [Los96].

Finally, in chapter 7, we consider free products which are hyperbolic and we consider the action of automorphisms on the boundary of such. We show that the free product points at infinity are contained in the boundary, and we manage to reduce the problem of studying the action of an automorphism on the boundary to that of the factor groups. As an easy application we therefore get an index theorem for certain virtually free groups, when considered as hyperbolic groups.

Chapter 1

Preliminaries

1.1 Relative Train Track Maps

Here we wish to present an outline of the theory of relative train track maps for free products. The results here were proven in [CT94], which in turn was a generalisation of the exposition in [BH92] of relative train track maps for free groups. As the free product case includes the free group situation we explain only the former theory so as to avoid repetition. Thus we will often give two references for facts pertaining to relative train track maps (in fact from the two sources above) as the proofs in the free product case are not dissimilar to the free group case. We note that if a graph of complexes $\mathcal{X}$, defined below, is actually a graph (so that we are dealing with a free group) then we write it X. The definitions for rank in the free product case also restrict to the usual definition of rank for free groups, so Theorem 1.1.2 actually has the Bestvina-Handel theorem as a special case.

We first give a notion of rank in the free product case (this is precisely the definition given in [CT94]). Let $G = *_{i=1}^{m} G_i$ be a free product where all the factors are freely indecomposable. Then if H is a subgroup of G we can write $H = F_s *_{j=1}^{t} H_i$, where the H_j are interserctions of H with conjugates of the factor groups of G and F_s is a free group of rank s which meets no such conjugate. (We note that it can be shown that the intersections of H with conjugates of the factor groups in the above decomposition can be indexed by double coset representatives. Also, either of s or t could be ∞, but since we are primarily concerned with finite ranks we do not distinguish infinite cardinals.) See [Kur34] for a proof or [Coh89] for a proof using Bass Serre theory.

Definition 1.1.1 If G is a group and $G = *_{i=1}^{m} G_i$, where each G_i is freely indecomposable, then m is the Kuroš rank of G – denoted K-rank(G). If H is a subgroup of G, so that $H = F_s *_{j=1}^{t} H_i$ as above, then $s + t$ is the Kuroš rank of H in G – denoted K-rank(G, H).

As in [CT94] we note that this number is well defined and invariant up to automorphism by [BL36]. The generalisation of the Scott Conjecture or Bestvina-Handel Theorem is then

Theorem 1.1.2 (3.11 [CT94]) *If $\phi : G \to G$ is an automorphism, then*

$$\text{K-rank}(G, \text{Fix}\phi) \leq \text{K-rank}(G)$$

Definition 1.1.3 *Graphs of Complexes:* A graph of complexes $\mathcal{X}$ is the union of a graph X with a family $\{C_j\}$ of 2-complexes called the *factor complexes*, where C_j is the multiplication complex of a freely indecomposable group G_j which is not isomorphic to an infinite cyclic group. Additionally we have an edge E_{C_j} joining the factor C_j to the graph X which is called the stem of C_j. We depart slightly from the terminology of [CT94] and term the edges of X along with the *stems* the *real* edges of $\mathcal{X}$. It is clear that the fundamental group $\pi_1(\mathcal{X})$ is the free product of the free group $\pi_1(X)$ with the groups G_j.

Definition 1.1.4 *Topological maps:* A map

$$f : \mathcal{X} \to \mathcal{X}'$$

between garphs of complexes is said to be *topological* if it is a continuous map satisfying the following conditions.

(1) f carries vertices to vertices.

(2) The restriction to each factor complex C_i is a homeomorphism onto a factor complex C_j sending edges to edges.

(3) Stems are carried to stems.

(4) Every edge of X can be subdivided as $[z_0, z_1, \ldots z_r]$ such that $f|_{[z_i, z_{i+1}]}$ maps $[z_i, z_{i+1}]$ to either an edge of X or to a path $E_{C_i} e \bar{E_{C_i}}$, where e is an edge in C_j. Furthermore, no two successive subintervals are mapped to inverse real edges or to paths meeting the same factor complex.

Topological maps are locally injective on real edges, but it can be shown that any continuous map between graphs of complexes is homotopic to a unique topological map, and the process of moving to a topological map is called *tightening*. In the special case where every positive iterate of our map is locally injective the map is called a *train–track map*. We shall discuss these for a free group below.

If $f : \mathcal{X} \to \mathcal{X}$ is a topological homotopy equivalence and μ is a path in $\mathcal{X}$ from $f(v)$ to v then the *path induced automorphism* $\pi_1(f, \mu) : \pi_1(\mathcal{X}, v) \to \pi_1(\mathcal{X}, v)$ is defined as:

$$\pi_1(f, \mu)([\alpha])] = [\mu f(\alpha) \bar{\mu}].$$

We shall in general omit the square brackets around paths (used to denote the homotopy equivalence class) and also use the notation $\bar{\alpha}$ to denote the inverse of the path α. We shall also use this 'bar' notation for elements of the free product. We note that if the path μ above is the trivial path then we call $\pi_1(f, \mu)$ a point induced automorphism and write it $\pi_1(f, v)$.

Definition 1.1.5 If $\phi \in \mathrm{Aut}G$, then a representative of ϕ is a topological self homotopy equivalence of a graph of complexes $\mathcal{X}$, a path μ in $\mathcal{X}$ and an isomorphism

$\tau : G \to \pi_1(\mathcal{X}, v)$ so that the following diagram commutes:

$$\begin{array}{ccc} G & \xrightarrow{\tau} & \pi_1(\mathcal{X}, v) \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \pi_1(f,\mu)} \\ G & \xrightarrow{\tau} & \pi_1(\mathcal{X}, v) \end{array}$$

The basic tool for studying automorphisms of G is by topological representatives, so one needs to ensure that some exist.

Proposition 1.1.6 ([CT94] 1.6) *Every automorphism of G has topological representatives.*

It is important to note that varying the path μ corresponds to varying an automorphism through its outer automorphism class. Thus results using topological representatives are often more naturally stated in terms of outer automorphisms.

It is further shown in [CT94] that every automorphism of G has an efficient representative. We shall outline the properties of such a representative, $f : \mathcal{X} \to \mathcal{X}$, below. Firstly one may define the *reduced* Kuroš rank of a graph of complexes $\mathcal{Z}$ which is not neccessarily connected. The reduced Kuroš rank of $\mathcal{Z}$, $\tilde{\text{K}}\text{-rank}(\mathcal{Z})$, with components $\mathcal{Z}_1, \ldots, \mathcal{Z}_k$ is defined to be:

Definition 1.1.7

$$\tilde{\text{K}}\text{-rank}(\mathcal{Z}) = 1 + \sum_{i=1}^{k} \{\text{Max}(0, ((\text{K-rank}(\pi_1(\mathcal{Z}_i)) - 1)\}$$

(If we have a graph X with no factor complexes this quantity is called *reduced rank* and is written $\tilde{rk}X$.)

It is then shown that there is an f-invariant stratification $\mathcal{X}_0 \subset \mathcal{X}_1 \subset \ldots \mathcal{X}_m$ of subgraphs of complexes of $\mathcal{X}$, where $\mathcal{X}_0$ consists of the factor complexes and the stems. The r^{th} stratum, $\mathcal{H}_r$, is the closure of $\mathcal{X}_r - \mathcal{X}_{r-1}$.

Given this stratification one may define the following:

Definition 1.1.8 A path, α, is called r-legal if on reduction of $f(\alpha)$, no edge of $\mathcal{H}_r$ is cancelled.

If $\alpha, \beta \subset \mathcal{X}_r$ are two r-legal paths, with common terminal vertex then the ordered pair (α, β) is called a turn. The turn is called r-legal if the path $\alpha\bar{\beta}$ is reduced as written and r-legal. It is called illegal otherwise.

There is an associated stratified graph of complexes $\Sigma_0 \subseteq \Sigma_1 \subseteq \ldots \Sigma_m = \Sigma$ and a map $p : \Sigma \to \mathcal{X}$ inducing maps $p : \Sigma_m \to \mathcal{X}_m$. Σ_0 is composed of factor complexes. The images of vertices of Σ are vertices of $\mathcal{X}$ fixed by f. The notation, Σ_r^v, is used to denote the component of Σ_r containing the vertex v.

If $\Sigma_r - \Sigma_{r-1}$ is non-empty then it has a single edge ϵ_r (and possibly a factor complex in case L(iii) below) such that $p(\epsilon_r)$ is an *Indivisible Nielsen path* (INP), which is a path in $\mathcal{X}$ which is fixed by f and cannot be decomposed as a reduced concatenation of fixed paths.

We may form the *transition* matrix, M, of f in $\mathcal{X}$ where after numbering the real edges of $\mathcal{X}$ we set the (i, j) entry of M to be the number of times the image of the i^{th} edge crosses the j^{th}. Each stratum determines a submatrix of M which is irreducible and we label the Perron-Frobenius eigenvalue of the $\mathcal{H}_r$, λ_r. (For a fuller discussion see [CT94], [BH92].) The stratum is called $\mathcal{H}_r$ growing, level or descending as λ_r is greater than one, equal to one or zero respectively. All efficient representatives have the relative train track properties. Namely,

Definition 1.1.9 (Relative Train Track Properties)

(1) The image of each edge of the r^{th} stratum under f begins and ends with edges from the r^{th} stratum.

(2) If $\alpha \subseteq \mathcal{X}_{r-1}$ has endpoints in $\mathcal{H}_r$, then $[f(\alpha)] \neq 1$.

(3) The f-image of an r-legal path is r-legal.

If λ_r is greater than one a length is assigned to each edge in $\mathcal{H}_r$ a length; this will define an r-length on paths in $\mathcal{X}$ (only non-zero for paths meeting $\mathcal{H}_r$). Then on applying the map f to an r-legal path, $\alpha \subset \mathcal{X}_r$, the r-length of α will be multiplied by a factor λ_r. If $\lambda_r = 1$ then it turns out that the associated transition matrix is a permutation matrix. In particular $\mathcal{H}_r$ consists of m edges, $\{E_1, \ldots E_m\}$, and $f(E_i) = a_i E_{\sigma(i)} b_i$, where a_i, b_i are paths in $\mathcal{X}_{r-1}$ and σ is a permutation of $\{1, \ldots, m\}$.

Fixed paths for efficient maps have several good properties, some of which may be deduced from the stratification and associated graph of complexes Σ, but which are actually used to prove the existence of Σ in [CT94][Proposition 3.11] and hence 1.1.2.

Proposition 1.1.10 ([CT94],2.11,3.5,3.6) *INP's of height r are classified as follows:*

(a) $\mathcal{H}_r$ is growing and there is a unique INP, ρ_r, of height r. Furthermore, $\rho_r = \alpha_r \bar{\beta}_r$, where α_r, β_r are r-legal paths starting and ending with edges in $\mathcal{H}_r$.

(b) $\mathcal{H}_r$ is level and is composed of a single edge E, and

L(i): ${\rho_r}^{\pm 1} = E\alpha$ for some $\alpha \subseteq \mathcal{X}_{r-1}$, or

L(ii): ${\rho_r}^{\pm 1} = E\beta\bar{E}$ for some $\beta \subseteq \mathcal{X}_{r-1}$ but not of Type L(iii), or

L(iii): ${\rho_r}^{\pm 1} = E\gamma x \bar{\gamma} \bar{E}$ for some $\gamma \subseteq \mathcal{X}_{r-1}$ and x an edge of a factor complex.

Additionally the important properties of reduced Kuroš rank are:

1.1.11

(1) $\tilde{\text{K}}$-rank$(\Sigma_0) \leq \tilde{\text{K}}$-rank$(\mathcal{X}_0)$

(2) $\tilde{\text{K}}$-rank$(\Sigma_{r-1}) \leq \tilde{\text{K}}$-rank$(\Sigma_r) \leq \tilde{\text{K}}$-rank$(\Sigma_r) + 1$

(3) $\tilde{\text{K}}$-rank$(\Sigma_{r-1}) < \tilde{\text{K}}$-rank$(\Sigma_r) \Rightarrow \tilde{\text{K}}$-rank$(\mathcal{X}_{r-1}) < \tilde{\text{K}}$-rank$(\mathcal{X}_r)$

(4) $\tilde{\text{K}}$-rank$(\mathcal{X}_{r-1}) \leq \tilde{\text{K}}$-rank$(\mathcal{X}_r)$.

Furthermore, $\pi_1(\Sigma, v) \cong \text{Fix}(\pi_1(f, v))$.

Thus an efficient map along with the associated graph of complexes encodes all the information on the fixed subgroups of the point induced automorphisms. Thus,to study automorphisms in this way, one need only show that:

Proposition 1.1.12 ([CT94],proposition 3.3) *If* $\phi \in \mathrm{Aut}G$ *is represented by a topological map* f *then* $\text{K-rank}(G; \mathrm{Fix}\phi) \geq 2$ *implies that* ϕ *is conjugate to a point induced automorphism by* f.

1.2 Stable Train Track Maps

To conclude our preliminary discussion we focus our attention on free groups and in particular on irreducible automorphisms. An outer automorphism $\Phi \in \mathrm{Out}F_n$, is called irreducible if every topological representative with no valence one vertices and no invariant forests (defined below) has an irreducible transition matrix.

A more algebraic characterisation is given in [BH92, lemma 1.3,1.16]:

Lemma 1.2.1 *If there are free factors* F_{n_i}, $1 \leq i \leq k$, $n_1 < n$, *such that* $F_{n_1} * F_{n_2} * \ldots * F_{n_k}$ *is a free factor of* F_n *and* Φ *cyclically permutes the* $F'_{n_i}s$, *then* Φ *is reducible. Conversely if* Φ *is reducible then there exist free factors of* F_n *permuted by* Φ *as above.*

Let $f : G \to G$ be a topological self homotopy equivalence of a graph. A forest in G is a contractible subgraph; it is not required to be connected. A subgraph $G_0 \subset G$ is called invariant if $f(G_0) \subseteq G_0$. Also note that an invariant forest is called pretrivial if some positive iterate of f applied to each edge of the forest is a vertex. One may collapse invariant and pretrivial forests by simply using the homotopy equivalence that maps each component to a point. The map induced by f will again be a topological representative.

As with relative train tracks, given a topological representative, $f : G \to G$, of an irreducible automorphism one can find the Perron Frobenius eigenvalue of the matrix λ and assign to each edge of G a length corresponding to the eigenvector associated to λ. If λ is greater than one, G then becomes a metric graph (paths have lengths). (Note that if λ is one then f is a finite order homeomorphism.) The length of the f-image of an edge e, is λ multiplied by the length of e (see [BH92]). When discussing lengths in the graph G we will always use the length given here. Also note that when applying this procedure one requires that f has no invariant forests.

So far we have concentrated on describing the properties of relative train tracks without reference to the proofs. However for irreducible automorphisms we will need to make use of the 4 basic operations defined in [BH92, pp. 11–16]. Throughout this $f : G \to G$ is a homotopy self equivalence of the graph G, which is a topological representative of Φ.

Suppose that $w \in G$ is not a vertex of G but that $f(w)$ is, then the process of changing the graph structure of G to make w into a vertex is called *subdivision*. This is the first basic operation.

If G has a valence one vertex then there is a homotopy equivalence that collapses the edge incident to the valence one vertex to a point. Then f induces a topological representative on the new graph after tightening and collapsing maximal invariant and pretrivial forests. This operation is called a valence one homotopy.

Suppose that G has a valence two vertex, v, at the end of the edge e_1 and at the start of e_2, and that e_1 is at least as long as e_2 (with respect to the lengths defined above). Perform the homotopy that stretches e_1 across e_1e_2 and remove the vertex that was v. Hence we end up with a graph, G', with one less edge and one less vertex. Call the new edge of the graph e. f induces a map on G' as follows. Any edge of G' other than e can be regarded as an edge of G so has a 'natural' f-image - we need only replace every occurence of e_1 with e and delete every occurence of e_2. Finally we tighten the induced map and collapse any pretrivial and invariant forests. This operation is called a valence two homotopy.

Now we come to the fourth operation. Suppose that two edges e_1 and e_2 start at the same vertex with nontrivial maximal initial segments e_i' respectively such that $f(e_1') = f(e_2')$ and these paths end in vertices of G. We then subdivide at the endpoints of e_i', identify e_1' with e_2' and then collapse any pretrivial and invariant forests. This last operation is called folding. It is called a *full fold* if either of e_i' is the whole edge and a *partial fold* otherwise.

A train track map is a relative train track, but with only one (trivial) invariant subgraph. Thus the concept of being legal applies to the whole graph and every edge in the graph is legal. (That is, every positive iterate of the train track map is locally injective on the interior of edges). It is shown in [BH92, Theorem 1.7] that any topological representative may be turned into a train track map by use of the four basic operations. Additionally, applying the basic operations to a train track map yields another train track map. (This fact also pertains to stable train track maps defined below.)

Just as in the relative train track case any indivisible Nielsen path, ρ, can be written $\alpha\bar{\beta}$ (reduced), where α and β are legal paths and the turn between α and β is illegal. The folding of this turn is termed folding the indivisible Nielsen path ρ. Given a train track map, f, the set $\mathcal{W}(f)$ consists of those train track maps that may be obtained by repeatedly folding indivisible Nielsen paths. Actually we do not wish to distinguish maps and graphs that are combinatoriallyequivalent. Two topological maps, f_1, f_2 of graphs G_1, G_2 respectively, are said to be *projectively equivalent* if, after rescaling the matric, there is an isometry $h : G_1 \to G_2$ taking edges to edges and vertices to vertices such that the following diagram commutes:

$$\begin{array}{ccc} G_1 & \xrightarrow{h} & G_2 \\ \downarrow{\scriptstyle f_1} & & \downarrow{\scriptstyle f_2} \\ G_1 & \xrightarrow{h} & G_2. \end{array}$$

Then $\mathcal{W}(f)$ is actually the set of projective equivalence classes of the train tracks obtained from f by the repeated folding of INP's.

One may then define stable train track maps although the fact of their existence is not obvious. In fact it turns out that a train track map which has the minimal number of INP's amongst all train tracks representing the same outer automorphism is stable. This, however, is not taken as the definition.

Definition 1.2.2 A train track map, $f : G \to G$, is called *stable* if no partial folds occur in the construction of $\mathcal{W}(f)$.

It is then shown that stable train track maps are very well behaved.

Proposition 1.2.3 *If $f : G \to G$ is a stable train track map then f has at most one Nielsen path and furthermore that the existence of a Nielsen path implies that there is exactly one illegal turn in X. (Neccessarily the illegal turn occuring in the Nielsen path).*

In fact the properties of relative train track maps given in the previous section are those of stable relative train track maps although we will not define this. Whenever dealing with irreducible outer automorphisms, stable train track maps turn out to be much more convenient to deal with than the ordinary kind. We make extensive use of them in later chapters.

Chapter 2

Maximal Rank Outer Automorphisms

2.1 Bases and Trees

In this section we wish to describe certain types of automorphism in a manageable combinatorial way starting from some algebraic information on fixed subgroups.

Our starting point is to define an equivalence relation on the set of automorphisms of the free group of rank n. In fact this equivalence relation is very natural when considering fixed subgroups and arises in [CM89] and [BH92].

Definition 2.1.1 Two automorphisms ϕ and ψ of F_n are said to be *similar* if there is a commuting diagram

$$\begin{array}{ccc} F_n & \xrightarrow{\gamma} & F_n \\ \downarrow{\scriptstyle\phi} & & \downarrow{\scriptstyle\psi} \\ F_n & \xrightarrow{\gamma} & F_n \end{array}$$

where γ is an inner automorphism. We term the equivalence classes under this equivalence relation, *similarity classes*, denoting by $[\phi]$ the similarity class containing ϕ.

Note that similar automorphisms are not only conjugate in AutF_n but also give rise to the same outer automorphism. We adopt the notation rkϕ, to stand for the rank of the fixed subgroup of ϕ, where the rank of a subgroup is the minimal number of generators. (Strictly speaking the rank of a subgroup H of F_n could be infinite in which case we would write rank $H = \infty$.) Also if G is a connected graph we use the notation rk$G$ to denote rank$\pi_1(G)$. As in the previous section, if G is not connected we use the quantity,

$$\tilde{\text{rk}}G = 1 + \sum(\text{rk}G_i - 1)$$

where the sum ranges over the non contractible components of G.

We now state the main theorem of [BH92]

Theorem 2.1.2 (The Scott Conjecture or Bestvina Handell Theorem)

Let $\phi_1, \phi_2, ..., \phi_k$ be automorphisms of F_n all belonging to the same outer automorphism class but to distinct similarity classes. Then,

$$\sum_1^k \max(0, \text{rk}\phi_i - 1) \leq n - 1$$

Bearing in mind the above theorem and the fact that all the automorphisms in a similarity class are conjugate and hence have isomorphic fixed subgroups, we make the following definition:

Definition 2.1.3 Let $[\phi]$ be a similarity class. We then call $[\phi]$ significant if rk$\phi \geq 2$.

By 2.1.2 any outer automorphism possesses only finitely many significant similarity classes and hence the following definition makes sense:

Definition 2.1.4 Let Φ be an outer automorphism of F_n, then we define

$$\text{rk}\Phi = 1 + \Sigma\ (\text{rk}\phi_i - 1)$$

where the sum ranges over representatives of the significant similarity classes of Φ.

Theorem 2.1.2 can then be reformulated thus:

Theorem 2.1.5 *Let Φ be an outer automorphism of F_n then,*

$$\text{rk}\Phi \leq n$$

We wish to consider the case where $\text{rk}\Phi = n$ and to this end we use the train track methods of [BH92] in the same way that the case $\text{rk}\phi = n$ was considered for a single automorphism in [CT96]. Thus we start with an outer automorphism of maximal rank ($\text{rk}\Phi = n$) and we represent this by a relative train track f on a stratified graph X. Then by [BH92] if Φ has k significant similarity classes $[\phi_1], ..., [\phi_k]$, then there are vertices $v_1, ..., v_k$ of X which are fixed by f and isomorphisms τ_i, such that the following diagrams commute:

$$\begin{array}{ccc} \pi_1(X, v_i) & \xrightarrow{\tau_i} & F_n \\ \Big\downarrow{\scriptstyle \pi_1(f,v_i)} & & \Big\downarrow{\scriptstyle \phi_i} \\ \pi_1(X, v_i) & \xrightarrow{\tau_i} & F_n \end{array} \tag{2.1}$$

The following proposition is precisely that of [CT96, proposition 3]

Proposition 2.1.6 *Let H_r be a growing stratum of X, then* $\tilde{\text{rk}}X_r \geq \tilde{\text{rk}}X_{r-1} + 2$.

This, taken with the fact from [BH92] that every stratum contains at most one indivisible Nielsen path shows, in the same way as [CT96], that a relative train track map representing a maximal rank *outer* automorphism cannot possess any growing strata.

Now from [BH92] every indivisible Nielsen path is of the form (using the terminology of [CT94])

$$E\alpha, \quad \alpha \subseteq X_{r-1} \quad \text{Type L(i)},$$
$$E\beta\bar{E}, \quad \beta \subseteq X_{r-1} \quad \text{Type L(ii)}$$

where $H_r = \{E\}$ and $f(E) = E\tau$ for some $\tau \subseteq X_{r-1}$.

However in [CT96] it is shown that we may find a relative train track f for Φ in which every Type L(i) stratum has α the trivial path and E a loop. From now on we only consider such representatives. Hence in relative train track maps for Φ which we consider, all indivisible Nielsen paths will be closed.

Now if for some vertex v which is fixed by f, we have that $\text{rk}\pi_1(f, v) \geq 2$ then there must be a Nielsen path joining v to some v_i (one of the vertices from 2.1). This follows by the properties of reduced rank. By definition of the vertices v_i, the only components of Σ with rank at least 2 are in fact the components containing the vertices v_i, Σ^{v_i}. As every edge of Σ maps to an indivisible Nielsen path in X, this demonstrates our claim. Note that since all our Nielsen paths are closed then such a v must equal v_i,for some i. Thus the only point induced automorphisms which are significant occur at one of the vertices $v_1, ..., v_k$ which we term Nielsen vertices.

These preliminary remarks aside, we now wish to construct a natural maximal tree for our relative train track map. First we prove two lemmas.

Lemma 2.1.7 *Suppose that ρ is a Nielsen path of height $r \geq 0$ and the component of X_r containing ρ has rank at least two, then the initial vertex of ρ is a Nielsen vertex.*

Proof: As ρ must be closed, $\tilde{\text{rk}}X_r > \tilde{\text{rk}}X_{r-1}$ and so $\tilde{\text{rk}}\Sigma_r = \tilde{\text{rk}}\Sigma_{r-1} + 1$. This last statement follows by the properties of reduced rank given in 1.1.11. As $\tilde{\text{rk}}\Sigma = \tilde{\text{rk}}X$ we get that $\tilde{\text{rk}}\Sigma_r = \tilde{\text{rk}}X_r$ for all r, and since $\tilde{\text{rk}}\Sigma_r \leq \tilde{\text{rk}}\Sigma_{r-1} + 1$ the claim follows. In fact the component of Σ_r contaning the initial vertex of ρ must have rank at least two, as the vertices of Σ map to fixed vertices of X. Thus, as in the above discussion there must be a Nielsen path from the initial vertex of ρ to a Nielsen vertex. As all Nielsen paths are closed, we see that the initial vertex of ρ is actually a Nielsen vertex.

The next lemma is an generalisation of the result of [CT93]. Our standing assumption is always that we are dealing with a relative train track map which represents a maximal rank outer automorphism and has no non closed Nielsen paths.

Lemma 2.1.8 *Let G be a component of X_r for some r, such that* $\text{rk}G = 1$. *Then there exists a connected $G' \supseteq G$ containing a Nielsen vertex v and so that $f(G') \subseteq G'$,* $\text{rk}G' = 1$ *and $\pi_1(f, v)$ is the identity automorphism on $\pi_1(G', v)$.*

Proof: Let s be the least integer such that the component of X_s containing G has rank at least 2. Call this component G_1. By definition of s,

$$\tilde{\text{rk}} X_s = \tilde{\text{rk}} X_{s-1} + 1$$

and hence

$$\tilde{\text{rk}} \Sigma_s = \tilde{\text{rk}} \Sigma_{s-1} + 1 \tag{2.2}$$

We thus have that there exists an indivisible Nielsen path of height s and hence H_s consists of a single edge E_s. Then by 2.2 we must have that for some Nielsen vertex v,

$$\text{rk}(\Sigma^v \cap \Sigma_s) = \text{rk}(\Sigma^v \cap \Sigma_{s-1}) + 1 \geq 2$$

By 2.1.7 the initial ~~edge~~ vertex of E_s is v. Now we let G_2 be the component of $G_1 - int(E_s)$ containing G if there is more than one component, and just $G_1 - int(E_s)$ otherwise.

First we consider the case where $v \in G_2$; in this case, combined with the fact that $\text{rk}(\Sigma^v \cap \Sigma_{s-1}) \geq 1$ we must have a Nielsen path in G_2, and as $\text{rk}G_2 = 1$ we get the result with $G' = G_2$.

Next we suppose that $v \notin G_2$, in which case E_s must be a separating edge whose terminal vertex lies in G_2. In this case if we set $G' = G_2 \cup E_s$, then $\text{rk}G' = 1$ and we again get the conclusion of the lemma as there is an indivisible Nielsen path of the form $E_s \beta \bar{E}_s$, where $\beta \subseteq X_{s-1}$ and hence $\beta \subseteq G_2$ as E_s is separating. Clearly we also have that $\pi_1(G', v) =< [E_s \beta \bar{E}_s] >$.

Consider now the collection C of all rank one subgraphs of X which are components of some stratum X_r. Clearly given any two such, either one is contained in the other or they are disjoint. Ordering them then by inclusion, let M be the set of maximal elements of C. So the elements of M are disjoint as subsets of X and also any element of C is contained as a subgraph in some element of M. Then for each element G of M we choose a single edge E of G so that $G - E$ is contractible. We also choose E to be in the highest stratum amongst the edges of G although this still does not give us a unique choice for E. The collection of edges obtained in this way we term the Type I edges.

Suppose that for some r, $\tilde{\mathrm{rk}} X_r = \tilde{\mathrm{rk}} X_{r-1} + 1$. Then the stratum H_r consists of a single edge E_r. We term the collection of all such edges E_r, where E_r is non separating when removed from X_r, the Type II edges. (Note it may be that in the graph X, some edge E is non separating whilst if we restrict our attention to some X_r the same edge E is separating.)

We then set B to be the collection of all Type I and Type II edges and we set $T = X - B$.

Proposition 2.1.9 *For any component C of some stratum X_s of X, $T \cap C$ is a maximal tree in C. In particular T is a maximal tree in X.*

Proof: Let C be as above and let v, w be vertices of C. First we will show that there is a path in $T \cap C$ from v to w. Define $B_r = X_r \cap B$. Now take some path γ in C joining the two edges v and w. Let r_0 be the least integer such that $\gamma \cap B = \gamma \cap B_{r_0} \neq \gamma \cap B_{r_0 - 1}$. In other words if we only consider the edges of γ which occur in B, γ has 'height' r_0.
case (i): $\tilde{\mathrm{rk}} X_{r_0} = \tilde{\mathrm{rk}} X_{r_0-1} + 1$, but H_{r_0} does not consist of a Type II edge. Here we have that H_{r_0} still consists of a single edge, but that this edge is separating in X_{r_0}. In this case $\gamma \cap B = \gamma \cap B_{r_0-1}$. So we reach a contradiction.
case (i)': H_{r_0} consists of a single Type II edge. Then by definition there must be a $\delta \subseteq X_{r_0-1}$ (possibly trivial) joining the endpoints of the Type II edge. Then replace in γ all occurrences of the Type II edge with δ. We note that as the Type II edge lies

in C and C is a component of some stratum we must have that $\delta \subseteq C$.

case (ii): $\tilde{\text{rk}} X_{r_0} = \tilde{\text{rk}} X_{r_0 - 1}$. In this case we either get that γ is a path which lies entirely in $T \cap C$ or that $H_{r_0} \cap B$ consists of a collection of Type I edges $E_1, ...E_k$. These each lie in disjoint members of M, that is rank one subgraphs of X. By definition there are paths $\delta_1, \ldots \delta_k$, where each δ_i joins the endpoints of E_i and lies in $X_{r_0 - 1}$. Thus we may replace each occurrence of an E_i with a δ_i, and as before we get that each $\delta_i \subseteq C$.

This describes a process of taking a path γ in C and replacing it after a finite number of steps with a path which lies entirely in $T \cap C$. Hence we have shown that $T \cap C$ is connected and contains every vertex of $T \cap C$. Our next step is to show that $T \cap C$ is contractible, and here we can omit the component C and just show that T is contractible.

Start with a closed path β of shortest length which lies in T. Let r_0 be the least integer such that $\beta \subseteq X_{r_0}$.

Let A be the component of X_{r_0} containing the path β.

case (i): $\text{rk}A \geq 2$. As $\beta \not\subseteq X_{r_0 - 1}$ we must have that $\tilde{\text{rk}} X_{r_0} = \tilde{\text{rk}} X_{r_0 - 1} + 1$. Now H_{r_0} must consist of a single edge. If this edge is separating then there must be a closed subpath of β which contradicts the minimality of β. If this edge is non separating then it must occur in β by the minimality of r_0 and this is a contradiction since we assumed that $\beta \subseteq T$.

case (ii): $\text{rk}A = 1$. Here A must be contained in some subgraph G which is an element of M. Thus there is a Type I edge E which lies in G such that $G - E$ is contractible. Thus β must contain the edge E which is a contradiction again since $\beta \subseteq T$.

Hence $T \cap C$ is a maximal tree in C.

Now we get back to our automorphisms $\pi_1(f, v_i)$. Fix one of these, $\pi_1(f, v_1)$ along with its isomorphism τ_1 from $\pi_1(X, v_1)$ to F_n. We then let δ_i be the unique (reduced) path in our maximal tree T from v_i to v_1. Define the isomorphism σ_i from $\pi_1(X, v_i)$

to $\pi_1(X, v_1)$ to be that induced by the map which takes a closed path α at the vertex v_i to the closed path $\bar{\delta}_i \alpha \delta_i$. We then have the following commuting diagram:

$$\begin{array}{ccccc} \pi_1(X, v_i) & \xrightarrow{\sigma_i} & \pi_1(X, v_1) & \xrightarrow{\tau_1} & F_n \\ \downarrow{\scriptstyle \pi_1(f, v_i)} & & \downarrow{\scriptstyle \pi_1(f, \mu_i)} & & \downarrow{\scriptstyle \phi_i' = \gamma_{g_i} \circ \phi_1} \\ \pi_1(X, v_i) & \xrightarrow{\sigma_i} & \pi_1(X, v_1) & \xrightarrow{\tau_1} & F_n \end{array}$$

where $\mu_i = f(\bar{\delta}_i)\delta_i$, $g_i = \tau_1(\mu)$ and γ_g denotes conjugation by g. We also call the composition of σ_i and τ_1, τ_i'.

The following proposition allows us to replace our random selection of representatives for the significant similarity classes of Φ with a more convenient choice.

Proposition 2.1.10 *The set $\{\phi_1, \phi_2', \phi_3', \ldots, \phi_k'\}$ form a complete set of representatives for the distinct significant similarity classes of Φ.*

Proof: As we know that there are k significant similarity classes of Φ and that each ϕ_i is conjugate to ϕ_i', it is enough to show that the members of $\{\phi_1, \phi_2', \phi_3', \ldots, \phi_k'\}$ are in pairwise distinct similarity classes. In fact, it is enough to show that the the members of $\pi_1(f, \mu_1), \pi_1(f, \mu_2), \ldots, \pi_1(f, \mu_k)$ are in pairwise distinct similarity classes. (μ_1 is the trivial path)

Suppose instead that we have for some $i \neq j$ a commuting diagram of the form:

$$\begin{array}{ccc} \pi_1(X, v_i) & \xrightarrow{\gamma_\beta} & \pi_1(X, v_j) \\ \downarrow{\scriptstyle \pi_1(f, \mu_i)} & & \downarrow{\scriptstyle \pi_1(f, \mu_j)} \\ \pi_1(X, v_i) & \xrightarrow{\gamma_\beta} & \pi_1(X, v_j) \end{array}$$

where by a slight abuse of notation we consider β as both a group element and a closed path in X starting at v_1.

The commutivity of the above diagram then implies that

$$f(\alpha)^{f(\beta)\mu_j} \simeq f(\alpha)^{\mu_i \beta}$$

for all closed paths α starting at v_1. Clearly this can only be the case if

$$f(\beta)\mu_j \simeq \mu_i\beta.$$

Recall that $\mu_i = f(\bar{\delta}_i)\delta_i$ (δ_1 is the trivial path) and so a little manipulation gives us that

$$f(\delta_i\beta\bar{\delta}_j) \simeq \delta_i\beta\bar{\delta}_j$$

which gives us a Nielsen path with endpoints v_i and v_j. This is a contradiction to the fact that all our Nielsen paths are closed and more generally to the fact that there cannot be a Nielsen path between v_i and v_j unless $\pi_1(f, v_i)$ and $\pi_1(f, v_i)$ represent similar automorphisms in F_n.

From now on we take the $\phi_i{}'$ and the $\tau_i{}'$ as our standard representatives of the similarity classes and isomorphisms and we drop the dash.

We define τ_{ij} to be ${\tau_j}^{-1} \circ \tau_i$ so that we have a commuting diagram:

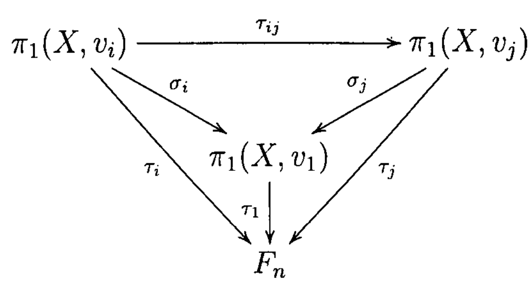

Thus it is clear that τ_{ij} is induced by the map which takes a closed path α at v_i to the closed path $\alpha^{\delta_i\bar{\delta}_j}$, where the path $\delta_i\bar{\delta}_j$ is (after possible reduction) the unique path in T from v_i to v_j.

Now each $\pi_1(X, v_i)$ has a 'natural' basis with respect to the maximal tree T. Namely if E is an edge of X not in T then there are unique reduced paths α_E, β_E in T joining v_i to the initial vertex of E and to the terminal vertex of E respectively. Then $\{\alpha_E E \bar{\beta}_E | E \in B\}$ defines a basis for $\pi_1(X, v_i)$. (Remember that $T = X - B$.) Let Y be the image of this basis in F_n. By the above remarks this image is the same for any choice of basepoint from the v_i. We shall denote by the basis element in Y corresponding to the edge E in B, τE. (Actually $\tau E = \tau_i \alpha_E E \bar{\beta_E}$ where the paths α_E, β_E are as above. The image we get is independent of the basepoint chosen.) We also note that if C is a component of some stratum X_r then as $T \cap C$ is a maximal tree in C,by 2.1.9, we also get a 'natural' basis for C. The image of this basis is then a subset of Y.

The isomorphisms τ_{ij} act not only as isomorphisms but also as bijections between these sets of natural basis elements even preserving the edge of B corresponding to a specific basis element. In other words, these bases and isomorphisms make it very easy to deduce what is happening at the group and automorphism level given information about our relative train track.

Before we prove the main result of this section we need two lemmas.

Lemma 2.1.11 *Let $H, K \leq F_n$ and $\phi \in \mathrm{Aut}F_n$ and suppose that $\langle H, K\rangle \cong H{*}K$ and also that $\phi(H) = H, \phi(K) = K$. Then if $\phi(v) = v^w$, for some $v, w \in \langle H, K\rangle$* but v is not in a conjugate of H or K, then some ~~$v \notin H \cup K$ we have that there is~~ *cyclically reduced conjugate of v is fixed by ϕ^m for some positive integer m.*

Proof: Given such a v we may find a cyclically reduced conjugate of the form

$$v' = h_1 k_1 \ldots h_s k_s$$

where $h_i \in H$ and $k_i \in K$ and we may assume that $h_1, k_1 \neq 1$.

Then $\phi^r(v')$ is a cyclically reduced conjugate of v' for all r. Thus for some integer m, $\phi^m(v') = v'$.

Lemma 2.1.12 *Let X be a basis for F_n and $\{x\} \cup X' \subseteq X$. Suppose that for some $\phi \in \mathrm{Aut}F_n$, $\phi(\langle X'\rangle) = \langle X'\rangle$ and $\phi x = xw'$ for some $w' \in \langle X'\rangle$. If ~~for some~~ $v \in \langle\{x\} \cup X'\rangle - \langle X'\rangle$,* and *$\phi v$ is a conjugate of v* but not of x *then there is a cyclically reduced conjugate v' of v which is fixed by ϕ^m for some positive integer m.*

Proof: Without loss of generality we may assume that we can take a cyclically reduced conjugate of the form

$$v' = xv_1x^{\epsilon_1}v_2x^{\epsilon_2}\dots x^{\epsilon_m}v_m$$

where $v_i \in \langle X' \rangle$ for all i and $v_m \neq 1$ and the $\epsilon_i = \pm 1$. It suffices to show that $\phi^r v'$ is cyclically reduced for all r.

For $w \in \langle \{x\} \cup X' \rangle$ let $|w|_x$ denote the number of occurrences of x in w (reduced). Now if w is cyclically reduced then

$$|w|_x \leq |w^g|_x \qquad \text{for any } g \in \langle \{x\} \cup X' \rangle$$

$$|\phi w|_x \leq |w|_x$$

(we note that the second property does not depend on w being cyclically reduced).

Note now that as v' begins with x, then $\phi v'$ is not cyclically reduced if and only if $|\phi v'|_x < |v'|_x$. However $\phi v'$ is a conjugate of v' and hence this is not possible. Moreover we may apply exactly the same argument to $\phi v''$, since we have now deduced that it is both cyclically reduced and begins with x. Hence inductively we have shown that $\phi^r v'$ is cyclically reduced for all r and again for some integer m, $\phi^m(v') = v'$.

We note that by the result of [DS75], for any automorphism ϕ of F_n and any integer r, Fixϕ is a free factor of Fixϕ^r and hence if we have a set of automorphisms $\{\phi_1, \dots, \phi_k\}$,

$$\sum_1^k \mathrm{rk}\phi_i - 1 \leq \sum_1^k \mathrm{rk}\phi^r - 1$$

Hence in the case of a maximal rank outer automorphism Φ, we get the following lemma as in [CT96].

Lemma 2.1.13 *If ϕ is a representative of a significant similarity class of Φ (that is* rk$\phi \geq 2$*) and for some integer r, $\phi^r v = v$ we actually get $\phi v = v$.*

We have up to now defined the set B which consists of edges which are either of Type I or of Type II. We further define the Type III edges analogously to the Type

II edges as follows. We call E a type III edge if it is of height r where $\tilde{\text{rk}}X_r = \tilde{\text{rk}}X_{r-1}$ and additionally, E is separating in X_r. Let B' denote the set of Type III edges and let P denote the union of B and B'.

We now describe a partial order on P. Let $E \in P$, and suppose that E has height r. Let C be the component of X_r containing E. We say that an edge in V is strictly less than E if it occurs as edge in $C - E$. It is clear that this defines a partial order on P.

Definition: If v, w are two elements of P we say that v *covers* w, denoted $v \rightarrowtail w$ if whenever $v \geq z > w$ then we have that $z = v$. (In other words there are no elements between v and w.)

Definition: For $v \in P$ we define

$$(i) \downarrow v = \{w \in P | w < v\}$$
$$(ii) \downarrow\!\!\!\!\rule{0pt}{0pt}\underline{\;} v = \{w \in P | w \leq v\}$$

Proposition 2.1.14 *All Type I edges in P are minimal, all Type II edges cover exactly one element of P and all Type III edges cover exactly two elements of P.*

Proof: Let E be an element of P of height r in X and let C be the component of X_r containing E. Now if E is a Type I edge then $C - \{E\}$ is contractible and hence cannot contain an edge from $B \cup B'$. If E is a Type II edge then $C - \{E\}$ is connected and non contractible thus must contain an edge from $B \cup B'$. Either rk$C - \{E\}$ is one, in which case $(C - \{E\}) \cap (B \cup B')$ consists of a single Type I edge or rk$C - \{E\} \geq 2$ in which case there is a maximal integer s such that $\text{rk}(X_s \cap C - \{E\}) = \text{rk}X_{s-1} \cap C - \{E\}$. Here we have that H_s must consist of a single edge (of $B \cup B'$) which is contained $C - \{E\}$. Thus this element is maximal in $(C - \{E\}) \cap (B \cup B')$ and thus a Type II edge covers exactly one edge. Finally if E is a Type III edge then $C - \{E\}$ consists of two non contractible components each of which (as before) contain maximal elements of $B \cup B'$ and thus we get that Type II edges cover exactly two other edges.

Given any partially ordered set we may construct a graph corresponding to the order. The vertices of the graph are the elements of the ordered set and given two elements v and w we say that there is an oriented edge from v to w if $v < w$. Bearing this in mind we have the following proposition.

Proposition 2.1.15 *The graph corresponding to P is a tree.*

Proof: It is enough to show that no element of P is covered by more than one other element of P. (This will show that there are no cycles in P, and we know that P must be connected since it has a maximal element. Namely this element is the unique edge of the highest stratum in which the rank increases.) So suppose that for some edge $E \in P$ we have that E is covered by two edges E_1 and E_2 of P and that E_1, E_2 are of height r_1, r_2 in X respectively. Without loss we may assume that $r_1 > r_2$. Let C_i be the component of X_{r_1-1} containing E_i for $i = 1, 2$. Then $E \in C_1 \cap C_2 \neq \emptyset$ and thus $C_2 \subseteq C_1$ as $r_1 > r_2$ and in fact we must have that $C_2 \subseteq C_1 - \{E_1\}$. Thus $E_1 > E_2$ as elements of P, which contradicts the fact that they both cover E.

We now wish to describe a labelling of the elements of P. It is important to note that this labelling need not be unique. That is, we may have two distinct elements of P with the same label.

Given a Type I edge the label we give it consists of a pair (y, ϕ) where $y \in Y$, the basis we defined earlier and $\phi \in \{\phi_1, \ldots, \phi_k\}$. We arrive at the label by use of lemma 2.1.8 which supplies us with the automorphism ϕ which fixes the basis element y corresponding to the Type I edge.

Type II edges are similarly labelled by pairs (y, ϕ) where the basis element y corresponds to the Type II edge as usual (via one of the isomorphisms τ_i or equivalently, via the map τ). The automorphism ϕ is supplied to us as the point induced automorphism from the initial vertex of the given Type II edge. This is clearly a Nielsen vertex by 2.1.7 as by properties of reduced rank there is an indivisible Nielsen path corresponding to the Type II edge.

Type III edges are edges E of height r in X such that $\tilde{\text{rk}} X_r = \tilde{\text{rk}} X_{r-1} + 1$ but which are separating in X_r. Again, the initial vertex of E must be a Nielsen vertex and there must be an indivisible Nielsen path of height r. We label a Type III vertex with the image in $\text{Aut}F_n$ of the corresponding point induced automorphism. (This is again an element of $\{\phi_1, \ldots, \phi_k\}$.)

We note that each element of Y occurs exactly once as part of a label of a member of P. It is precisely with the Type III vertices that non uniqueness of labelling arises.

Before we describe what information there is in this labelling we need some notation.

Let S be a subset of P. We shall call the elements of P vertices in deference to the graph structure of P. We say that ϕ_i occurs in S if there is a Type III vertex in S labelled by ϕ_i or there is a Type I or II vertex labelled by (y, ϕ_i) for some $y \in Y$. Similarly we say that $y \in Y$ occurs in S if there is a Type I or II vertex labelled by (y, ϕ) for some $\phi \in \{\phi_1, \ldots, \phi_k\}$.

Definition: For $S \subseteq P$, $\langle S \rangle = \langle y \in Y | y \text{ occurs in } S \rangle$

Lemma 2.1.16 *Let v be a Type I vertex labelled by (y, ϕ). Then v is covered by a unique vertex w where either w is a Type II vertex labelled by (y', ϕ) for some $y' \in Y$ or w is a Type III vertex labelled by ϕ. In other words ϕ occurs in the vertex covering v.*

Proof: The Type I vertex is actually a Type I edge, E,in X which belongs to some rank one subgraph, G, which in turn is a component of some stratum. Without loss (in fact as part of the definition of a Type I edge) we may assume that G is maximal amongst such. Following the construction of 2.1.8 we get an edge E' which is either a Type II or III edge such that E' covers E (this is clear from the description). However if we review the definition we see that the automorphism which occurs in v (actually the edge E) is the image of the point induced automorphism from the initial vertex of E'. This is also the automorphism which labels the element in P which is E'.

We also get the following result which tells us that when we can say if vertices of P are comparable.

Lemma 2.1.17 *Let z_1, z_2 be vertices of P in which the automorphism ϕ occurs. Then these are not comparable only if one of them is a Type I vertex covered by a Type III vertex.*

Proof: We first show that if neither of these vertices are Type I, then they must be comparable. In this case there are edges E_1, E_2 which give rise to the vertices z_1, z_2 of P (in fact they are the same object viewed in a different light). The fact that z_1, z_2 are not Type I and that the automorphism ϕ occurs in both means that E_1, E_2 must have the same initial vertex. (This is the Nielsen vertex whose corresponding point induced automorphism has image ϕ in F_n). Without loss we assume that E_1 is in a higher stratum than E_2. From this we immediately deduce that E_1 is greater than E_2, as they share a common endpoint, and hence z_1 is greater than z_2. Thus we have proved the lemma in the case where neither vertex is Type I. We now proceed to prove the lemma by contradiction. Suppose that z_1 and z_2 are not comparable and that neither is a Type I vertex covered by a Type III vertex. We may then replace these vertices by z_1', z_2' where $z_i' = z_i$ if z_i is not Type I and otherwise z_i' is a Type II vertex covering z_i in which the automorphism ϕ occurs. (We have made use of 2.1.16 here.) Now it is clear that if z_1, z_2 are not comparable then neither are z_1', z_2'. We then get the result by applying the argument above.

Corollary 2.1.18 *Let z be a Type III vertex labelled by ϕ and covering vertices v, w. Then no automorphism except ϕ may occur in both $\downarrow v$ and $\downarrow w$.*

Proof: A reiteration of the above argument yields the result.

Proposition 2.1.19 *If (y, ϕ) labels a Type I vertex then $\phi y = y$*

If (y, ϕ) labels a type II vertex z, then $\phi y = yw$ where $w \in \langle \downarrow z \rangle$. (In particular w contains no occurrences of y). Also $\phi \langle \downarrow z \rangle = \langle \downarrow z \rangle$ and hence $\phi \langle \updownarrow z \rangle = \langle \updownarrow z \rangle$.

We have that

$$\mathrm{Fix}\phi|_{\langle \updownarrow z \rangle} = \begin{cases} \mathrm{Fix}\phi|_{\langle \downarrow z \rangle} * \langle yv\bar{y} \rangle & \text{if } w \neq 1 \text{where } v \in \langle \downarrow z \rangle \\ \mathrm{Fix}\phi|_{\langle \downarrow z \rangle} * \langle y \rangle & \text{if } w = 1 \end{cases}$$

and $\mathrm{rk}\phi|_{\langle \updownarrow z \rangle} \geq 2$. If $\mathrm{rk}\phi|_{\langle \updownarrow z \rangle} = 2$ then ϕ does not occur in $\downarrow z$ unless $\downarrow z$ consists of a single type I vertex labelled by (y', ϕ) for some $y \in Y$.

Proof: If (y, ϕ) labels a Type I vertex then that $\phi y = y$ is the content of 2.1.8. Suppose then that (y, ϕ) labels a Type II vertex. In this case there is a Type II edge E in X whose initial point is the Nielsen vertex v_i and there is a path β in our maximal tree from the endpoint of E to v_i. Then E has a height r in X, so let C be the component of X_r containing E. By 2.1.9 β is a path in C. Thus all the elements of $\beta \cap B$ are strictly less than E. (Taking the order of our poset P.) This order is also preserved by our relative train track map f and so the same is true of $f(\beta)$. Then as $f(E\beta) \simeq E\tau f(\beta)$, where τ is a path in C not containing any occurrence of E we see that $\phi y = yw$ where w corresponds to the path $\bar{\beta}\tau f(\beta)$. We have shown that any occurrence of an edge of B in this path is necessarily strictly less than E. This is equivalent to the statement that $w \in \langle \downarrow z \rangle$. The train track properties of f give us that $f(C - \{E\} \subseteq C - \{E\}$ as E is the highest edge (with respect to the stratification of X) in C. Also there is an indivisible Nielsen path, of the same height as E, which is either of the form $E\gamma\bar{E}$ or it is a Type L(i) indivisible Nielsen path, in which case the edge E is a loop which is fixed by f. This proves the statement about the nature of the fixed subgroup. To see that $\mathrm{rk}\phi|_{\langle \updownarrow z \rangle} \geq 2$ it suffices to observe that $\tilde{\text{rk}}X_r = \tilde{\text{rk}}X_{r-1} + 1$ and hence $\text{rk}X_r \geq 2$. If $\text{rk}X_r > 2$ (and hence $\text{rk}\phi|_{\langle \bar{\downarrow} z\rangle} > 2$) then ϕ is point represented in $C - \{E\}$ by f. However, as all Nielsen paths are closed in our relative train track map there must be an indivisible Nielsen path of height lower than r at the vertex v_i. This is equivalent to saying that ϕ occurs in $\downarrow z$.

If $\text{rk}\phi|_{\langle \downarrow z\rangle} = 2$ then in general ϕ does not occur in $\downarrow z$. If it did occur in $\downarrow z$, then by considering the rank of the fixed subgroup we conclude that it must occur at a Type I vertex. By 2.1.16 we must have that $\downarrow z$ consists precisely of this Type I vertex.

We now proceed to prove a similar looking proposition for the Type III vertices.

Proposition 2.1.20 *Let z be a vertex of P which is of Type III. Then z covers exactly two vertices of P, z_1 and z_2 one of which will be distinguished, say z_1, and we have that $\phi\langle \bar{\downarrow} z_1\rangle = \langle \bar{\downarrow} z_1\rangle$, $\phi\langle \bar{\downarrow} z_2\rangle = \langle \bar{\downarrow} z_2\rangle$ and hence that $\phi\langle \bar{\downarrow} z\rangle = \langle \bar{\downarrow} z\rangle$. (Note that $\langle \bar{\downarrow} z\rangle = \langle \bar{\downarrow} z_1\rangle * \langle \bar{\downarrow} z_2\rangle = \langle \downarrow z\rangle$) We also have that*

$$\text{Fix}\phi|_{\langle \downarrow z\rangle} = \text{Fix}\phi|_{\langle \bar{\downarrow} z_1\rangle} * \text{Fix}\phi|_{\langle \bar{\downarrow} z_2\rangle}$$

that $\text{rk}\phi|_{\langle \bar{\downarrow} z_2\rangle} = 1$ and that ϕ does not occur in $\bar{\downarrow} z_2$ unless z_2 is a type I vertex labelled by (x, ϕ) for some $x \in X$. We have that $\text{rk}\phi|_{\langle \bar{\downarrow} z_1\rangle} \geq 1$ and as above that if $\text{rk}\phi|_{\langle \bar{\downarrow} z_1\rangle} = 1$ that ϕ does not occur in $\bar{\downarrow} z_1$ unless z_1 is a type I vertex labelled by (x, ϕ) for some $x \in X$. Also if $\text{rk}\phi|_{\langle \bar{\downarrow} z_1\rangle} > 1$ then ϕ will occur in $\bar{\downarrow} z_1$.

Proof: A Type III vertex of P is really a Type III edge, E, in X. Thus E has a height r, $\tilde{\text{rk}}X_r = \tilde{\text{rk}}X_{r-1} + 1$ and E is separating in X_r. Let C be the component of X_r containing E and let C_1 and C_2 be the two components of $C - \{E\}$. We suppose that the initial vertex of E, which is a Nielsen vertex v_i, lies in C_1. Then the distinguished vertex, z_1, corresponds to the to the topmost edge in C_1 which is in P. The statements of the proposition all follow easily as in the previous proposition. For instance $\pi_1(C, v_i) = \pi_1(C_1, v_i) * \pi_1(C_2 \cup \{E\}, v_i)$ is equivalent to the statement that $\phi \langle \updownarrow z_1 \rangle = \langle \updownarrow z_1 \rangle$, $\phi \langle \updownarrow z_2 \rangle = \langle \updownarrow z_2 \rangle$. We note that $\text{rk}\phi|_{\langle \updownarrow z_2 \rangle} = 1$ since otherwise ϕ is point represented in C_2, which would mean that v_i is a vertex in C_2 as all our Nielsen paths are closed. This clearly cannot happen as E is separating. The other statements follow as in the previous proposition.

The following was proved in [GJLL]. We state a slightly stronger version, which was also accessible with their methods.

Corollary 2.1.21 *Let $\Phi \in \text{Out}F_n$ be a maximal rank outer automorphism then there are elements $a, b \in F_n$ which generate a free factor of F_n and a $\phi \in \Phi$ such that:*

$$\phi a = a$$
$$\phi b = ba^m \qquad \text{for some integer } m.$$

Proof: One needs only find, by 2.1.16 and 2.1.19, a Type I vertex which is covered by a Type II vertex. If however we find that our Type I vertex is covered by a Type III vertex, z, then we just look for a minimal element in $\updownarrow z_1$ the distinguished vertex covered by z. We know from 2.1.20 that the rank of $\langle \updownarrow z_1 \rangle \geq 2$ so we may proceed in this way until we find the required Type I vertex covered by a Type II vertex.

The next proposition allows us to determine the nature of the fixed subgroups.

Proposition 2.1.22 *Let $w \in F_n$ and suppose that $\Phi \in \text{Out}F_n$ is a maximal rank outer automorphism which fixes the conjugacy class of w. Then there is a conjugate of w which is fixed by some $\phi \in \text{Aut}F_n$ with $\text{rk}\phi \geq 2$.*

Proof: We start by assuming that w is already cyclically reduced. There is no loss of generality in doing this. We then write w in terms of our natural basis and form the subset Y' of Y consisting of all those elements of Y which are used in the normal form of w. The set Y' corresponds to a set of vertices of P, namely those vertices for which some element of Y' occurs in the vertex. (We note that each element of Y occurs in exactly one vertex of P.) We now set z to be the vertex of P which is the least upper bound for the set of vertices defined by Y'. Such an element exists since P has a greatest element and is a tree by 2.1.15. Now if z is a type III vertex, then it is labelled by an automorphism ϕ, where rk$\phi \geq 2$. We may then apply lemmas 2.1.11 and 2.1.13 to get a (cyclically reduced) conjugate of w which is fixed by ϕ. Similarly if z is a Type II vertex labelled by (x, ϕ), then again rk$\phi \geq 2$ and we apply lemmas 2.1.12 and 2.1.13 to get the result. If z is a Type I vertex labelled by (x, ϕ) then by definition of z, $w \in \langle x \rangle$ and so is fixed by ϕ.

Note that the conditions of the lemmas 2.1.11 and 2.1.12 are satisfied precisely because z is a least upper bound.

It is useful to note that our constructions all respect the structure of the relative train track map and so whenever we prove a proposition like 2.1.22 we may be more explicit about which of our 'natural' basis elements we use in, for this example, the conjugation of w. In fact, suppose that we take a stratum X_r where $\tilde{\text{rk}} X_r = \tilde{\text{rk}} X_{r-1} + 1$ and we look at a component, C, of X_r. The relative train track map fixes C and in fact induces a maximal rank automorphism. We can produce a poset as before. However, we also have that the r^{th} stratum will consist of a single Type II or Type III edge which will correspond to a vertex z of our original poset. It is an easy exercise and a consequence of 2.1.9 that the poset and labelling that we get by restricting our map to the subgraph C is exactly the same as the 'sub-poset' $\updownarrow z$ of P. The following proposition is an easy consequence of these remarks. As always Φ is an outer automorphism of maximal rank.

Proposition 2.1.23 *If z is a vertex in P, then $\Phi|_{\updownarrow z}$ is a maximal rank automorphism. In addition, the restriction of the automorphisms which occur in $\updownarrow z$ form a complete set of representatives of the significant similarity classes of $\Phi|_{\updownarrow z}$.*

It is now clear that we may improve on 2.1.22.

Corollary 2.1.24 *Let $w \in \langle \downarrow z \rangle$ be conjugate fixed by Φ. Then there is a $u \in \langle \downarrow z \rangle$ and an automorphism ϕ, with $\mathrm{rk}\phi|_{\langle \downarrow z \rangle} \geq 2$ which occurs in $\downarrow z$ and fixes w^u.*

Proof: We may either reiterate the argument of 2.1.22 or apply 2.1.23.

2.2 Refining the structure

The description given in the last section will allow us to prove almost anything we want about maximal rank outer automorphisms of F_n, but there are ways to make things a little bit more convenient. That is the goal of this section. Our starting point as always is an outer automorphism of maximal rank, Φ. Suppose now that we have a poset P, a basis $\{x_1, \ldots, x_n\}$ of F_n and a complete set of representatives of the significant similarity classes of Φ, $\{\phi_1, \ldots, \phi_k\}$. Suppose that every vertex of P is either minimal or covers exactly one or two vertices and refer to these vertices as Type I, II and III respectively. We also suppose that there is a labelling of the vertices of P, where Type I and Type II vertices are labelled by (x, ϕ) where x is an element of our basis and $\phi \in \{\phi_1, \ldots, \phi_k\}$, and Type III automorphisms are labelled by one of the automorphisms in $\{\phi_1, \ldots, \phi_k\}$. We require that a basis element occurs in exactly one vertex. We then say that P is an allowable poset (for Φ) if it satifies the conclusions of 2.1.15, 2.1.16, 2.1.19 and 2.1.20. The results 2.1.18 and 2.1.23 must therefore also be satified by an allowable poset. Actually all of these results are consequences of 2.1.19 and 2.1.20 as the description on how the ranks of the fixed subgroups go up has many implications for the allowable poset.

We note that it can be shown that given an allowable poset P, for Φ it is possible to construct a relative train track map representing Φ which gives rise to the poset, P as in the last section. So, in a sense, what follows is as much a dicussion about different relative train track representatives for Φ as it is about allowable posets.

Suppose that we start with an allowable poset P and we replace one of the elements of the basis x_i with $x'_i = x_i w_i$ where $w_i \in \langle \downarrow z \rangle$ and z is the unique vertex of P which x_i occurs in. This has the effect of producing a new object with the same underlying poset as P but with a different basis and so a different labelling. We call this process right multiplication and it is clear that the resulting object is an allowable poset. This is because the underlying poset is unchanged and the contents of 2.1.16,2.1.19 and 2.1.20 are essentially unaffected by right multiplication as an easy check will show. Hence we have:

Lemma 2.2.1 *Performing right multiplication on an allowable poset results in another allowable poset.*

Right multiplication allows us to change an element of the basis but we would also like to be able to change the representatives of our significant similarity classes in the following way. Suppose we have a Type III vertex z in our allowable poset P which covers the vertices v and w. Let ϕ be the automorphism which labels z and suppose that ϕ does not occur in $↧ w$. Let g be an element in $\langle ↧ w \rangle$. We then replace any automorphism ψ which occurs in $↧ w$, by $\psi' = \gamma_g \psi \gamma_{\bar{g}}$ and we also replace any of our basis elements x which occur in a label with one of these ψ by $x' = x^g$. Note that x does not neccessarily occur in $↧ w$, we only require that there is some vertex of P labelled by (x, ψ) for some ψ which does occur in $↧ w$.

Similarly, suppose that we have a Type II vertex z, in which the automorphism ϕ occurs and that ϕ does not occur in $\downarrow z$. Let $g \in \langle \downarrow z \rangle$. We then replace ϕ with $\phi' = \gamma_g \phi \gamma_{\bar{g}}$ and any basis element x which occurs with ϕ by $x' = x^g$. We call either of these operations, changing the representative of $[\phi]$.

The effect of changing the represenative of $[\phi]$ is to change the conjugating element between one of our automorphisms and the automorphisms which are 'below' it. This will enable us to choose a convenient set of represenatives of the similarity classes of

Φ without disturbing the allowable poset structure. For this to be valid we need the following:

Lemma 2.2.2 *Changing the representative of* $[\phi]$ *results in another allowable poset.*

Proof: This is another tedious but easy check.

So far we have two operations which change the basis and the representatives for Φ. We now give an operation which gives a very small change to the underlying poset. Namely, suppose we have a Type III vertex, z, which covers a Type I vertex v. Then if ϕ is the label for z we know from 2.1.16 that there is a basis element x so that (x, ϕ) is the label for v. Also, from 2.1.19 $\phi x = x$. Our operation is then to delete the Type I vertex v thus making z into a Type II vertex which we label (x, ϕ). We call this deleting an unneccssary Type I vertex. It is trivial to show that:

Lemma 2.2.3 *The operation of deleting an unneccssary Type I vertex results in an allowable poset.*

We now wish to use these operations to get a better allowable poset. Let us describe the situations that we wish to avoid.

Suppose that we have a Type II vertex z labelled by (x, ϕ). Then $\phi x = xw$ for some $w \in \langle \downarrow z \rangle$, by 2.1.19. Consider the situation where $w \neq 1$ and so $\mathrm{Fix}\phi|_{\downarrow z} = \mathrm{Fix}\phi|_{\downarrow z^*} < xv\bar{x} >$ for some $v \in \langle \downarrow z \rangle$. By 2.1.24 there is an automorphism ψ which occurs in $\downarrow z$ and an element $g \in \langle \downarrow z \rangle$ such that ψ fixes v^g. We call this vertex a bad vertex if we cannot take this g to be 1.

There are two other kinds of vertices which we shall term bad. The first is where we have a Type II vertex, z, in which the automorphism ϕ occurs but for which ϕ does not occur in $\downarrow z$. By 2.1.19 we know that $\mathrm{rk}\phi|_{\langle \downarrow z \rangle} = 1$, and so there is a $v \in \langle \downarrow z \rangle$ fixed by ϕ. Again by 2.1.24 we may find a $g \in \langle \downarrow z \rangle$ and an automorphism ψ which occurs in $\downarrow z$ (and hence is different from ϕ) so that ψ fixes v^g. We call z bad if we cannot take g to be 1. Similarly if we have a Type III vertex z labelled by ϕ and covering the vertex w such that ϕ does not occur in $\updownarrow w$ then $\mathrm{rk}\phi|_{\langle \updownarrow z \rangle} = 1$, by 2.1.20. As above we find an automorphism ψ different from ϕ and occuring in $\updownarrow w$ so that ψ fixes v^g for some $g \in \langle \updownarrow w \rangle$. Again we call z bad if we cannot take g to be 1.

Definition: We call an allowable poset strong if it has no bad vertices and no Type I vertex is covered by a Type III vertex.

We now apply the operations above to get a strong allowable poset. Such an object is an analogue of a strong right layered basis for a single maximal rank automorphism as produced in [CT96].

Theorem 2.2.4 *Given a maximal rank outer auotmorphism Φ of F_n, there exists a strong allowable poset for Φ.*

Proof: We are given an allowable poset, P, for Φ which comes from a relative train track representative of Φ. We may assume that in this poset no Type I vertex is covered by a Type III vertex by repeatedly deleting unneccessary Type I vertices. We note that our other two operations do not change the underlying poset so this condition will continue to be satisfied if we apply them. Suppose that z is a bad vertex of P but that $\downarrow z$ is a strong allowable poset for $\Phi|_{\langle \downarrow z \rangle}$ (if z is a Type III vertex we assume the corresponding statement for each of the vertices that z covers). We note that this makes sense in the light of 2.1.23. We will then show that we may perform one of our operations on z so that $\updownarrow z$ is a strong allowable poset for $\Phi|_{\langle \updownarrow z \rangle}$. Thus by starting at the Type I vertices and working up we will have shown the result.

So suppose that z is a bad Type II vertex. Take the case where z is labelled by (x, ϕ), ϕ fixes $xv\bar{x}$ and there is a ψ fixing v^g. Perform a right multiplication on P which replaces x with $x' = x^g$. Then ϕ fixes $x'v^g\bar{x'}$ and we have got rid of a bad Type II vertex. We note that right multiplication leaves $\downarrow z$ unaffected and so assuming that $\downarrow z$ was a strong allowable poset gives us that $\bar{\downarrow} z$ is a strong allowable poset after the right multiplication.

Now suppose that z is a bad Type III vertex labelled by ϕ and covering the vertices v, w. Suppose that ϕ does not occur in $\bar{\downarrow} w$ and that $\bar{\downarrow} v$ and $\bar{\downarrow} w$ are strong allowable posets for $\Phi|_{\langle \bar{\downarrow} v\rangle}$ and $\Phi|_{\langle \bar{\downarrow} w\rangle}$ respectively. Thus we are given a ψ occurring in $\bar{\downarrow} w$ and elements $v, g \in \langle \bar{\downarrow} w\rangle$ such that ϕ fixes v and ψ fixes v^g. We then change the represenative of ϕ using the element g. This actually has the effect of leaving ϕ unchanged and replacing ψ with $\psi' = \gamma_g \psi \gamma_{\bar{g}}$ (along with some other changes). Thus both ϕ and ψ' fix the element v. As we have relabelled everything in $\bar{\downarrow} w$ in a consistent way, and that by 2.1.18 $\bar{\downarrow} v$ is unaffected we get that $\bar{\downarrow} z$ is now a strong allowable poset.

The remaining case of a bad Type II vertex is similar, so that we have shown what we set out to do and we get the result.

Corollary 2.2.5 *If (x, ϕ) labels a Type II vertex z in a strong allowable poset then*

$$\phi x = xv\bar{h}$$

where $\psi = \gamma_h \phi$ occurs in $\downarrow z$, $\psi v = v$ and ϕ fixes $xv\bar{x}$.

Proof: This is a consequence of having no bad Type II vertices. We know that for some $\psi = \gamma_h \phi$ and some g, ϕ fixes $xg\bar{x}$ and ψ fixes g. If $\phi x = xw$ then

$$g = \psi g = g^{wh}$$

and hence wh is equal to some power of the root of g, so that $w = v\bar{h}$ for some v which is fixed by ψ and so that $xv\bar{x}$ is fixed by ϕ.

Corollary 2.2.6 *Let P be a strong allowable poset for Φ. Then for every automorphism ϕ_i which occurs in P there is an automorphism $\phi_j = \gamma_g \phi_i$ which occurs in P and such that $g \in \mathrm{Fix}\phi_i \cap \mathrm{Fix}\phi_j$*

Proof: Consider the subset S of vertices of P in which ϕ_i occurs. Suppose that there is a Type II or Type III vertex which is minimal amongst these. In other words, there is a vertex z in which ϕ_i occurs and such that ϕ_i does not occur in $\downarrow z$. We then satisfy the conclusions of the above corollary as z cannot be a bad vertex. On the other hand if we cannot find such a vertex then ϕ_i must label a Type I vertex. We then concentrate on the unique vertex z above this Type I vertex such that no automorphism other than ϕ_i occurs in $\downarrow z$ but such that $\phi_j \neq \phi_i$ occurs in z. Again the fact that z is not a bad vertex gives us the conclusion.

In [CT96] a maximal rank automorphism (not outer) is considered and a good basis is found for it. That is, a strong right layered basis.

Definition 2.2.7 Let ϕ be a maximal rank automorphism of F_n, then the basis $\{x_1, \ldots, x_n\}$ is called a strong right layered basis if:

(i) $\phi x_i = x_i w_i$ with $w_i \in < x_1, \ldots x_{i-1} >$

(ii) For each i, $w_i \in$ Fixphi.

(iii) If $w_i \neq 1$ then there does not exist a $u \in < x_1, \ldots x_{i-1} >$ with $w_i = {u_i}^{-1} \phi u_i$.

Corollary 2.2.8 *If ϕ is a maximal rank automorphism then there is a strong right layered basis for ϕ.*

Proof: The outer auotmorphism Φ corresponding to ϕ gives rise to a strong allowable poset, P, with no Type III vertex, since such a Type III vertex would have to cover a Type I vertex. (As Φ has only one significant similarity class.) Thus P is linear (has a linear order) and this is the same as having a right layered basis. The previous corollary gives us a strong right layered basis.

Proposition 2.2.9 *If Φ has exactly two significant similarity classes then we may find a strong allowable poset for Φ which is linear.*

Proof: Take a strong allowable poset P. Suppose that z is a Type III vertex labelled by ϕ and covering verices v, w. Suppose that the only automorphism to occur in $\downarrow v$ is ϕ and that there is some other automorphism ψ which is the only automorphism to occur in $\downarrow w$. Then since no Type III vertex can cover a Type I vertex, $\downarrow v$ and $\downarrow w$ must be linear. Change the poset by deleting the vertex z and by placing the (unique) minimal vertex of $\downarrow v$ above w. This results in another strong allowable poset and after finitely many stages we get a linear strong allowable poset.

We may rephrase the previous proposition in less technical language as follows. Given an outer automorphism of F_n of maximal rank and with precisely two siginificant similarity classes, we may find a basis $\{x_1, \ldots, x_n\}$ and two representatives of the siginficant similarity classes, ϕ_1, ϕ_2, with the following properties.

(i) $\phi_2 = \gamma_g \phi_1$ where $g \in \text{Fix}\phi_1 \cap \text{Fix}\phi_2$

(ii) There is a function $c : \{1, 2, \ldots, n\} \to \{1, 2\}$ such that

$$c(1) = c(2)$$

$$\phi_{c(i)} < x_1, \ldots x_i > = < x_1, \ldots x_i > \qquad \text{and in fact,}$$

$$\phi_{c(i)} x_i = x_i w_{i-1} \qquad \text{where } w_{i-1} \in < x_1, \ldots x_{i-1} >$$

In addition $w_{i-1} \in \text{Fix}\phi_1 \cup \text{Fix}\phi_2$.

We can easily compute that $\text{rk}\phi_1 + \text{rk}\phi_2 = n + 1$. For if $\phi_1 x_i = x_i w_{i-1}$ then either $w_{i-1} \in \text{Fix}\phi_1$ then we have that $x_i \widehat{w_{i-1}} \bar{x}_i$ is fixed by ϕ_1. Otherwise, $\phi_1 w_{i-1} g = g w_{i-1}$ and hence ϕ_1 fixes $x_i (\widehat{w_{i-1} g}) \bar{x}_i$. A similar computation for ϕ_2 gives the formula on the ranks, on adding in the contribution from g. The w_{i-1} are also chosen so that we do not have a word $x_i u$ fixed, with $u \in < x_1, \ldots x_{i-1} >$ unless $w_{i-1} = 1$. This is a consequence of 2.1.19. We note that g lies in the subgroup $< x_1, \ldots x_j >$ where j is chosen maximally such that the function c is constant on $\{1, \ldots, j\}$. In particular $j \geq 2$.

A natural question to ask is whether we can extend the above formulation to the general case. An examination of 2.2.4 shows that we can get something similar, but with some extra complications. This arises due to the fact that in general we cannot get rid of our Type III vertices, so that there is no natural linear order to put on the basis.

However we can do the following. Given an outer automorphism of F_n maximal rank we first find a strong allowable poset for it. This supplies us with a complete set of representatives $\{\phi_1, \ldots, \phi_k\}$ of the siginificant similarity classes of Φ along with a basis which has a partial order on it. We then impose any linear order which contains said partial order. In general this is not unique unless we have no Type III vertices. Using the linear order we then write the basis as $\{x_1, \ldots, x_n\}$. Now we define two functions c, p from $\{1, 2, \ldots, n\}$ to $\{1, 2, \ldots, k\}$, the idea being that on given a basis element we get an appropriate automorphism, satisfying the following conditions.

$$c(1) = c(2)$$

$$\phi_{c(i)} x_i = x_i w_i \qquad \text{where } w_i \in < x_1, \ldots x_{i-1} >$$

$$\phi_{p(i)} < x_1, \ldots x_i > = < x_1, \ldots, x_i >$$

Define the elements g_{ij} so that $\phi_i = \gamma_{g_{ij}} \phi_j$. Then by 2.2.5 for each i we can find a j so that $w_i = v\bar{g_{ij}}$ for some $v \in \text{Fix}\phi_j$ and such that $x_i v \bar{x}_i$ is fixed by ϕ_i. By 2.2.6 for every i there must be at least one j such that $g_{ij} \in \text{Fix}\phi_i \cap \text{Fix}\phi_j$ and by considerations of rank there must be exactly $k-1$ such pairs. (The trained eye will spot a tree lurking somewhere. In the next chapter we construct Dehn twist automorphisms from maximal rank outer automorphisms. These are certain graph of group isomorphisms and the 'lurking' tree is a maximal tree in the graph of groups.)

We conclude with an example of a maximal rank outer automorphism and an allowable poset for it. Consider $\phi \in \mathrm{Aut}F_4 =< x_1, x_2, x_3, x_4 >$ defined as follows:

$$\begin{aligned}\phi x_1 &= x_1 \\ \phi x_2 &= x_2 {x_1}^2 \\ \phi x_3 &= {x_3}^{{x_1}^{-1}} = {x_3}^{(x_3 {x_1}^{-1})} \\ \phi x_4 &= (x_4 {x_3}^2)^{(x_3 {x_1}^{-1})}.\end{aligned}$$

Let $\chi = \gamma_{x_1}\phi$ and $\psi = \gamma_{x_3^{-1}}\chi$. If Φ is the outer automorphism corresponding to ϕ (thus including χ and ψ) then we claim that Φ is of maximal rank. Simply considering the abelianisation will show that ϕ, χ, ψ lie in distinct similarity classes. Also at the very least we see that $< x_1, x_2 x_1 {x_2}^{-1} > \leq \mathrm{Fix}\phi$, $< x_1, x_3 > \leq \mathrm{Fix}\chi$ and $< x_3, x_4 x_3 {x_4}^{-1} > \leq \mathrm{Fix}\psi$. Hence by 2.1.2 there is actually equality here and Φ is of maximal rank. The strong allowable poset for Φ with respect to this basis and these representatives of the significant similarity classes is as below. (We note that the fact that the allowable poset is strong follows by elementary checking and no general reason).

A Strong Allowable Poset for Φ:

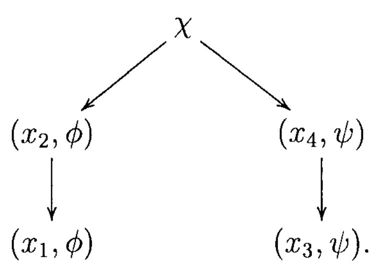

Chapter 3

Dehn Twists

In [CL] it was shown that certain graph of group automorphisms give rise to maximal rank outer automorphisms. In this section we would like to prove the converse of this result, thus starting with purely algebraic information about a free group outer automorphism we end up with a great deal of topological information.

3.1 Graphs of Groups and Dehn Twists

A graph of groups $\mathcal{G}$ consists of a graph Γ with a verex group G_v attached each vertex v of Γ, an edge group G_e attached to each edge e of Γ and monomorphisms f_e from G_e to $G_{\tau(e)}$. A dehn twist consists of a graph of groups $\mathcal{G}$ and a graph of groups automorphism $\mathcal{D}$, which is the identity on the underlying graph, the vertex groups and the edge groups and for every edge e,

$$\mathcal{D}(e) = ef_e(z_e)$$

where $f_e : G_e \to G_e$ is the graph of groups monomorphism, $z_e \in$ Centre G_e and $z_{\bar{e}} = {z_e}^{-1}$. z_e is called the twistor for the edge e.

$\mathcal{D}$ restricts to an automorphism

$$D_v :\to \pi_1(\mathcal{G}, v) \to \pi_1(\mathcal{G}, v)$$

for every vertex v of $\mathcal{G}$ and hence induces an outer automorphism.

We further require that there be a specified maximal tree in $\mathcal{G}$ which induces the isomorphisms,

$$\pi_1(\mathcal{G}, v) \xrightarrow{\delta_{vw}} \pi_1(\mathcal{G}, w)$$

$$\gamma \longmapsto \bar{\alpha}\gamma\alpha$$

where γ is a closed path at v and α is the unique path in the maximal tree from v to w.

We also require isomorphisms $\tau_v : \pi_1(\mathcal{G}, v) \to F_n$, (we are only interested in the case where the fundamental group is free) with the following commuting diagrams:

$$\begin{array}{ccc} \pi_1(\mathcal{G}, v) & \xrightarrow{\tau_v} & F_n \\ \downarrow D_v & & \downarrow \phi_v \\ \pi_1(\mathcal{G}, v) & \xrightarrow{\tau_v} & F_n \end{array}$$

so that $\phi_v, v \in V(\mathcal{G})$ all lie in the same outer automorphism class, denoted $\Phi_{\mathcal{G}}$, and we say that $\mathcal{D}$ represents $\Phi_{\mathcal{G}}$ and also that D_v represents ϕ_v. We also require that the automorphisms $\{\phi_v | v \in V(\mathcal{G})\}$ lie in different similarity classes and that some subset of these form a complete set of representatives of significant similarity classes for $\Phi_{\mathcal{G}}$. (We allow for the case where some of the ϕ_v have fixed subgroup of rank less than two.)

Our dehn twists will also be *standardised* , that is to say that $\tau_v(G_v) = \mathrm{Fix}\phi_v$ whenever $\mathrm{rk}\phi_v \geq 2$ and that the following diagrams commute for all $v, w \in V(\mathcal{G})$.

$$\begin{array}{ccc} \pi_1(\mathcal{G}, v) & \xrightarrow{\delta_{vw}} & \pi_1(\mathcal{G}, v) \\ & {\scriptstyle \tau_v} \searrow \quad \swarrow {\scriptstyle \tau_w} & \\ & F_n. & \end{array}$$

3.2 Relation to Maximal Rank

We wish to prove that any outer automorphism of maximal rank of F_n can be represented as a dehn twist. We wish to prove this by induction and by using the results of the previous section. In order to do this we need to be able to restrict our automorphism to subgroups in a controlled way.

Proposition 3.2.1 *Let P be a strong allowable poset for the outer automorphism Φ and z a vertex of P. Then $\downarrow z$ is a strong allowable poset for $\Phi|_{\langle \downarrow z \rangle}$ with basis consisting of those basis elements occurring in $\downarrow z$ and with the automorphisms occurrring in $\downarrow z$, the representatives of the significant similarity classes of $\Phi|_{\langle \downarrow z \rangle}$.*

Proof: This is just the same result as 2.1.23 but with a strong allowable poset. We note that we cannot create bad vertices by restricting to a subposet and hence the above proposition holds.

Theorem 3.2.2 *Every maximal rank outer automorphism of F_n can be represented by a dehn twist.*

Proof: We start with an outer automorphism Φ of maximal rank of F_n and find a strong allowable poset, P, for Φ. We wish to prove by induction on n that Φ is represented by a dehn twist and that also the automorphisms induced at the vertices of $\mathcal{G}$, namely ϕ_v, include those automorphisms which occur in P. The result is clear for ranks one and two, so our induction hypothesis is that the theorem holds for all ranks less than n. Let z be the top vertex of P.

case(i) z is a Type II vertex labelled by (x, ϕ).

By 3.2.1 $\downarrow z$ is a strong allowable poset for $\Phi|_{\langle\downarrow z\rangle}$. Thus we may apply the induction hypothesis to $\Phi|_{\langle\downarrow z\rangle}$.

(a) rk$\phi \geq 3$

In this case we must have that rk$\phi|\langle\downarrow z\rangle \geq 2$ and so ϕ must occur in $\downarrow z$. Hence if $\mathcal{G}$ and $\mathcal{D}$ are the graph of groups and dehn twist representing $\Phi|_{\langle\downarrow z\rangle}$, then $\phi|_{\langle\downarrow z\rangle}$ must be represented by some D_a, for some vertex a of $\mathcal{G}$ by the induction hypothesis.

Now by 2.2.5 $\phi x = xv\bar{h}$, where $\psi = \gamma_h \phi$ occurs in $\downarrow z$, fixes v and so ϕ fixes $xv\bar{x}$. As ψ occurs in $\downarrow z$ we must have that there is a vertex b of $\mathcal{G}$ such that D_b represents ψ.

We define a new dehn twist $(\mathcal{G}', \mathcal{D}')$ as follows. Change $\mathcal{G}_a$ by adding a free generator:

$$\mathcal{G}'_a = \mathcal{G}_a * < y >$$

Then we add an edge e from a to b and put $G_e \cong \mathbb{Z} \cong < c >$. Then we need the monomorphism from this new edge group to the vertex groups at its endpoints.

$$f_e c = \tau_b{}^{-1}(\hat{v}), \qquad f_{\bar{e}} c = y$$

Recall that $\hat{v}$ is the root of v and note that $\tau_b{}^{-1}(\hat{v} \in G_b$ by the conditions above and hence that in our new fundamental group $y = e\tau_b{}^{-1}(\hat{v}\bar{e}$.

Now for some positive integer m, $\hat{v}^m = v$ so let the twistor for e be c^m. (In other words $\mathcal{D}'e = ef_e(c^m)$.)

We leave the rest of the graph of groups unchanged and we let $\mathcal{D}'|_{\mathcal{G}} = \mathcal{D}$. $\mathcal{G}'$ inherits a maximal tree from $\mathcal{G}$ and by setting $\tau_a e\alpha = x$, we can extend the isomorphisms τ_u defined for $\mathcal{G}$ to isomorphisms for $\mathcal{G}'$. (Here α is the unique reduced path in the maximal tree from b to a.) We have thus defined a new standardised dehn twist and to prove the inductive step in this case it suffices to prove that ${D'}_a$ represents ϕ. This is sufficient as the maximal tree lies entirely within $\mathcal{G}$. In fact it suffices to show that ${\tau'}_a({D'}_a(e\alpha)) = xv\bar{h}$ since we already know that ${D'}_a|_{\pi_1(\mathcal{G},a)} = D_a$ which represents ϕ by the induction hypothesis.

(It is clear that the new dehn twist is standardised since the appropriate diagrams commute by definition and a calculation of the ranks of the vertex groups gives us that ${\tau'}_u(G'_u) = \mathrm{Fix}\phi_u$ for all vertices u of $\mathcal{G}'$.)

Consider the commuting diagrams:

$$\begin{array}{ccccc}
\pi_1(\mathcal{G},b) & \xrightarrow{\delta_{ba}} & \pi_1(\mathcal{G},a) & \xrightarrow{\tau_a} & F_{n-1} \\
\downarrow{\scriptstyle D_b} & & \downarrow{\scriptstyle \gamma_{(\bar{\mathcal{D}}\alpha)\alpha}\circ D_a} & & \downarrow{\scriptstyle \gamma_h\circ\phi|_{F_{n-1}} = \psi|_{F_{n-1}}} \\
\pi_1(\mathcal{G},b) & \xrightarrow{\delta_{ba}} & \pi_1(\mathcal{G},a) & \xrightarrow{\tau_a} & F_{n-1}
\end{array}$$

where $F_{n-1} = \langle \downarrow z \rangle$, α is the unique path in the maximal tree from b to a and as usual γ denotes conjugation. Note that by using 2.1.19 and 2.1.20 we know that $h \in F_{n-1}$. From the above we deduce that $h = \tau_a((\bar{\mathcal{D}}\alpha)\alpha)$.

Hence,

$$\begin{aligned}
{D'}_a(e\alpha) &= ef_e(c^m)D_a(\alpha) \\
&= e{\tau_b}^{-1}(v)\alpha{\tau_a}^{-1}(h^{-1}) \\
&= e\alpha{\tau_a}^{-1}(v)\alpha^{-1}\alpha{\tau_a}^{-1}(h^{-1}) \\
&= e\alpha{\tau_a}^{-1}(v\bar{h})
\end{aligned}$$

Thus ${\tau'}_a{D'}_a(e\alpha) = xv\bar{h}$ and we have completed the inductive step in this case.

(We note that if it were the case that $\phi x = x$ we could just add a free generator to the vertex group G_a and this would define a standardised dehn twist.)

(b) $\mathrm{rk}\phi = 2$.

If ϕ occurs in $\downarrow z$ then by 2.1.19 we must have that $n = 2$, which we have dealt with. Otherwise ϕ does not occur in $\downarrow z$ and by our induction hypothesis $\Phi|_{\langle \downarrow z \rangle}$ has a standardised dehn twist $\mathcal{D}$ on the graph of groups $\mathcal{G}$, representing it. In addition each automorphism occurring in $\downarrow z$ is represented by some D_u. Now if ϕ is represented by some D_a then we may use the same construction as in case (a). Otherwise, since we know that z is not a bad vertex we may find an automorphism ψ occurring in $\downarrow z$ and an $h \in \langle\downarrow z\rangle \cap \text{Fix}\phi \cap \text{Fix}\psi$ so that $\gamma_h\phi = \psi$. Modify $(\mathcal{G}, \mathcal{D})$ as follows: add a valence one vertex a to $\mathcal{G}$ and an edge e from a to b, where D_b is the automorphism which represents $\psi|_{\langle\downarrow z\rangle}$.

Then put

$$G_a =< y >\cong \mathbb{Z}, \quad G_e =< c >\cong \mathbb{Z}$$

$$f_e(c) = {\tau_b}^{-1}(\hat{g}), \qquad f_{\bar{e}}(c) = y$$

Note that f_e is well defined since ${\tau_b}^{-1}\text{Fix}\psi|_{\langle\downarrow z\rangle} = G_b$. Hence $y = e{\tau_b}^{-1}(\hat{g})\bar{e}$ in the new fundamental group. Also $\hat{g}$ is the root of g, so there is an intger m such that $\hat{g}^m = m$. The twistor for e is then c^m.

We extend the given maximal tree of $\mathcal{G}$ to one in $\mathcal{G}'$ by adding the edge e, and define $\mathcal{D}'|_{\mathcal{G}} = \mathcal{D}$ remembering that $\mathcal{D}'(e) = ef_e(c^m) = e{\tau_b}^{-1}(g)$. In a calculation similar to that of 3.2 we will get that D'_a represents $\gamma_{\bar{h}}\psi = \phi$ and that we have a standardised dehn twist and the completion of this case now runs as in (a).

case(ii) z is a Type III vertex labelled by ϕ. Then z covers the vertices z_1 and z_2, one of which is distinguished, say z_1. By our induction hypothesis, we have standardised dehn twists $(\mathcal{G}_1, \mathcal{D}_1)$ and $(\mathcal{G}_2, \mathcal{D}_2)$ representing $\Phi|_{\langle\downarrow z_1\rangle}$ and $\Phi|_{\langle\downarrow z_2\rangle}$ respectively. Now ϕ does not occur in $\downarrow z_2$ since a Type III vertex may not cover a Type I vertex and by use of 2.1.20. However as z is not a bad vertex there must be a $g_2 \in \langle\downarrow z_2\rangle$ and an automorphism $\psi_2 = \gamma_{g_2}\phi$ such that $g_2 \in \text{Fix}\phi \cap \text{Fix}\psi_2$. Similarly, either ϕ occurs in $\downarrow z_1$ or there is a $g_1 \in \langle\downarrow z_1\rangle \cap \text{Fix}\phi \cap \text{Fix}\psi_2$ such that $\psi_1 = \gamma_{g_1}$. In every case we can assume, by using the argument in case (i) (b) if neccessary, that $\phi|_{\langle\downarrow z_1\rangle}$ and $\phi|_{\langle\downarrow z_2\rangle}$ are represented in $\mathcal{G}_1$ and $\mathcal{G}_2$ at the vertices a_1 and a_2. We then form the graph of groups $\mathcal{G}$ by taking $\mathcal{G}_1$ and $\mathcal{G}_2$, identifying $a_1 = a_2 = a$ and putting the vertex group at a to be the free product of the vertex groups G_{a_1} and G_{a_2}. Since every closed loop at a in $\mathcal{G}$ can be written uniquely as a concatenation of closed loops in $\mathcal{G}_1$ and $\mathcal{G}_2$ at a_1 and a_2, we may define the isomorphism τ_a from $\pi_1(\mathcal{G}, a)$ to F_n by its restriction to $\pi_1(\mathcal{G}_1, a_1)$ and $\pi_1(\mathcal{G}_2, a_2)$. Taking the maximal tree in $\mathcal{G}$ to be the union, with one vertex identified, of the maximal trees in $\mathcal{G}_1$ and $\mathcal{G}_2$ we have then defined the isomorphisms τ_u at the remaining vertices thus giving us a standardised dehn twist representing Φ. This completes the induction.

We note that our proof is constructive but that the dehn twist we actually end up with is not neccessarily *efficient* as defined in [CL].

Chapter 4

Conjugacy Problems

4.1 A Combinatorial Approach

In [CT96] a normal form is given for maximal rank outer automorphisms which is very neat and admits a great deal of combinatorial information. However, this form is not actually 'normalised' in the sense that there is no discussion of when two forms give rise to conjugate automorphisms. In [CL] the conjugacy problem is solved for dehn twist automorphisms and it is noted in that paper that maximal rank automorphims are dehn twists. This gives a solution to the problem of deciding whether two normal forms represent conjugate automorphisms, but in rather a circuitous manner. We wish to give here a solution that is more direct and is also strictly combinatorial.

We start with two automorphisms (not outer automorphisms), ϕ and ψ of F_n of maximal rank. We also have strong right layered bases for these automorphisms and our goal is to recognise if there is another automorphism χ of F_n and a commuting diagram:

$$\begin{array}{ccc} F_n & \xrightarrow{\chi} & F_n \\ {\scriptstyle\phi}\downarrow & & \downarrow{\scriptstyle\psi} \\ F_n & \xrightarrow{\chi} & F_n. \end{array}$$

To remind ourselves, we say that $\{x_1, \ldots, x_n\}$ is a strong right layered basis for ϕ if

$$\phi x_i = x_i w_i$$

where for each i $w_i \in < x_1, \ldots, x_{i-1} > \cap \mathrm{Fix}\phi$ and also that w_i is cyclically reduced and that there does not exist a $u \in < x_1, \ldots, x_{i-1} >$ such that $w_i = u\phi\bar{u}$ unless $w_i = 1$. The w_i are called the right multipliers. (Technically the requirement that the multipliers be cyclically reduced is not part of the definition. However, the proof, in [CT96], of the existence of a strong right layered basis does produce cyclically reduced multipliers and we shall henceforth assume that this is part of the definition.)

We also have a strong right layered basis $\{y_1, \ldots, y_n\}$ for ψ with right multipliers v_i.

The next two lemmas are elementary and state that right layered bases are well behaved. From now on we assume that ϕ and ψ are maximal rank automorphisms with strong right layered bases as above.

Lemma 4.1.1 *Let $w \in < x_1, \ldots, x_i >$ then the number of occurences of x_i in w is equal to the number of occurences of x_i in ϕw. (As always we count these occurrences in the reduced words.)*

Proof: The lemma is clear if the number of these occurrences is 0 or 1. If the lemma is false then on applying ϕ to w there must be cancellation between an adjacent pair of x_i. For cancellation to occur this adjacent pair must be opposite in sign and hence we must have a (reduced) subword of w, $x_i v x_i^{\pm 1}$, such that $v \in < x_1, \ldots, x_{i-1} >$ in whose image one of the x_i cancel. The only way for this to happen is if $\phi v^{w_i} = 1$ which implies that $v = 1$ and so that $x_i v x_i^{\pm 1}$ is not reduced.

The next lemma is proved in [CT96]. We restate it here for convenience. Recall that $\hat{w}$ denotes the root of w.

Lemma 4.1.2 *Let $b \in \mathrm{Fix}(\gamma_{w_i}\phi) \cap \{x_1, \ldots, x_{i-1}\}$, for some $w_i \neq 1$. Then $b \in < \hat{w}_i >$.*

Proof: It is clear that $\mathrm{rk}\phi|_{<x_1,\ldots,x_{i-1}>} = i-1$, by inspection and that since w_i is fixed, $w_i \in \mathrm{Fix}(\gamma_{w_i}\phi) \cap \{x_1, \ldots, x_{i-1}\}$. By the strong form of the Scott Conjecture 2.1.5 if $\mathrm{rk}(\gamma_{w_i}\phi)|_{\{x_1,\ldots,x_{i-1}\}} \geq 2$ then $(\gamma_{w_i}\phi)|_{\{x_1,\ldots,x_{i-1}\}}$ is similar to $\phi|_{\{x_1,\ldots,x_{i-1}\}}$ and so that there is a $u \in \{x_1, \ldots, x_{i-1}\}$ such that $w_i = u\phi\bar{u}$. Hence the result.

We now prove the key ingredient of our solution. For the purposes of the next proposition we assume that ϕ and ψ are conjugate under the automorphism χ an that x_i is a member of the right layered basis for ϕ with $w_i \neq 1$. (Equivalently $x_i \notin \mathrm{Fix}\phi$.)

Proposition 4.1.3 *$\chi x_i = \alpha g \beta$ for some $\alpha, \beta \in \mathrm{Fix}\psi$ and where g is either ${y_j}^{\pm 1}$ or $g = (y_j a y_k)^{\pm 1}$ for some $k < j$ with $\hat{v}_j{}^a = \hat{v}_k{}^{\pm 1}$ and $a \in \{y_1, \ldots, y_{k-1}\}$.*

Proof: Consider the word χx_i in terms of the basis $\{y_1, \ldots, y_n\}$. We can decompose this as a *reduced* word in the following manner, for some integer j.

$$\chi x_i = \alpha t {y_j}^{\pm 1} u \beta$$

where $\alpha, \beta \in \mathrm{Fix}\psi$, $t \in < y_1, \ldots, y_{j-1} >$ and $u \in < y_1, \ldots, y_j >$. Additionally we choose α, β so that $|\alpha| + |\beta|$ is maximal. We thus consider two cases depending on whether y_j has positive or negative exponent in the above decomposition.

case(i) We have a commuting diagram as follows:

$$\begin{array}{ccc} x_i & \xrightarrow{\chi} & \alpha t y_j u \beta \\ \downarrow \phi & & \downarrow \psi \\ x_i w_i & \xrightarrow{\chi} & \alpha\psi(t) y_j v_j \psi(u)\beta = \\ & & \alpha t y_j u \beta \chi(w_i) \end{array}$$

Now as we know that $w_i \in \mathrm{Fix}\phi$ we have that $\chi(w_i) \in \mathrm{Fix}\psi$ and hence that the following word is fixed by ψ:

$$\bar{u}\bar{y}_j\bar{t}\psi(t)y_jv_j\psi(u).$$

By 4.1.1, after reduction, the word $y_jv_j\psi(u)$ must begin with y_j and thus we get that if $\bar{t}\psi(t) \neq 1$ then as $t \in \{y_1, \dots, y_{j-1}\}$ we get that $yu \in \mathrm{Fix}\psi$ which is a contradiction to the maximality of α and β. Hence $t \in \mathrm{Fix}\psi$ and so $t = 1$, again by the maximality of α and β.

Thus we get that the following word is fixed by ψ:

$$\bar{u}v_j\psi(u) \tag{4.1}$$

We wish to show now that u does not contain any occurrence of y_j. We prove this by contradiction and we split the proof into two cases depending on whether the first occurrence of y_j in u has positive or negative exponent.

(1) $u = ay_jb$ where $a \in \{y_1, \dots y_{j-1}\}, b \in \{y_1, \dots y_j\}$.

By 4.1 we know that

$$\bar{b}\bar{y}_j\bar{a}v_j\psi(a)y_jv_j\psi(b)$$

is fixed by ψ. Now if $\bar{a}v_j\psi(a) = 1$ then ya is fixed by ψ which is a contradiction to the maximalty of α. Hence $\bar{a}v_j\psi(a) \neq 1$. However, by 4.1.1, $y_jv_j\psi(b)$ begins with y_j after reduction and hence y_jb is fixed whcih is again a contradiction to the maximality of β. So this situation cannot arise and we move on to the next case.

(2) $u = ay_j{}^{-1}b$ where $a \in \{y_1, \dots y_{j-1}\}, b \in \{y_1, \dots y_j\}$.

By 4.1 we have that

$$\bar{b}\bar{y}_j\bar{a}v_j\psi(a)y_jv_j\psi(b)$$

is fixed by ψ and as above we deduce that $\bar{a}v_j\psi(a) \neq 1$ and hence that $b \in \mathrm{Fix}\psi$ which implies that $b = 1$.

Now by 4.1.2 we get that

$$\bar{a}v_j\psi(a) \in <\hat{v}_j> \leq \mathrm{Fix}\psi. \tag{4.2}$$

However consider the following,

$$\begin{aligned}\psi(v_j\psi(a)\bar{a}) &= \psi((\bar{a}v_j\psi(a))^{\bar{a}}) \\ &= (\bar{a}v_j\psi(a))^{\bar{\psi a}} \text{ by 4.2} \\ &= \psi(a)\bar{a}v_j\end{aligned}$$

Thus we see that $v_j\psi(a)\bar{a} \in \mathrm{Fix}(\gamma_{v_j}\psi)|_{\{y_1,\dots,y_{j-1}\}}$ and hence by 4.1.2 that

$$v_j\psi(a)\bar{a} \in <\hat{v}_j> .$$

We immdiately get that

$$<\hat{v}_j> \ni \bar{a}v_j\psi(a) = (v_j\psi(a)\bar{a})^a \in <\hat{v}_j>^a$$

and hence that $a \in <\hat{v}_j>$. This means that $y_ja\bar{y}_j \in \mathrm{Fix}\psi$ which is in contradiction to the maximality of α. Thus we conclude that u cannot contain any ocurrence of the letter y_j and thus that $u \in \{y_1, \dots, y_{j-1}\}$.

Now if $u = 1$ we would be done as then the image of x_i under χ would satisfy the conditions of the proposition so we assume that $u \neq 1$ and we attempt to show that there is a $k < j$ such that v_k and v_j are conjugate by a word a, and that $u = a\bar{y}_k$.

We continue in the same manner by looking at the highest letter that occurs in u and almost as before we consider separately the cases where the last occurrence of said letter has positive or negative exponent.

(1) $u = ay_kb$ where $k < j$, $a \in \{y_1, \ldots, y_k\}$ and $b \in \{y_1, \ldots, y_{k-1}\}$.

Now consider,

$$\begin{aligned} \psi(v_j\psi(u)\bar{u}) &= \psi((\bar{u}v_j\psi(u))^{\bar{u}}) \\ &= (\bar{u}v_j\psi(u))^{\psi(\bar{u})}, \text{ by 4.1} \\ &= \psi(u)\bar{u}v_j \end{aligned} \tag{4.3}$$

and hence that $v_j\psi(u)\bar{u} \in \mathrm{Fix}(\gamma_{v_j}\psi) \cap \{y_1, \ldots, y_{j-1}\}$. So that by 4.1.2 $v_j\psi(u)\bar{u} \in< \hat{v}_j >$ and in fact that $\psi(u)\bar{u} \in< \hat{v}_j >$. In terms of the decomposition of u this means that,

$$\psi(a)y_kv_k\psi(b)\bar{b}\bar{y}_k\bar{a} \in< \hat{v}_j >\leq \mathrm{Fix}\psi.$$

If $v_k\psi(b)\bar{b} = 1$ then we can easily compute that $y_kb \in \mathrm{Fix}\psi$ which is a contradiction to the maximality of β. On the other hand if $v_k\psi(b)\bar{b} \neq 1$ then by 4.1.1 $a \in \mathrm{Fix}\psi$ and hence v_j cannot be cyclically reduced. Thus u cannot have this form.

(2) $u = a\bar{y}_kb$ where $k < j$, $a \in \{y_1, \ldots, y_k\}$ and $b \in \{y_1, \ldots, y_{k-1}\}$.

As above we get that $\psi(u)\bar{u} \in< \hat{v}_j >$ and writing this in terms of the decomposition,

$$\psi(a)\bar{v}_k\bar{y}_k\psi(b)\bar{b}y_k\bar{a} \in< \hat{v}_j >\leq \mathrm{Fix}\psi.$$

If $\psi(b)\bar{b} \neq 1$ we again get that v_j is not cyclically reduced, so we conclude that $b \in \mathrm{Fix}\psi$ and so $b = 1$, by the maximality of β.

Rewriting the above equation we get,

$$\psi(a)\bar{v}_k\bar{a} \in< \hat{v}_j >\leq \mathrm{Fix}\psi. \tag{4.4}$$

We may reiterate the above arguments to show that a can have no occurrence of a letter y_m with $m \geq k$, in the following manner.

Suppose that $a = a_1 y_m b_1$ where $a_1 \in \{y_1, \dots, y_m\}$ and $b \in \{y_1, \dots, y_{m-1}\}$. Then,

$$\psi(a_1)y_m v_m \psi(b_1)\bar{v_k}\bar{b_1}\bar{y_m}\bar{a_1} \in <\hat{v_j}> \leq \text{Fix}\psi.$$

If $v_m\psi(b_1)\bar{v_k}\bar{b_1} = 1$ then $y_m b_1 \bar{y_k} \in \text{Fix}\psi$ which contradicts the maximality of β. (Remember that $b = 1$.) Thus we get that $a_1 \in \text{Fix}\psi$ and thus that v_j is not cyclically reduced. Hence a may not have this form.

If on the other hand $a = a_1\bar{y_m}b_1$ where $a_1 \in \{y_1, \dots, y_m\}$ and $b \in \{y_1, \dots, y_{m-1}\}$. Then,

$$\psi(a_1)\bar{v_m}\bar{y_m}\psi(b_1)\bar{v_k}\bar{b_1}y_m\bar{a_1} \in <\hat{v_j}> \leq \text{Fix}\psi.$$

Here we argue that if $\psi(b_1)\bar{v_k}\bar{b_1} = 1$ then $b_1\bar{y_k} \in \text{Fix}\psi$ which is in contradiction to the maximality of β, (Remember that we are still dealing with the situation where $u = a\bar{y_k}$) and that otherwise v_j is not cyclically reduced. Hence we conclude that $a \in \{y_1, \dots, y_{k-1}\}$

By repeating the calculation of 4.3, using that $av_k\bar{\psi(a)} \in \text{Fix}\psi$ and $a \in \{y_1, \dots, y_{k-1}\}$ we get that $\bar{\psi(a)}a \in <\hat{v_k}>$. However we know from 4.4 that there is an integer r with $\psi(a)\bar{v_k}\bar{a} = \hat{v_j}^r$ and so,

$$(\hat{v_j}^r)^a = \bar{a}\psi(a)\bar{v_k} \in <\hat{v_k}>$$

and hence that

$$\hat{v_j}^a = \hat{v_k}^{\pm 1}$$

This concludes case (i).

case(ii) This is similar to case (i) but with the exponent of y_j being negative. In other words our commuting diagram is now,

$$
\begin{array}{ccc}
x_i & \xrightarrow{\chi} & \alpha t\bar{y}_j u\beta \\
\downarrow \phi & & \downarrow \psi \\
x_i w_i & \xrightarrow{\chi} & \alpha\psi(t)\bar{v}_j\bar{y}_j\psi(u)\beta = \\
& & \alpha t\bar{y}_j u\beta\chi(w_i)
\end{array}
$$

where, as before, $t \in \{y_1, \ldots, y_{j-1}\}$, $u \in \{y_1, \ldots, y_j\}$ and $\alpha, \beta \in \text{Fix}\psi$. Additionally, $\alpha t\bar{y}_j u\beta$ is reduced as written and $|\alpha| + |\beta|$ is maximal.

As before we know that $w_i \in \text{Fix}\phi$ and thus that the following word is fixed by ψ:

$$\bar{u}y_j\bar{t}\psi(t)\bar{v}_j\bar{y}_j\psi(u).$$

Retreading now familiar arguments, we see that if $\bar{t}\psi(t)\bar{v}_j = 1$, then $t\bar{y}_j \in \text{Fix}\psi$ in contradiction of the maximality of α and so by use of 4.1.1 we conclude that $u \in \text{Fix}\psi$ and so $u = 1$ by the maximality of β.

By 4.1.2 we get that $\bar{t}\psi(t)\bar{v}_j \in < \hat{v}_j > \leq \text{Fix}\psi$. We also have that $tv_j\psi(\bar{t}) \in \text{Fix}\psi$ by the following calculation:

$$
\begin{aligned}
tv_j\psi(\bar{t}) &= \psi((v_j\psi(\bar{t})t)^{\bar{t}}) \\
&= (\psi(\bar{t})tv_j)^{\psi(\bar{t})} \text{ as } tv_j\psi(\bar{t}) \in < \hat{v}_j > \leq \text{Fix}\psi \\
&= tv_j\psi(\bar{t})
\end{aligned}
$$

We may then reiterate the argument of the previous section to show that either $t = 1$ or $t = y_k a$ for some $k < j$ and where $\hat{v}_j{}^a = \hat{v}_k{}^{\pm 1}$ and $a \in \{y_1, \ldots, y_{k-1}\}$. This completes the proposition.

We wish now to get rid of the unpleasantness in the previous proposition and deduce that if two maximal rank automorphisms are conjugate then the image of a non fixed basis element is, up to sign and fixed elements, another basis element. First we prove a two lemmas, the first of which was proved in [CT96] and which we reproduce here for completeness.

Lemma 4.1.4 *Let $g \in \text{Fix}\phi \cap \{x_1, \dots, x_i\}$ and suppose that $x_i \notin \text{Fix}\phi$. Then $g = u_1 u_2 \dots u_m$ where each $1 \neq u_j \in (\{x_1, \dots, x_{i-1}\} \cap \text{Fix}\phi) \cup < x_i \hat{w}_i \bar{x}_i >$ and this product is reduced as written.*

Proof: We note that this is actually a consequence of train track properties and the strong right layered basis, but we give a combinatorial proof. If $g \in \{x_1, \dots, x_{i-1}\}$ we are done. If not let u_1 be the shortest initial subword of ψ which either ends in $\bar{x}_i$ or such that the next letter of ψ is x_i. If no such nontrivial subword exists, put $u_1 = \psi$. If $u_1 \in \{x_1, \dots, x_{i-1}\}$ then by 4.1.1 we also get that $u_1 \in \text{Fix}\phi$. If $u_1 = x_i u \bar{u}$, with $u \in \{x_1, \dots, x_{i-1}\}$ then by 4.1.2 $u \in < \hat{w}_i >$. The only other possibilities for u_1 are that $u_1 = x_i u$ or $u_1 = \bar{u}\bar{x}_i$ where $u \in \{x_1, \dots, x_{i-1}\}$. However in each of these cases, by the definition of u_1 and an application of 4.1.1 we see that $x_i u \in \text{Fix}\phi$. This is turn implies that $\bar{u} w_i \psi(u) = 1$ in contradiction of our hypotheses, as we have a strong right layered basis. Hence the lemma.

Lemma 4.1.5 *If $v^a = w$, where v and w are cyclically reduced fixed words, then $a \in \text{Fix}\psi$.*

Proof: Since we know that both v and w are cyclically reduced, we must be able to write $a = v^r a'$ for some r and some a' which is a subword of v. Therefore it is enough to show that a' is fixed by ψ. In addition to the fact that a' is a subword of v, we also have that $v^{a'} \in \text{Fix}\psi$ and is cyclically reduced.

Now look at the integer i such that, $v \in \{x_1, \dots, x_i\} - \{x_1, \dots, x_{i-1}\}$.

Suppose that $x_i \in \text{Fix}\psi$. Then we may write $v = u_1 \dots u_m$, where the u_j are alternately in $< x_i >$ and $\{x_1, \dots, x_{i-1}\}$. However using 4.1.1 we see that if $u_j \in \{x_1, \dots, x_{i-1}\}$ then u_j is fixed. Then $a' = u_1 \dots u_k g_1$ where $u_{k+1} = g_1 g_2$. Hence,

$$g_2 u_{k+2} \dots u_m u_1 \dots u_k g_1 \in \text{Fix}\psi.$$

If $u_{k+1} \in < x_i >$ then we are done. Otherwise $u_{k+1} \in \{x_1, \dots, x_{i-1}\}$ and $1 \neq u_{k+2} \in < x_i >$. Thus $g_2 \in \text{Fix}\psi$ and so g_1 and a' are fixed by ψ.

Suppose that instead, $x_i \notin \mathrm{Fix}\psi$. Using 4.1.4, we may write,

$v = u_1 \ldots u_m$ where each $1 \neq u_j \in (\{x_1, \ldots, x_{i-1}\} \cap \mathrm{Fix}\phi) \cup < x_i \hat{w}_i \bar{x}_i >$ for some i.

Then $a' = u_1 \ldots u_k g_1$ where $u_{k+1} = g_1 g_2$. If $u_{k+1} \in \{x_1, \ldots, x_{i-1}\}$ then as before we get that a' is fixed. Otherwise $u_{k+1} = x_i h_1 h_2 \bar{x}_i$, where $h_1 h_2 \in < \hat{w}_i >$ and $g_1 = x_i h_1$. However as $v^{a'}$ is fixed we get a contradiction in this case, using either 4.1.1 or 4.1.4. Thus we have shown that a is fixed.

Now we have the improvement of 4.1.3,

Proposition 4.1.6 *Suppose that ϕ and ψ are conjugate under χ. If $x_i \notin \mathrm{Fix}\phi$, then for some j,*

$$\chi x_i = \alpha y_j^{\pm 1} \beta \text{ where } \alpha, \beta \in \mathrm{Fix}\psi.$$

Proof: Our first step is to show that if the proposition does not hold, then there exist i_1, i_2 such that $\chi x_{i_1} = \alpha y_j a \bar{y}_k^{\pm 1} \beta$ and $\chi x_{i_2} = \alpha' y_k^{\pm 1} \beta'$, for some $j > k$. Let Π denote the natural abelianisation map from F_n to the free abelian group of rank n. Divide the basis elements $y_1, \ldots, y_n$ into two disjoint sets $A = \{y_j | v_j = 1\}$ and $B = \{y_j | v_j \neq 1\}$. Then by 4.1.4 (and a trivial inductive argument), $\Pi(\mathrm{Fix}\psi) = \Pi(A)$. Now $\Pi(\chi(x_i))$, which we label z_i, as i goes from 1 to n is a basis for the free group of rank n as is $\Pi(y_j)$ $1 \leq j \leq n$. By 4.1.3 we have that for each i, either $z_i \in \Pi(A)$ or $z_i = \Pi(y_j) + g_i$ for $y_j \in B$ and $g_i \in \Pi(\mathrm{Fix}\psi) = \Pi(A)$ or $z_i = \Pi(y_j) - \Pi(y_k) + g_i$, where $j > k$, $y_j, y_k \in B$ and $g_i \in \Pi(\mathrm{Fix}\psi) = \Pi(A)$. This last statement is a consequence of 4.1.5.

It is now clear by inspection that we can find the required x_{i_1} and x_{i_2} with,

$$\begin{aligned} \chi(x_{i_1}) &= \alpha y_j a \bar{y}_k^{\pm 1} \beta \\ \chi(x_{i_2}) &= \alpha' y_k^{\pm 1} \beta' \end{aligned}$$

where $j > k$ and $\alpha, \beta, \alpha', \beta' \in \mathrm{Fix}\psi$. Also $a \in \mathrm{Fix}\psi \cap \{y_1, \dots y_{k-1}\}$. Now remembering that $v_j{}^a \in < \hat{v}_i >$ we have,

$$\begin{aligned} \chi(w_{i_1}) &\in < \bar{\beta}(y_k \hat{v}_k \bar{y_k})^{\pm 1} \beta > \\ \chi(x_{i_2} w_{i_2} \bar{x_{i_2}}) &\in < \alpha'(y_k \hat{v}_k \bar{y_k})^{\pm 1} \bar{\alpha}' > . \end{aligned}$$

Plainly then,

$$w_{i_1} \in < \gamma(x_{i_2} \hat{w_{i_2}} \bar{x_{i_2}})^{\pm 1} \bar{\gamma} >,$$

where $\gamma \in \mathrm{Fix}\psi$. Now if $\gamma \notin \{x_1, \dots, x_{i_2-1}\}$ then clearly w_{i_1} will not be cyclically reduced and so we get a contradiction. Thus $\gamma \in \{x_1, \dots, x_{i_2}\}$, but by 4.1.4 we know that occurences of x_{i_2} in γ must come in adjacent pairs of opposing sign. Hence in all cases, w_{i_1} will not be cyclically reduced in contradiction to the properties of a right layered basis. Thus we have proved the proposition.

Now for the main theorem,

Theorem 4.1.7 *Given two maximal rank automorphisms, ϕ and ψ with strong right layered bases, there exists an algorithm to decide whether they are conjugate in* $\mathrm{Aut}F_n$.

Proof: For convenience we assume that (using our previous notation) for some r, $w_1 = w_2 = \dots = w_r = 1 = v_1 = \dots = v_r$. Also $w_j \neq 1$ for $j > r$ (and hence $v_j \neq 1$ for $j > r$). First we consider permutations σ of $r+1, \dots, n$. For each $i > r$ we choose an $\epsilon_i = \pm 1$ and then attempt to find, using Whitehead's algorithm, an isomorphism, χ_1, from $\mathrm{Fix}\phi$ to $\mathrm{Fix}\psi$ which sends, for some $\alpha, \beta \in \mathrm{Fix}\psi$, w_i to $v_{\sigma(i)}{}^\beta$ and $x_i w_i \bar{x}_i$ to $y_{\sigma(i)} v_{\sigma(i)} \bar{y_{\sigma(i)}}{}^{\bar{\alpha}}$ if $\epsilon_i = 1$. If $\epsilon_i = -1$ then we require that χ_1 sends w_i to to $y_{\sigma(i)} \bar{v_{\sigma(i)}} \bar{y_{\sigma(i)}}{}^\beta$ and sends $x_i w_i \bar{x}_i$ to $\bar{v_{\sigma(i)}}{}^{\bar{\alpha}}$. We then define $\chi(x_i)$ to be $\chi_1(x_i)$ if $i \leq r$ and to $\alpha y_{\sigma(i)} \beta$ otherwise. If we have found χ_1 with the above properties, then ϕ and ψ are conjugate under this χ.

Conversely if ϕ and ψ are conjugate, then by 4.1.6 for each $i > r$, $\chi(x_i) = \alpha y_{\sigma(i)}{}^{\pm 1} \beta$. We can easily see that for $\chi(x_i)$ to be a basis of F_n, σ must be a permutation. We can also see that χ induces an isomorphism between $\mathrm{Fix}\phi$ and $\mathrm{Fix}\psi$, χ_1 as above. Thus by considering all the finitely many possibilities of σ and ϵ_i we will find an automorphism which conjugates ϕ to ψ. So we are done.

The attentive reader will have noticed that the above Theorem does not quite solve the conjugacy problem for maximal rank automorphisms since it assumes that we have a strong right layered basis. Given a maximal rank automorphism it is algorithmically possible to find a basis, $\{x_1, \ldots, x_n\}$ such that,

$$\begin{aligned} \phi x_1 &= x_1 \\ \phi x_2 &= x_2 w_2 \\ &\vdots \\ \phi x_n &= x_n w_n \end{aligned}$$

where $w_i \in \mathrm{Fix}\phi \cap < x_1, \ldots, x_{i-1} >$ and w_i is cyclically reduced. We know that there is such a basis, so we may search through all possible bases (if we do not care for the speed of our algorithm!) until we find one. The point being that we do not know for any w_i whether or not there exists a $u \in < x_1, \ldots, x_{i-1} >$ with $\bar{u} w_i \phi(u) = 1$. (We note that given a right layered basis, one in which the multipliers just lie in the correct subgroup and no more, we can find a basis as above. The process is descibed in [CT96] and is also the content of 2.1.19.) We would like to show that the above basis is *strongly* right layered, but this is not neccessarily the case. For suppose that for some i, $w_i = g_1 g_2$ where both g_1 and g_2 are fixed by ϕ and that for some $j > i$, $w_j = g_2 g_1$. Then the basis will fail to be strongly right layered but in a very easily correctible way. If we have that the above situation occurs (where the products are reduced as written) then we replace x_j, by $x_j{}' = x_j \bar{g}_1 \bar{x}_i \in \mathrm{Fix}\phi$. As we only look for reduced products there is clearly an algorithm to do this, so we assume that our basis has this additional property.

We now claim that the basis above is a strong right layered basis.

Proposition 4.1.8 *The basis described above is a strong right layered basis for* ϕ.

Proof: We merely need to show that for each $w_i \neq 1$ there does not exist a $u \in \{x_1, \ldots, x_{i-1}\}$ with $u\phi(\bar{u}) = w$. We prove this by induction noting that $w_1 = 1$, so the induction does start. Suppose then that $\{x_1, \ldots, x_{i-1}\}$ is a strong right layered basis for $\phi|_{<x_1,\ldots,x_{i-1}>}$, we will show that the same is true for i. If $w_i = 1$, then we are done. Suppose then that $w_i \neq 1$ and that there exists a $u \in < x_1, \ldots, x_{i-1} >$ with $u\phi(\bar{u}) = w_i$. As we have done before we look for the last occurence of the highest letter occuring in u. We additionally assume that u is the shortest amongst such words with the above propery.

case (i) $u = ax_jb$, where $a \in < x_1, \ldots, x_j > b \in < x_1, \ldots, x_{j-1} >$ and $j < i$. Then,

$$w_i = u\phi(\bar{u}) = ax_jb\phi(\bar{b})\bar{w}_j\bar{x}_j\phi(\bar{a}).$$

We note that by the induction hypotheses $b\phi(\bar{b})\bar{w}_j \neq 1$ and so by 4.1.1 (which does not actually require a strong right layered basis) we get that $a \in \text{Fix}\phi$ and that w_i is not cyclically reduced, which is a contradiction. Hence u cannot be written in this form.

case (ii) $u = a\bar{x}_jb$, where $a \in < x_1, \ldots, x_j > b \in < x_1, \ldots, x_{j-1} >$ and $j < i$. This time,

$$w_i = u\phi(\bar{u}) = a\bar{x}_jb\phi(\bar{b})x_jw_j\phi(\bar{a}).$$

Now if $b\phi(\bar{b}) \neq 1$ then w_i is not cyclically reduced, by 4.1.1. If, on the other hand $b\phi(\bar{b}) = 1$, then we deduce that $b = 1$ by the minimality of u. Hence,

$$w_i = aw_j\phi(\bar{a}).$$

Again by the minimality of u, we deduce that $w_j \neq 1$. We repeat a calculation we have done before,

$$\begin{aligned}
\phi(w_j\phi(\bar{a})a) &= \phi((aw_j\phi(\bar{a}))^{\bar{a}} \\
&= (aw_j\phi(\bar{a}))^{\phi(\bar{a})} \text{ ,as} aw_j\phi(\bar{a}) \in \text{Fix}\phi \\
&= \phi(\bar{a})aw_j.
\end{aligned}$$

Thus by 4.1.2 and our inductive hypotheses, $\phi(\bar{a})a \in <\hat{w}_j>$. Hence ${w_i}^a \in <\hat{w}_j>$ and so by 4.1.5 $a \in \text{Fix}\phi$. Thus we have that $w_i = {w_j}^{\bar{a}}$. By the minmality of u and as w_i and w_j are cyclically reduced we see that a is an initial subword of w_j. Thus $w_j = aa'$, where $a, a' \in \text{Fix}\phi$ and $w_i = a'a$. This is precisely the situation which we have avoided, as discussed in the comments before this propostion. Thus we have proved the proposition.

We have thus shown the following:

Corollary 4.1.9 *The conjugacy problem for automorphisms of maximal rank is solvable.*

As commented previously, it is shown in [CL] that the conjugacy problem for outer automorphisms of maximal rank is solvable. (Actually, there it is proven for dehn twists, but by 3.2.2 these are the same.) This, together with some partial results indicate that the above methods will work for relative train track maps which represent maximal rank outer automorphisms, or at least some prudent choice of relative train tracks.

What is interesting about our proof is that it is entirely combinatorial and so one might hope to extend the result to maximal rank automorphisms of *free products*. In [CT96], it is shown that if ϕ is a maximal rank automorphism of the free product G. Then we can write $G = G_1 * \ldots * G_n$ such that,

$\phi G_i = G_i$ if $G_i \not\cong \mathbb{Z}$

$\phi(x_i) = x_i w_i$ where $w_i \in G_1 * \ldots * G_{i-1}$, if $G_i \cong \mathbb{Z} =< x_i >$

Additionally, $\text{Fix}\gamma_{w_i}\phi|_{\cap G_1 * \ldots * G_{i-1}} \cong \mathbb{Z}$ or $\subset \bar{\lambda}_i G_j \lambda_i$

We wish to discuss the conjugacy problem for maximal rank automorphisms of free products. Firstly we must have that the conjugacy problem be solvable in each of the automorphism groups of the factor groups. For the next condition consider the following illustrative set of examples.

Let $G = H * \mathbb{Z}$ where H is indecomposable and not infinite cyclic, and $\mathbb{Z} =< x >$. Define two maximal rank automorphisms ϕ_1, ϕ_2 defined as follows. $\phi_1|_H = \phi_2|_H = identity$ and for some $h_1, h_2 \in H$, $\phi_i x = x h_i$. Suppose that there is an $a \in H$ with $h_1 = \bar{a} h_2 \phi_2(a)$, then by defining $\chi_H = identity$ and $\chi x = xa$ we see that ϕ_1 and ϕ_2 are conjuagate under χ.

We say that the *twisted* conjugacy problem is solvable in H if given $h_1, h_2 \in H$ and $\phi \in \text{Aut}H$ with $\text{Fix}\phi, \text{Fix}(\gamma_{h_1}\phi), \text{Fix}(\gamma_{h_2}\phi) \neq \{1\}$ then it is decidable whether there is an $a \in H$ with $h_1 = \bar{a} h_2 \phi(a)$.

We believe it can then be shown, using methods similar to those above, that if the conjugacy problem is solvable for maximal rank automorphisms of G then the twisted conjugacy problem is solvable in each of the factor groups. We conjecture that this is also a sufficient condition.

Chapter 5

Index Theorem

5.1 The Attractivity of Infinite Fixed Words

The following theorem is the free product analogue of the index theorem for finitely generated free groups by [GJLL]. It can be seen as a refinement of the Bestvina-Handel theorem [BH92] later extended by Collins and Turner to Free Products in [CT94].

Recall that for a subgroup H of a free product, $G = G_1 * G_2 * ... * G_k * F_r$, where the G_i are freely indecomposable groups and F_r is a finitely generated free group of rank r with basis $\{y_1, \ldots, y_r\}$ we have a notion of Kuroš rank of H in G written $K(H; G)$. Collins and Turner then prove that for $\phi \in \mathrm{Aut}G$ $K(Fix\phi; G) \leq k + r$. (1.1.2)

Our goal is to study the dynamics of an automorphism of G on its infinite reduced words, that is an infinite sequence of letters of G, $a_1 a_2 a_3 \ldots$ such that every product $a_i a_{i+1}$ is reduced as written. This means that either a_i and a_{i+1} belong to different factors or they are both basis elements in the free group and $a_i \neq {a_{i+1}}^{-1}$.

As in [Coo87] we can define a topology on G using the word metric. This space is not, however, compact although automorphisms of G still act uniformly continuously.

Given two possibly infinite words, Y and Z, we define $Y \wedge Z$ to be the longest common initial subword of Y and Z. This allows us to turn our free product G into a metric space. Namely,

$$d(Y,Z) = \begin{cases} (1+|Y \wedge Z|)^{-1} \text{ if } Y \neq Z \\ 0 \text{ otherwise} \end{cases}.$$

We set ∂G to be the set of reduced infinite fixed words and for any $H \subset G$, ∂H to be the set of $X \in \partial G$ such that for every integer N there is a $w \in H$ and $|w \wedge X| \geq N$.

Then, as in the free group case, free product automorphisms have the bounded cancellation and uniform continuity properties (see [Gol90]).

Theorem 5.1.1 (Bounded Cancellation) *Given $p \geq 0$ there is a $q \geq 0$ such that for some $x, x' \in G$, if $|x \wedge x'| \leq p$ then $|\phi x \wedge \phi x'| \leq q$. Note that we set B_ϕ to be the constant we get on setting $p = 0$ and $B = max(B_\phi, B_{\phi^{-1}})$.*

We also have the related property of uniform continuity:

Theorem 5.1.2 (Uniform Continuity) *Given $q \geq 0$ there is a $p \geq 0$ such that if $|x \wedge x'| \geq p$ then $|\phi x \wedge \phi x'| \geq q$.*

Hence it is possible to define the action of an automorphism on the infinite reduced words. Moreover, the properties of Uniform continuity and Bounded cancellation still apply when infinite reduced words are considered. If X is such a word and $\phi \in \text{Aut}G$, then ϕX is defined to be $\lim_{i \to \infty} \phi X_i$ where X_i is the initial subword of X of length i. That this limit exists is a consequence of Uniform Continuity. Additionally for any sequence of words $\{w_i\}$ such that $w_i \to X$, then $\lim_{i \to \infty} \phi w_i = \phi X$. (Again a consequence of uniform continuity.)

Suppose we are given an infinite word X which is fixed by ϕ. We shall always let X_i denote the finite initial word of X of length i and let x_i denote the ith letter of X

or equivalently the terminal letter of X_i. Now we have that $\phi X_i = X_{k(i)}.Z_i$ where the $k(i)$ is chosen maximally and by bounded cancellation we have that $|Z_i| \leq B$ where B is the bounded cancellation constant. Similarly $\phi^{-1}X_i = X_{h(i)}Y_i$ where $|Y_i| \leq B$. We put $w_i = X_i^{-1}\phi X_i$ and $v_i = X_i^{-1}\phi^{-1}X_i$.

Since X is fixed, $k(i) \to \infty$ and we say that X is algebraically attracting for ϕ if $k(i) - i \to \infty$. We say that X is algebraically repelling for ϕ if $h(i) - i \to \infty$.

We also have a notion of topological attractivity. We say that X is topologically attracting for ϕ if there exists an integer N such that if $|Y \wedge X| \geq N$ then $\lim_{n\to\infty} \phi^n Y = X$, and we call X topologically repelling if the corresponding statement is true of ϕ^{-1}.

That these notions do not in general coincide is an interesting departure from the free group case.

We may define the index of an automorphism, $i(\phi)$, to be ~~$K(Fix\phi, G)$~~ K-rank(G, Fixφ) $+ \frac{a(\phi)}{2}$ where $a(\phi)$ denotes the number of equivalence classes of topologically attracting fixed infinite words for the automorphism. Two infinite attracting reduced words, X, Y, for ϕ are said to be equivalent if there exists a (finite) word, $w \in \text{Fix}\phi$ with $wX = Y$. Then for $\Phi \in$ Out G we define $i(\Phi) = 1 + \sum(i(\phi) - 1)$ where the sum is taken over a set of representatives for the positive similarity classes of Φ. It is not a priori clear that this number is finite.

(We recall that two automorphisms ϕ, ψ of G are said to be *similar* if there is a commuting diagram:

$$\begin{array}{ccc} G & \xrightarrow{\gamma} & G \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \psi} \\ G & \xrightarrow{\gamma} & G \end{array}$$

where γ is an inner automorphism. A similarity class is an equivalence class under this relation. It is clear that the index of similar automorphisms is the same and hence a positive similarity class is one in which each automorphism has positive index.)

The above definition is in terms of topologically attracting fixed infinite words. For free products this is not the same as its algebraic counterpart.

Consider the group $G = H_1 * H_2 * \mathbb{Z}$. Where both H_1 and H_2 are isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$, and we have that $' : H_1 \to H_2$ is a fixed isomorphism between them. Let θ be the automorphism of H_1 defined by the matrix $\binom{2\,1}{3\,2}$.

Labelling the generator of the infinite cyclic factor x, we define the automorphism α of G as follows.

For all $h \in H_1$, $\alpha h = (\theta h)' \in H_2$ and $\alpha h' = h$, for all $h' \in H_2$. On the infinite cyclic factor we have $\alpha x = xg$, where $g = (1, 1) \in H_1$.

It is clear that α possesses no fixed words but the following reduced infinite word is fixed by α:

$$xg(\theta g)'\theta g(\theta^2 g)'\theta^2 g...$$

This word is neither algebraically attracting (applying α to a subword of length k always results in a subword of length $k + 1$) nor is a limit of fixed words. We note that we are using the free product length here which ignores the fact that the words $\theta^n(g)$ are arbitrarily long in terms of the *generators of* H_1 and assigns to each of them a length of one. In fact our proof will show that this is indeed always the case. It may thus be tempting to abandon the free product length for a length based entirely on the generators of the factors. We do not do this since for this to work we would first require that our groups be finitely generated (we make no such requirement) and moreover we would need the length with respect to the generators to be in some sense well behaved. For example we would need to assume that the group were hyperbolic or automatic (as in the above example). The free product length has the advantage of having the bounded cancellation and uniform continuity properties regardless of the factor groups involved. We finally note that in the above example we could have used any group in place of $\mathbb{Z} \oplus \mathbb{Z}$ as long as it had an automorphism of infinite order.

Hence the notion of topological attractivity is weaker than its algebraic counterpart but will be sufficient for our purposes. From now on when we use the term *attracting* it will be assumed that we refer to the topological notion. In this way we have the important analogue of [GJLL, proposition I.1]. However, before we prove this we need some technical lemmas.

Now it is proved in [CT88, proposition 2.4] that there are certain restrictions on the images of free product words. The following is a restatement of their result. Although the statement there concerns only finite words Y, the result is easily deduced for infinite words by Uniform continuity.

Lemma 5.1.3 *Given a non cyclic factor of G say G_i, then $\phi(G_i) = G_{i_k}^{u_i}$ for some non cyclic factor G_{i_k} and some $u_i \in G$. Suppose for some possibly infinite word Y,*

$$\phi(Y) = u_i^{-1} x Y'$$

where this product is reduced as written and $x \in G_{i_k}$.

Then either

(i) Y begins with a letter from G_i, or

(ii) there is a finite set of letters L, and a constant, r, depending only on ϕ such that x is a product of at most r letters from L.

Note that for our purposes we asumme that the set L contains all the letters which occur in the normal form of the u_i in the lemma and also all the letters which occur in the words $\phi y_i^{\pm 1}$, where $\{y_1, \ldots, y_r\}$ is the basis of F_r and that $y_i^{\pm 1} \in L$ for all i. We also assume that L is closed under taking inverses.

The following lemma deals with controlling the words Z_i.

Lemma 5.1.4 *If $k(i) > k(i-1)$, then Z_i occurs as a (terminal) subword of ϕx_i.*

Proof: Now, $\phi x_i = Z_{i-1}^{-1} x_{k(i-1)+1} \ldots x_{k(i)} Z_i$. This product is not necessarily reduced as written, however $x_{k(i-1)+1} \ldots x_{k(i)} Z_i$ is reduced and also by definition, the first letter of Z_{i-1} cannot be the same as $x_{k(i-1)+1}$ although it could be in the same factor. Hence $Z_{i-1}^{-1}x_{k(i-1)+1}...x_{k(i)}$ ends with a letter in the same factor as $x_{k(i)}$ and thus no cancellation can occur with Z_i.

Now it is convenient to deal with a well behaved sequence i_r which is defined as follows. Let i_0 be the least integer such that $k(i_0) > 0$ and inductively let $k(i_{r+1})$ be the least integer such that $k(i_{r+1}) > k(i_r)$. We immediately have that $i_r > i_{r-1}$ and $k(i_r) > k(i_r - 1)$ and so lemma 5.1.4 applies. We also have

Lemma 5.1.5 $i_{r+1} - i_r \leq C$ *for some constant* C *independent of* r.

Proof: By uniform continuity there is a constant C such that if $|w| \geq C$ then $|\phi w| > 3B$ where B is the bounded cancellation constant. Suppose that $i_{r+1} - i_r > C$, then $|\phi X_{i_{r+1}-1}| = |\phi(X_{i_r})\phi(x_{i_r+1}...x_{i_{r+1}-1})| > k(i_r) - 2B + 3B = k(i_r) + B$, using bounded cancellation and the fact that $|\phi(x_{i_r+1}...x_{i_{r+1}-1})| > 3B$. However $k(i_{r+1} - 1) \leq k(i_r)$ and hence $|Z_{i_{r+1}-1}| > B$, which is a contradiction.

Now a little more notation. Suppose we have a finite set of letters S. Then we define another finite set of letters $S^1 = S$ and inductively we define $S^{i+1} = S^iS^i$. where we mean SS to be all the products xx' where $x, x' \in S$ and in the same factor. We assume that $1 \in S$ and so $S \subset S^1$ and that our sets are closed under taking inverses.

Definition 5.1.6 We also need the following notation. If G_i is a non cyclic factor of G, then by the Kuroš subgroup theorem, the image of G_i, under an automorphism is the conjugate of another non cyclic factor. If x is a letter in a non cyclic factor of G then we write $\phi x = [\phi x]^{\mu x}$, where we have that $[\phi x]$ will be another letter in a non cyclic factor and this product is reduced as written. This decomposition is then unique.

The next two lemmas will be used later to set up an inductive step that will enable us to show that certain fixed infinite words are attracting.

Lemma 5.1.7 *Let S be a finite set of letters such that $L \subset S$. Suppose for some k, $x_k \notin S^C$. Then for some j we have that*

(i) $[\phi x_j] \notin S$.

(ii) $|\phi(X_{j-1})\mu(x_j)^{-1} \wedge X| \geq k-1$.

(iii) $|\phi(X_{j-1})\mu(x_j)^{-1}[\phi x_j] \wedge X| \geq k-1$.

(iv) If for some integer i, $k(i) \geq k$ then, $j \leq i$.

Proof: Let m be the least integer such that $k(i_m) \geq k$. Then $k(i_m - 1) \leq k(i_{m-1}) < k \leq k(i_m)$. Now,

$$\phi(x_{i_m}) = Z_{i_m-1}{}^{-1} x_{k(i_m-1)+1} \dots x_{k(i_m)} Z_{i_m}.$$

If this is reduced as written or if $k > k(i_m - 1) + 1$ then, as in lemma 5.1.4, x_k occurs in the image of x_{i_m} and hence as $L \subset S^C$ we have that x_{i_m} is a letter from a non cyclic factor, that $x_k = [\phi x_{i_m}]$ and that conditions (ii) and (iii) in the statement of the theorem hold. So suppose that neither of these is the case. Then $k = k(i_m - 1) + 1 = k(i_{m-1}) + 1$ and we can write $Z_{i_m-1} = zZ'$ where this product is reduced and z is in the same non cyclic factor as $x_k = x_{k(i_m-1)+1}$, but not equal to it.

Thus $(z^{-1}x_k)$ occurs in the image of ϕx_{i_m} and we have two possibilities:

(a) $(z^{-1}x_k) \notin S^{C-1}$ or,

(b) $(z^{-1}x_k) \in S^{C-1}$ which implies that $z^{-1} \notin S^{C-1}$ since we know that $x_k \notin S^C$.

If (a) holds then we are done with $j = i_m$. Condition (i) is immediate and conditions (ii) and (iii) follow on noting that $\phi(X_{i_m-1})\mu(x_{i_m})^{-1} = X_{k(i_m-1)} Z_{i_m-1} Z'^{-1} = X_{k-1}z$ and similarly $\phi(X_{i_m-1})\mu(x_{i_m})^{-1}[\phi x_{i_m}] = X_{k(i_m-1)} Z_{i_m-1} Z'^{-1} z^{-1} x_k = X_k$.

So we suppose that (b) holds and so $z^{-1} \notin S^{C-1}$. Now if $k(i_m - 1) > k(i_m - 2)$ then by lemma 5.1.4, z occurs in the image of x_{i_m-1} and so much as before we get that conditions (i)-(iii) hold with $j = i_m - 1$. Otherwise $k(i_m - 1) \leq k(i_m - 2)$. Note that in this case we must have that $i_m - 1 > i_{m-1}$ since as noted $k(i_{m-1}) > k(i_{m-1}-1)$, by the comments after 5.1.4 . However we claim that here $k(i_m - 1) = k(i_m - 2)$ since if $k(i_m - 1) < k(i_m - 2)$ then $k - 1 = k(i_m - 1) < k(i_m - 2) \leq k(i_{m-1}) = k - 1$.

To review, $k(i_m - 1) = k(i_m - 2) = k - 1 = k(i_{m-1})$ and $i_{m-1} \leq i_m - 2$. The argument then proceeds as before. If z occurs in the image of x_{i_m-1} then we are done. Otherwise we must have that $Z_{i_m-2} = z'Z''$ where z and z' are in the same factor. If $z = z'$ then we shift attention to x_{i_m-2} remembering that $z = z' \notin S^{C-1} \supset S^{C-2}$. If $z \neq z'$ then $z'^{-1}z$ occurs in the image of x_{i_m-1} (since $\phi x_{i_m-1} = Z''^{-1}z'^{-1}zZ'$). As before we get two possibilities:

(a) $z'^{-1}z \notin S^{C-2}$ in which case we are done with $j = i_m - 1$, or

(b) $z'^{-1}z \in S^{C-2}$ and hence $z'^{-1} \notin S^{C-2}$

It is clear that we may repeat this argument, at each stage lowering the candidate for our possible j and reducing the index, l of our set S^l. However we always reach a positive conclusion if at some stage we get that $k(j) > k(j-1)$ but we know that $k(i_{m-1}) > k(i_{m-1}-1)$ and hence we reach a positive conclusion after at most C steps. This is also why our sets S^l are always well defined.

The next lemma is central in describing how little cancellation occurs under certain assumptions and will be key in determining attractivity properties. First we need to make the definition that ensures the 'small cancellation'.

Let x be a letter of G. We say that x has property P if the following conditions hold.

a) $x \notin L$

b) If $y, y' \in L^r$ and both are in the same factor as $[\phi x]$ then $y[\phi x]y' \notin L^r$. Or equivalently $[\phi x] \notin L^{3r}$

c) Again, if $y, y' \in L^r$ and both are in the same factor as $[\phi x]$, then $[\phi(y[\phi x]y')] \notin L^{2r}$.

(Note that these are the same L and r as in lemma 5.1.4.)

Lemma 5.1.8 *Suppose for some j that x_j has property P and that there is an $i \geq j$, such that*

$|\phi(X_{j-1})\mu(x_j)^{-1} \wedge X| \geq i$, *and* $|\phi(X_{j-1})\mu(x_j)^{-1}[\phi x_j] \wedge X| \geq i$

Then there is a $j_1 > i$ such that

(i) $|\phi(X_j Z) \wedge X| \geq j_1 - 1$, and $|\phi^2(X_j z) \wedge X| \geq j_1$, whenever Z is a possibly infinite word such that $X_j Z$ is reduced.

(ii) $|\phi(X_{j_1-1})\mu(x_{j_1})^{-1} \wedge X| \geq j_1$, and $|\phi(X_{j_1-1})\mu(x_{j_1})^{-1}[\phi x_{j_1}] \wedge X| \geq j_1$,

(iii) There are $x, x' \in L^r$ such that $x_{j_1} = x[\phi x_j]x'$.

Proof: Let Z be a possibly infinite word such that $X_j Z$ is reduced as written. Consider $\phi(X_j Z) = \phi(X_{j-1})\mu(x_j)^{-1}[\phi x_j]\mu(x_j)\phi(Z)$.

Now by lemma 5.1.3 and property P, $[\phi x_j]$ is not entirely cancelled in this product and so we can find $y, y' \in L^r$ (possibly trivial) so that

$\phi(X_j Z) = X_j w(y[\phi x_j]y')Z'$, for some words w and Z' so that this product is reduced (if we count $(y[\phi x_j]y')$ as a single letter). Note that X_j occurs as an initial subword by our hypotheses and in fact $|X_j w \wedge X| \geq i$ since $[\phi x_j]$ does not entirely cancel. Also keep in mind that y is independent of Z and that, $\phi(X_{j-1})\mu(x_j)^{-1}[\phi x_j] = X_j w(y[\phi x_j])$. Now we apply ϕ again,

$$\begin{aligned}\phi^2(X_j Z) &= \phi(X_{j-1})\mu(x_j)^{-1}[\phi x_j]\mu(x_j)\phi(w)\phi(y[\phi x_j]y')\phi(Z') \\ &= X_j w(y[\phi x_j])\mu(x_j)\phi(w)\phi(y[\phi x_j]y')\phi(Z')\end{aligned}$$

Applying lemma 5.1.3 again we can find a $y'' \in L^r$ such that

$$X_j w(y[\phi x_j])\mu(x_j)\phi(w) = X_j w(y[\phi x_j]y'')w'$$

for some w', where this product is reduced on counting the bracketed term as a single letter. Note that property P ensures that the bracketed term is non trivial. Now we observe that if $w' = \mu(y[\phi x_j]y')$ and $(y[\phi x_j]y'')$ and $[\phi(y[\phi x_j]y')]$ are in the same factor then by lemma 5.1.4 $(y[\phi x_j]y'') \in L^r$ which contradicts part (b) of property P and hence:

$$\phi^2(X_j) = X_j w(y[\phi x_j]y''')w_2(z[\phi(y[\phi x_j]y'])\mu((y[\phi x_j]y')$$

where either $w_2 \neq 1$ or $[\phi x_j]$ and $[\phi^2 x_j]$ are in different factors and otherwise the product is reduced as written with $y''', z \in L^r$.

Yet another use of lemma 5.1.3 yields that,

$$\phi^2(X_j Z) = X_j w(y[\phi x_j]y''')w_2(z[\phi(y[\phi x_j]y']z')Z''$$

where the bracketed terms are non trivial, by parts (b) and (c) of property P respectively, and they are either in different factors or $w_2 \neq 1$. We also have from lemma 5.1.3 that $z' \in L^r$ and otherwise this product is reduced.

Now we claim that the required x_{j_1} is just $(y[\phi x_j]y''')$. Firstly it is clear that this does not depend on Z. Then observe that $X_j w(y[\phi x_j]y''')$ is always a subword of $\phi^2(X_j Z)$ and hence a subword of $\phi^2(X) = X$, so we do have that $(y[\phi x_j]y''') = x_{j_1}$ for some $j_1 > i$. (Our hypotheses guarantee that $|X_j w \wedge X| \geq i$.)

The rest of the lemma follows on noting that

$$\phi(X_{j_1-1})\mu(x_{j_1})^{-1} = \phi(X_j w)\mu(x_{j_1})^{-1} = X_j w x_{j_1} w_2 z'$$

where in this last product either $w_2 \neq 1$ or x_{j_1} and z' are in different factors and is otherwise reduced and

$$\phi(X_{j_1-1})\mu(x_{j_1})^{-1}[\phi x_{j_1}] = X_j w x_{j_1} w_2(z'[\phi x_{j_1}]).$$

The following proposition is central to the discussion, however the proof is trivial.

Proposition 5.1.9 *There is a finite set of letters M such that if $x \notin M$ then $x \in G$ satisfies property P*

Let $M_0 = M$ and inductively $M_{i+1} = \phi^{i+1}LM_i\phi^{i+1}L$, where ϕL denotes the finite set of letters that occur in the reduced images of letters in L and much as before, $\phi^{i+1}LM_i\phi^{i+1}L$ is the set of all products xmx' where $x, x' \in \phi^{i+1}L$ and $m \in M_i$ and we assume that $1 \in M$ that $M \subset M_i$ for all i and that these sets are closed under taking inverses. Let N_i be the finite set of letters such that $[\phi^i y] \in M_i$ implies that $y \in N_i$. That is, N_i is the set of preimages of M_i under ϕ^i. Hence if $x \notin N_i$ then $\phi^i x \notin M_i$. Now we say that a letter z_1 is a *descendent* of z if there are $y, y' \in L$ such that $z_1 = y[\phi z]y'$, and in general that z_k is a k-descendent of z if there are letters $z = z_0, z_1, ...z_k$ such that z_{i+1} is a descendent of z_i. We have thus defined N_i so that if $x \notin N_i$ then all i-descendents of x have property P. Thus we get the following corollary to proposition 5.1.9.

Corollary 5.1.10 *There is a finite set N such that if $x \notin N$ then all k-descendents of x have property P for $1 \leq k \leq C+1$. ($C+1$ is the constant from lemma 5.1.5)*

Lemma 5.1.11 *Suppose that for some i, $k(i) \geq i$ and $k(i) > k(i-1)$ and also that some letter of w_i is not a member of N^{C+1}. Then there exist integers $i = i_0 < i_1 < ... < i_{C+1}$ such that*

$$|\phi(X_{i_k}Z) \wedge X| \geq i_{k+1} - 1$$

and,

$$|\phi^2(X_{i_k}Z) \wedge X| \geq i_{k+1}$$

for $0 \leq k \leq C$ whenever $X_{i_k}Z$ is reduced.

Proof: Now lemma 5.1.7 along with lemma 5.1.4 (if for some letter z of Z_i, $z \notin N$ then $[\phi x_i] = z \notin N$) provide us with the hypotheses necessary in order to apply lemma 5.1.8, given our choice of the set N and corollary 5.1.10, but in fact we can apply lemma 5.1.8 C times in order to reach the conclusion of the lemma.

So far we have given conditions to help show attractivity given we can exclude elements from a finite set. The next proposition shows how we can obtain this exclusion and is exactly the same as in [GJLL, proposition I.1].

Proposition 5.1.12 *Suppose that there is an integer B and a finite set of letters A such that for infinitely many w_i, $|w_i| \leq B$ and every letter of w_i is a member of A. Then $X \in \partial \text{Fix}\phi$.*

Proof: There are only finitely many words with the above properties and hence we have a sequence w_{j_k} which are all the same. Thus for every k, ${X_{j_0}}^{-1}\phi(X_{j_0}) = {X_{j_k}}^{-1}\phi(X_{j_k})$ and so $X_{j_k}{X_{j_0}}^{-1} \in \text{Fix}\phi$ and as j_0 is fixed $X_{j_k}{X_{j_0}}^{-1} \to X$.

Our final lemma before the main proposition of this section lets us deal with the case contrary to lemma 5.1.11 when $k(i) < i$.

Lemma 5.1.13 *Suppose that $k(i-1) < k(i) < i$ and that $h(k(i)) < k(i)$, then there is a finite set of letters A such that either every letter of $v_{k(i)}$ belongs to A or every letter of w_i belongs to A.*

Proof: Now $v_{k(i)} = {x_{k(i)}}^{-1}...{x_{h(k(i)+1}}^{-1}Y_{h(k(i))}$. By definition of Y_i we must have that $v_{k(i)}$ begins with a letter in the same factor as $x_{k(i)}$ and hence ${x_i}^{-1}...{x_{k(i)+1}}^{-1}v_{k(i)}$ is reduced as written. However, $\phi X_i = X_{k(i)}Z_i$ and $\phi^{-1}X_{k(i)} = X_{h(k(i))}Y_{h(k(i))}$ and so

$$\phi^{-1}{Z_i}^{-1} = {x_i}^{-1}...{x_{k(i)+1}}^{-1}v_{k(i)}$$

- Now by lemma 5.1.4 Z_i is a terminal subword of ϕx_i. Now either $x_i \in S$ or it belongs to a non cyclic factor of G, but in either case at most one letter of Z_i does not belong to L and this letter will be in a non cyclic factor. Then by enlarging our set suitably we see that we have a finite set of letters A' such that every letter of $\phi^{-1}{Z_i}^{-1}$ belongs to A' with the exception of at most one. Hence either every letter of $v_{k(i)}$ belongs to A' or each of $x_{k(i)+1}, ..., x_i$ belongs to A'. In the latter case there are only finitely many possibilities for Z_i and thus by enlarging our set once more, to a set A, we will have that every letter of w_i belongs to A.

The main proposition is now relatively easy to prove.

Proposition 5.1.14 *The following are equivalent:*

(1) X is a topologically attracting or repelling infinite fixed word for ϕ

(2) $X \notin \partial\mathrm{Fix}\phi$.

Proof: That (1) implies (2) is clear as $X \in \partial\mathrm{Fix}\phi$ means that for any $N \in \mathbb{N}$ there must exist a $Y \in \mathrm{Fix}\phi$ such that $|Y \wedge X| \geq N$ and so as $\lim_{n\to\infty} \phi^n Y = \lim_{n\to\infty} \phi^{-n} Y = Y$ we cannot satisfy (1).

The implication (2) implies (1) is quite different from the free group case since we have that in general there may be infinitely many words of any finite length.

So we are given an infinite fixed word X and we know that $X \notin \partial\mathrm{Fix}\phi$. We first fix a j_0, using proposition 5.1.12, such that for all $i > j_0$, if $|w_i| \leq 2B + C$ then some letter of w_i does not belong to N^{C+1}.

We then look for some $i > j_0 + C =$ such that $k(i) \geq i$. If there is such an i we claim that our word is attracting. Assume there is such an i and choose the greatest r such that $i_r \leq i$. Then we know that $i_r > j_0$ and that $k(i_r) \geq k(i) \geq i \geq i_r$. If $|w_{i_r}| > 2B + C$ then $k(i_r) - i_r > B + C$ and $|Y \wedge X| \geq i_r$ implies that $|\phi(Y) \wedge X| \geq i_r + C \geq i_{r+1}$ by bounded cancellation and lemma 5.1.5. Thus we would also get that $k(i_{r+1}) \geq i_{r+1}$ and this would form an inductive step. If however $|w_{i_r}| \leq 2B + C$ then by lemma 5.1.11 there are integers $i_r = m_0 < m_1 < ... < m_{C+1}$ such that if $|Y \wedge X| \geq m_s$ then $|\phi Y \wedge X| \geq m_{s+1} - 1$ and $|\phi^2 Y \wedge X| \geq m_{s+1}$. Thus if

$$|Y \wedge X| \geq i_r$$

then

$$|\phi^{2C} Y \wedge X| \geq i_r + C \geq i_{r+1}.$$

Also here we have that $k(i_{r+1}) \geq i_{r+1}$ since we may find an s with $m_s \leq i_{r+1} < m_{s+1}$ and so $k(i_{r+1}) \geq m_{s+1} - 1 \geq i_{r+1}$, using lemma 5.1.11.

We can proceed by induction to show that X is a topologicaly attracting word for ϕ.

If then we cannot show that X is attracting or repelling by this method then for some j' we have that for all $i \geq j'$, $k(i) < i$ and $h(i) < i$. By uniform continuity we may choose j' large enough so that for all $i > j'$ we also have that $h(k(i)) < k(i)$. We may also find a constant D such that if $|w| \leq B$ then $|\phi^{-1}w| \leq D$ (Remember that B is a cancellation constant for both ϕ and ϕ^{-1}.) Suppose for some $i > j'$ that $|v_{k(i)}| > B+D$. Then $k(i)-h(k(i)) > D$ and hence $|\phi^{-1}(X_{k(i)}Z_i) \wedge X| \leq h(k(i))+D < k(i) < i$ which is a contradiction. Thus $|v_{k(i)}|$ is bounded for all $i > j'$. Similarly, if $|w_i| > B + D$ then $i - k(i) > D$ and $|\phi^{-1}(X_{k(i)}Z_i) \wedge X| \leq h(k(i) + D < k(i) + D < i$. Hence $|w_i|$ is also bounded for all $i > j'$. Then by lemma 5.1.13 and proposition 5.1.12 we get that $X \in \partial Fix\phi$ which contradicts our hypothesis.

Thus for arbitrarily large i we have that either $i \leq k(i)$ or $i \leq h(i)$. Note that since an infinite word cannot be both attracting and repelling only one of these can occur.

Corollary 5.1.15 *Let X be an attracting fixed infinite word. Then there is a finite subword X' of X such that $\lim_{n\to\infty} \phi^n X' = X$ and if $|Y \wedge X| \geq |X'|$ then $|\phi Y \wedge X| \geq |X'|$. In particular X' is a subword of $\phi X'$.*

Proof We simply apply the techniques of proposition 5.1.14 noting that for some i we must have that $k(i) \geq i$ since X is attracting.

Corollary 5.1.16 *Let X be an attracting fixed infinite word, let S be a finite set of letters and let N be an integer. Then there is an integer i_0 such that for all $i \geq i_0$ either* *(i)* $k(i) \geq i + N$, *or*

(ii) There is a $j > i$ *and an element* $g \in G$ *such that both* x_j *and* g *lie in the same non free factor and so that* $X_{j-1}(x_j g) \leq \phi(X_i)$ *and* $x_j g \notin S$.

Proof This is clear on applying the argument of 5.1.14 and on noting the details of 5.1.11 and 5.1.8.

5.2 Main Theorem

We wish to consider a slightly different object from a free product, but which nevertheless holds many of the same properties. In this section we consider topological maps on graphs of complexes whose fundamenatl group is a finite free product as above and which induce an automorphism on the fundamental group. From now on we assume that all our topological maps are homotopy equivalences. In particular relative train track maps are such. In effect we are studying the fundamental groupoid of the graph of groups and we wish to state our results in terms of paths which are not necessarily closed.

So suppose that the we have a topological map, f then one still has a notion of path length and maximal common initial paths of paths which start from some common vertex v. We assign to each real edge and stem a length of one and to each element of one of the complexes we also assign a length of one. We write, $|p|$ for the length of the path p and $p \wedge q$ for the maximal common initial subpath of p and q when they start at the same vertex. We may also consider infinite reduced *paths* - an infinite sequence $x_0 x_1 ...$, where each x_i is either a real edge or a stem or an an element of a complex and additionally each pair $x_i x_{i+1}$ is reduced as a *path.* Note that an infinite path corresponds precisely to an infinite closed path in a finite graph, but we distinguish between the two as it is important to take possibly non closed subpaths of the former. We may deduce the properties of Bounded Cancellation and

Uniform Continuity for topological maps from the corresponding property for the automorphism they induce.

Proposition 5.2.1 *Let f be a topological homotopy self equivalence of a graph of complexes $\mathcal{G}$. Let v be a vertex of $\mathcal{G}$ fixed by f and p, q be paths in $\mathcal{G}$ starting at the vertex v and so that the product pq is defined. (p is neccessarily closed.) Then there is a constant B such that,*

$$|f(pq)| \geq |p| + |q| - B.$$

Proof: First suppose that both p and q are closed. Then they both induce elements of $\pi_1(\mathcal{G}, v)$ and we have the bounded cancellation constant, $B_{\pi_1(f)}$ for the automorphism induced by f. As $\mathcal{G}$ is finite, there is a constant M, such that any path, (not neccessarily closed) of length M in $\mathcal{G}$ gives rise to a word of length at least 1 in $\pi_1(\mathcal{G})$. The result then holds in this case with $B = MB_{\pi_1(f)}$. If q is not closed we may find a path s such that qs is closed and there is a constant C which does not depend on q such that $|f(s)| \leq C$. Then,

$$\begin{aligned} |f(pq)| + C \geq |f(pqs)| &\geq |f(p)| + |f(qs)| - MB_{\pi_1(f)} \\ &\geq |f(p)| + |f(q)| - MB_{\pi_1(f)} - C. \end{aligned}$$

So the result holds with $B = MB_{\pi_1(f)} + 2C$.

Proposition 5.2.2 *Let p, q be paths starting at the vertex v which is fixed by f. Given an integer N, there exists an intger K_N (not depending on p and q) such that whenever $|p \wedge q| \geq K_N$ then $|f(p) \wedge f(q)| \geq N$.*

Proof: Let C and M be the constants from the proof of the above proposition. Let s_p, s_q be paths in $\mathcal{G}$ so that ps_p and qs_q are closed. By possibly increasing C, we may assume that both s_p and s_q are paths of length at most C. (C will still only depend on f and $\mathcal{G}$.) To avoid confusion we shall use the notation, $|\gamma|_*$, to denote the *word* length of γ when it is a closed path in $\mathcal{G}$ but considered as an element of the fundamental group. Then by bounded cancellation for autmorphisms we know that there exists a constant K_1 such that if $|ps_p \wedge qs_q|_* \geq K_1$ then $|f(ps_p) \wedge f(qs_q)|_* \geq M(N+C)$. Then if $|p \wedge q| \geq M(K_1 + C)$ we get that $|ps_p \wedge qs_q|_* \geq K_1$. However when $|f(ps_p) \wedge f(qs_q)|_* \geq M(N+C)$ we deduce that $|f(p) \wedge f(q)| \geq N$. Thus the proposition holds with $K_N = M(K_1 + C)$.

Thus if v is a vertex fixed by f, then f acts on the infinite reduced paths starting at v, by arguments that are the same as for the case of an automorphism. An infinite reduced word X which is fixed by f is called topologically attracting if there exists an integer i such that whenever Y is a path starting at the same vertex as X with $|Y \wedge X| \geq i$ then $\lim_{n\to\infty} \phi^n Y = X$. So the notion of topologically attracting fixed infinite path is well defined but one must remember that it refers to subpaths and not only subwords. One may as in the group case define a topology on the set of finite and infinite paths and as there, if H is a subset of paths of $\mathcal{G}$, we use the notation ∂H to denote the infinite paths which are limits of paths ~~in~~ generated by H.

It is easy to see that we also have an analogue of 5.1.3 and hence the whole argument of the previous section may be used in this context and so we have an analogue of proposition 5.1.14 and of corollary 5.1.15.

Proposition 5.2.3 *Let f be a topological ~~map of groups~~ homotopy self-equivalence of $\mathcal{G}$ where all the factor groups are indecomposable and let v be a fixed vertex of f. Suppose that X is an infinite fixed path starting at v, then the following are equivalent:*

(1) X is a topologically attracting or repelling fixed infinite path.

(2) $X \notin \partial \mathrm{Fix} f$

We also have the analogue of corollary 5.1.15

Proposition 5.2.4 *Let X be an attracting fixed infinite path. Then there is a finite subpath X' of X such that* $\lim_{n\to\infty} f^n X' = X$ *and if Y is a path with $|Y \wedge X| \geq |X'|$ then* $\lim_{n\to\infty} f^n(Y) = X$ *and $|f(Y) \wedge X| \geq |X'|$. In particular X' is a subpath of X.*

Now suppose that we have a finite graph of complexes, $\mathcal{G}$ as above and that ϕ is an automorphism of the free product G which is isomorphic to the fundamental group of this graph of complexes. Further suppose that f is a topological self homotopy equivalence of $\mathcal{G}$,v is a vertex of $\mathcal{G}$, μ is a path from v to $f(v)$ and τ is an isomorphism between G and $\pi_1(\mathcal{G}, v)$ such that the following diagram commutes:

$$\begin{array}{ccc} G & \xrightarrow{\tau} & \pi_1(\mathcal{G}, v) \\ \downarrow \phi & & \downarrow \pi_1(f,\mu) \\ G & \xrightarrow{\tau} & \pi_1(\mathcal{G}, v) \end{array}$$

Let X be an attracting fixed infinite word for ϕ and let γ be the reduced path that X traces out in $\mathcal{G}$. Let γ_i denote the initial subpath of γ of length i.We then have the following:

Lemma 5.2.5 *There is an m_0 such that for all $m \geq m_0$, $f(\gamma_m) = (\mu\gamma_m)Y$ where after reduction of $(\mu\gamma_m)$, this product is reduced and $Y \neq 1$.*

Proof: Now $\phi X_i = X_{k(i)} Z_i$ and we know from 5.1.16 that for any choice of integer N and any finite set $S \subset G$ there is an i_0 such that for all $i \geq i_0$ either $k(i) \geq i + N$ or $k(i) \geq i$ or there is a $j > i$ such that x_j lies in a factor group and $X_{j-1}(x_j g) \leq \phi(X_j)$ where g is in the same factor group as x_j and $(x_j g) \notin S$. We assume that a maximal tree in $\mathcal{G}$ has been chosen. Then given a γ_m there is a unique path, p in the maximal tree so that $\gamma_m p$ is closed. Thus $\gamma_m p$ corresponds to a subword of X when regarded as a group element. We may choose an m_0 such that whenever $m \geq m_0$ then this subword of X has length at least i_0 (as a word). Given $m \geq m_0$, let i be the length of the subword corresponding to $\gamma_m p$. Now if N is large enough and $k(i) \geq i + N$ then since we know that

$$\mu f(\gamma_m p)\mu^{-1} \simeq X_{k(i)} Z_i$$

we must get that not only is γ_m a subpath of $X_{k(i)}Z_i$ (once we reduce it as a path), but also that γ_m is a subpath of $X_{k(i)}Z_i\mu f(p^{-1})$ as there are only finitely many choices for p. Suppose on the other hand that for some $j > i$,

$$\phi X_i = X_{j-1}(x_j g)Z \simeq \mu f(\gamma_m p)\mu^{-1}$$

as above, where x_j and g lie in the same vertex group and $(x_j g) \notin S$, where S is some finite set. Then if we write this a a reduced path for a start we must have that γ_m is a subpath of $X_{j-1}(x_j g)Z$ and in fact X_{j-1} corresponds to a path γ_n, where $n \geq m$ which does not get cancelled. In addition $(x_j g)$ will correspond to an element in a vertex group and will also not get cancelled. If we choose our set, S to be large enough then we will guarantee that in the product $(x_j g)Z\mu f(p^{-1})$ the element corresponding to $(x_j g)$ in the vertex group will not get cancelled when we reduce this as a free product.

Thus by a careful choice of N and S and a use of 5.1.16 we may find an m_0 such that for all $m \geq m_0$, $\gamma_m \leq \mu f(\gamma_m)$. We conclude the proof by noting that if $\gamma_m = \mu f(\gamma_m)$ for infinitely many m then γ and hence X may be written as the limit of of sequence of fixed paths (words) for f (ϕ).

Our starting point is thus representing our free product automorphism as a relative train track map. We call the graph of complexes $\mathcal{G}$ and the homotopy equivalence f. We then show that if the index is positive then our automorphism is point represented. That is, there is some fixed vertex $v \in \mathcal{G}$ and a homotopy equivalence τ such that the following diagram commutes.

$$\begin{array}{ccc} G & \xrightarrow{\tau} & \pi_1(\mathcal{G},v) \\ \downarrow \phi & & \downarrow \pi_1(f,v) \\ G & \xrightarrow{\tau} & \pi_1(\mathcal{G},v) \end{array}$$

Note that $\pi_1(f,v)$ denotes the point induced automorphism.

We then show that any attracting infinite path is equivalent to one which is generated by a single edge and hence using an euler characteristic argument we obtain the theorem.

Our first stage then is to show that any automorphism of positive index is point represented. As in [CT94] we can always represent our automorphism by $\pi_1(f,\mu)$ for some path μ. We pick a vertex v in $\mathcal{G}$ and a lift $\tilde{v}$ of v in the universal cover $\tilde{\mathcal{G}}$ of G. (We note that we can think of our graph of groups topologically, the universal cover is then the topological one, that is a 2 - complex. This universal cover is obtained by first finding the universal cover of the underlying graph and then to each vertex which is a lift of a factor group we attach in the universal cover the Cayley complex of the respective group. The Cayley complex just being the Cayley graph of a group with 2 - cells added for every relation in the group.) We call the lifts of vertex groups *countries* as in [CT94] . When referring to edges and vertices in $\tilde{\mathcal{G}}$ we shall always mean edges and vertices which do not lie in any country.

We let $\eta : \tilde{\mathcal{G}} \to \mathcal{G}$ be the covering map and if $\tilde{u}$ and $\tilde{w}$ are points of $\tilde{\mathcal{G}}$ then $[\tilde{u},\tilde{w}]$ be any path in $\tilde{\mathcal{G}}$ from $\tilde{u}$ to $\tilde{w}$

We identify the covering translation group $\mathbf{T}$ with $\pi_1(\mathcal{G},v)$ in the natural way. (An element $t \in \mathbf{T}$ is identified with the group element $\eta[\tilde{v},t(\tilde{v})]$.) We then define $\zeta_{\tilde{f}} : \mathbf{T} \to \mathbf{T}$. So that the following diagram commutes.

$$\begin{array}{ccc} \tilde{\mathcal{G}} & \xrightarrow{\tilde{f}} & \tilde{\mathcal{G}} \\ \downarrow{\scriptstyle t} & & \downarrow{\scriptstyle \zeta_{\tilde{f}}(t)} \\ \tilde{\mathcal{G}} & \xrightarrow{\tilde{f}} & \tilde{\mathcal{G}} \end{array}$$

If we choose a lift $\tilde{v}$ of v and the lift of f with $\mu = \eta[\tilde{v}, \tilde{f}(\tilde{v})]$, then $\zeta_{\tilde{f}}$ is conjugate to $\pi_1(f, \mu)$ under the natural isomorphism.

$$\begin{array}{ccc} \mathbf{T} & \longrightarrow & \pi_1(\mathcal{G}, v) \\ \downarrow{\scriptstyle \zeta_{\tilde{f}}} & & \downarrow{\scriptstyle \pi_1(f,\mu)} \\ \mathbf{T} & \longrightarrow & \pi_1(\mathcal{G}, v) \end{array}$$

Then as in [CT94, Proposition 3.2] ϕ is point represented by f if and only if $\tilde{f}$ has a fixed point.

We thus argue that if $\tilde{f}$ is fixed point free (that is having no fixed points outside countries) then $\pi_1(f, \mu)$ cannot have positive index. As in [CT94, Proposition 3.3] we can define an orientation on edges $\tilde{E}$ in $\tilde{\mathcal{G}}$ by choosing an interior point $\tilde{w} \in \tilde{E}$ and using the direction of the first partial edge in the reduced path from $\tilde{w}$ to $\tilde{f}(\tilde{w})$. This is well defined if $\tilde{f}$ is fixed point free. The following lemma is true even if we drop the hypothesis that $\tilde{f}$ is fixed point free, but we do not actually need this.

Lemma 5.2.6 *Let X be an ~~infinite~~ attracting fixed infinite word for $\pi_1(f, \mu)$ which we represent as a path $\tilde{\gamma}$ in $\tilde{\mathcal{G}}$ starting from the vertex $\tilde{v}$ and suppose that $\tilde{f}$ is fixed point free. Then all but a finite number of edges of $\tilde{\gamma}$ (in $\tilde{\mathcal{G}}$) are oriented along γ.*

Proof: By 5.2.5 there is an m_0 such that for all $m \geq m_0$, $f(\gamma_m) = (\mu^{-1}\gamma_m)Y$, where $Y \neq 1$. Let $\tilde{E}$ be an edge of $\tilde{\gamma}$ with endpoints $\tilde{u}$ and $\tilde{w}$. Then $\eta[\tilde{v}, \tilde{u}] = \gamma_n$ and $eta[\tilde{v}, \tilde{w}] = \gamma_{n+1}$ for some n. Suppose that $n \geq m_0$. Then

$$\tilde{f}(\tilde{\gamma_n}) = [\tilde{f}(\tilde{v}), \tilde{v}]\tilde{\gamma}_n P_1, \text{ and}$$

$$\tilde{f}(\gamma_{\tilde{n}+1}) = [\tilde{f}(\tilde{v}), \tilde{v}]\gamma_{\tilde{n}+1}P_2$$

where these are reduced paths and P_1, P_2 are non trivial. Hence $\tilde{f}(\tilde{E}) \simeq P_1{}^{-1}\tilde{E}P_2$ where this product is not necessarily reduced. It does however imply that either there is a fixed point in the interior of $\tilde{E}$ or $\tilde{E}$ is oriented along $\tilde{\gamma}$, and so we are done.

It is shown in [CT94, proposition 3.3] that this orientation (we still assume that $\tilde{f}$ is fixed point free) satisfies the following:

(i) At each vertex that is not in a country, exactly one edge is oriented out. (ii) Among the edges incident to a particular country, at most one is directed out. (iii) There is at most one country with no outwardly directed edges.

We then proceed as in [CT94, proposition 3.3],

Proposition 5.2.7 *If the index of $\zeta_{\tilde{f}}$ is positive then $\tilde{f}$ has a fixed point.*

Proof: The case where $K - rank(\mathbf{T}, Fix(\zeta_{\tilde{f}})) \geq 2$ is standard. So we shall assume that this is not the case. Now we consider $\hat{\mathcal{G}} = \tilde{\mathcal{G}}/Fix(\zeta_{\tilde{f}}$. Since $t\tilde{f} = \tilde{f}t$ for all $t \in Fix\zeta_{\tilde{f}}$, $\hat{\mathcal{G}}$ inherits the edge directions from $\tilde{\mathcal{G}}$. Let $\hat{G}$ be the graph obtained from $\hat{\mathcal{G}}$ by collapsing all the countries. Now (iii) guarantees that at most one country in $\hat{\mathcal{G}}$ is essential. Then $\pi_1(\hat{\mathcal{G}}) = \pi_1(\hat{G}) * \pi_1(C)$ where C is the essential country and otherwise $\pi_1(\hat{\mathcal{G}}) = \pi_1(\hat{G})$. In the former case we must have that the edge directions in $\hat{G}$ have a sink. Now an attracting fixed infinite word will produce an infinite oriented ray in $\hat{G}$. The orientation is given by lemma 5.2.6 and the fact that this is a ray is a consequence of the fact that the word is not a limit of fixed words. Thus since we assume that $\zeta_{\tilde{f}}$ has positive index we must have that $\hat{G}$ contains either a sink and an infinite outward oriented ray or two infinite oriented rays or a loop and an infinite outward oriented ray. In each case if we consider a path between the two objects in question we get a contradiction to property (i) above.

We have thus shown that any automorphism of positive index can be point represented by our relative train track map.(In fact our proof also shows that if our automorphism has two attracting fixed infinite words then it can also be point represented). We note that by our proof this fixed vertex is not contained in a country. This can be stated in a slightly different form.

Corollary 5.2.8 *Let the pair $(f, \mathcal{G})$ be a relative train track map and suppose that C is an f-invariant component of $\mathcal{G}$. Suppose further that for some path μ in C, $\pi_1(f|_C, \mu)$ has positive index or has two attracting fixed infinite words, then there is a path α in C such that $f(\alpha) = \mu^{-1}\alpha$.*

The next step is to look at a relative train track map and consider attracting fixed infinite paths which start at fixed points. It is important that although our paths are infinite that they are regarded as paths and not just words since it is important to consider subpaths. Our aim is to reduce the counting of equivalence classes of attracting fixed infinite paths to the counting of edges. Here we say that two attracting infinite fixed paths,X_1 and X_2, starting at possibly different vertices are equivalent if there is a fixed path, ρ, such that $\rho X_1 = X_2$. If X_1 and X_2 start from the same vertex then this restricts to the equivalence relation we have already encountered for attracting fixed infinite words. (We note that X_1 is attracting as a path if and only if it is attracting as a word. This is an easy consequence of 5.2.3 and 5.1.14.) Otherwise we note that ρ defines a natural conjugacy between $\pi_1(f, v)$ and $\pi_1(f, w)$ (where v and w are the initial vertices for X_1 and X_2 respectively) under which X_1 gets sent to X_2 and hence the equivalence class of attracting fixed words containing X_1 gets sent to that containing X_2. So if ϕ is an automorphism of positive index, there is vertex, v, fixed by f such that ϕ is represented by $\pi_1(f, v)$. We call two vertices *Nielsen equivalent* if there is a Nielsen path connecting them. Let U be the set of vertices which are Nielsen equivalent to v. (Note no element of U can lie in a factor complex.) From the above remarks, there is a bijection between the equivalence classes of fixed infinite paths starting from some vertex in U and the equivalence classes of fixed infinite words of ϕ.

First we note that an attracting fixed infinite path can be given a height with respect to the strata of the relative train track. The lowest stratum $C(\mathcal{G})$ consists just of factor complexes and stems which are permuted up to homeomorphism. It is thus clear that no attracting infinite path can have height 0, and must always contain real edges.

So start with the case that we are given an attracting fixed infinite path, X, of height r where $\mathcal{H}_r$ is a level stratum. Since the transition matrix of $\mathcal{H}_r$ is a permutation matrix and X is fixed we must have that $\mathcal{H}_r$ consists of a single edge and subdivision of fixed points means that $\mathcal{H}_r = E$ and $f(E) = E\tau$ for some path $\tau \subset \mathcal{G}_{r-1}$.

Using the attractivity properties of X we easily get that X is equivalent to an attracting fixed infinite word X_0 where either X_0 has height $\leq r-1$ or the first edge of X_0 is E and every other edge of X_0 lies in $\mathcal{G}_{r-1}$. Now suppose we have two attracting fixed infinite paths both of which are of the form E...... Let C be the component of $\mathcal{G}_{r-1}$ containing the endpoint of E. Then we get that $\pi_1(f|_C, \tau)$ either has two inequivalent attracting fixed infinite paths or there is an indivisible Nielsen path of height r. In either case we can apply corollary 5.2.8 to obtain a path $\alpha \subset C$ with $f(\alpha) = \tau^{-1}\alpha$. The following is then immediate:

Proposition 5.2.9 *Let X be an attracting fixed infinite path of height r where $\mathcal{H}_r$ is level then either (i) There is exactly one attracting infinite path of height r and no INP's of height r or, (ii) There is an INP of height r and any attracting fixed infinite paths of height r are equivalent to attracting infinite paths of height $\leq r-1$.*

It is interesting to note that in case (i) above we have that the infinite attracting fixed infinite path is actually $\lim_{n\to\infty} f^n(E)$.

Next we deal with the growing case:

Proposition 5.2.10 *Let X be an attracting infinite path of height r where $\mathcal{H}_r$ is a growing stratum, then X is equivalent to an attracting infinite word X_1 where either X_1 has height $\leq r-1$ or there is an edge $E \in \mathcal{H}_r$ such that E is the first edge of $f(E)$ and* $\lim_{n\to\infty} f^n(E) = X_1$.

Proof: Let X' be the subpath of X as provided by corollary 5.2.4. Write X' as a concatenation of maximal length r-legal paths. $X' = \alpha_1 ... \alpha_k$.

Then $f(\alpha_1) = \sigma_1\tau_1, ..., f(\alpha_i) = \tau_{i-1}^{-1}\sigma_i\tau_i, ..., f(\alpha_k) = \tau_{k-1}^{-1}\sigma_k$, where all these products are reduced and the τ_i are the paths that cancel. Thus $f(X') = \sigma_1 ... \sigma_k$ but $X' \leq f(X')$ hence each σ_i must be non trivial and each turn (σ_i, σ_{i+1}) must be an illegal turn in $\mathcal{H}_r$. Again using the fact that $f(X') \geq X'$ we get that $\alpha_1 = \sigma_1, \alpha_2 = \sigma_2, ..., \alpha_k \leq \sigma_k$. and hence

$$f(\alpha_k) = \tau_{k-1}^{-1}\alpha_k{\sigma_k}'$$

So we get a fixed point somewhere in the path α_k, and hence there is a $\alpha_k' \leq \alpha_k$ with $f(\alpha_k') = {\tau_{k-1}}^{-1}\alpha_k'$ and so $\alpha_1 ... \alpha_k'$ is a fixed path and X is equivalent to

$$X_1 = \lim_{n\to\infty} f^n(\alpha_k'')$$

where $\alpha_k = \alpha_k'\alpha_k''$. Write $\gamma = \alpha_k''$. Now if $\gamma \subset \mathcal{G}_{r-1}$ then we are done. Otherwise we know that γ is r-legal and without loss of generality we may assume that the first edge of γ is in $\mathcal{H}_r$ for otherwise γ will have an initial fixed subpath. (Remember that $\gamma \leq f(\gamma)$). So if E is the first edge of γ we get that

$$\lim_{n\to\infty} f(\gamma) = \lim_{n\to\infty} f(E)$$

since γ is an r legal path. It is then clear that $E \leq f(E)$ and we are done.

Bearing in mind the above proposition we have the following,

Definition 5.2.11 Suppose that E is the edge of a growing stratum and X is an attracting fixed infinite path. Then we say that E generates X if, $X = \lim_{n\to\infty} f^n(E)$ and E is the first edge in its image.

The following is an easy corollary to the previous two propositions

Corollary 5.2.12

$$a(\pi_1(f,v)) \leq \sum_{u\in U} E(u)$$

where U is the set of vertices Nielsen equivalent to v, and $E(u)$ is the number of oriented real edges starting at u.

So corollary 5.2.12 almost proves the index theorem. The only difficulty is in taking account of Nielsen paths, which should limit the number of possible attracting fixed infinite paths. Now by if we have an INP ρ of height r where $\mathcal{H}_r$ is a growing stratum then $\rho = \alpha\beta^{-1}$, where α and β are r-legal paths and the turn between α and β is the only r-illegal turn in ρ.

Furthermore we have:

Lemma 5.2.13 *Let ρ be an INP of height r where $\mathcal{H}_r$ is a growing stratum, $\rho = \alpha\beta^{-1}$*

(i) α begins and ends with edges in $\mathcal{H}_r$, and similarly for β.

(ii) α and β end with distinct edges.

(iii) α and β start with distinct edges.

Proof: (i) is easily proved since if α has an initial subpath contained in $\mathcal{G}_{r-1}$ then this subpath must be fixed by the properties of a relative train track. This however contradicts the fact that ρ is indivisible. We also know that α and β end with edges from $\mathcal{H}_r$ since the turn between them is r-illegal.

(ii) is clear since $\alpha\beta^{-1}$ is reduced as written.

It only remains to show that α and β start with distinct edges. Now, $f(\alpha) = \alpha\tau$ and $f(\beta) = \beta\tau$ for some path $\tau \subset \mathcal{G}_{r-1}$. Hence if they start with the same edge (which must be an edge of $\mathcal{H}_r$) then without loss of generality, α must be an initial subpath of β. This is because α and β are r-legal and $\mathcal{H}_r$ is growing.

So $\alpha = \beta\beta'$ for some r-legal path β' which ends with an edge of $\mathcal{H}_r$. However we must have that $f(\beta') = \tau^{-1}\beta'\tau$ and since $f(\beta) = \beta\tau$ we deduce that τ is an initial subpath of $\beta'\tau$ as $\beta\beta'$ is r-legal and $\beta'\tau$ is reduced. This tells us that the r length of $f(\beta')$ cannot be more than the r length of β' which is a contradiction as β' is r-legal and the last edge of β' lies in $\mathcal{H}_r$.

Now we can attach to an equivalence class of attracting infinite paths a height, namely the height of the lowest path. That the number of these equivalence classes is affected by the number of INPs is the content of the next proposition.

Proposition 5.2.14 *Let $\mathcal{H}_r$ be a growing stratum with an INP $\rho = \alpha\bar{\beta}$ then,*

(i) The initial edges of α and β do not generate inequivalent attracting fixed infinite paths.

(ii) The initial edges of $\bar{\alpha}$ and $\bar{\beta}$ do not both generate attracting fixed infinite paths.

Proof: Let i_α,t_α, i_β and t_β be the initial and terminal edges of α and β respectively. Now by lemma 5.2.13 all of these edges lie in $\mathcal{H}_r$, they are not however all necessarily distinct. We do know that $i_\alpha \neq i_\beta$ and $t_\alpha \neq t_\beta$. Now as $f(\alpha) = \alpha\tau$ and $f(\beta) = \beta\tau$ and we know by the relative train track properties that

(a) The terminal edge of τ lies in $\mathcal{H}_r$ and,

(b) The terminal edge of τ is the same as the terminal edges of $f(t_\alpha)$ and $f(t_\beta)$.

Thus not both of $\bar{t_\alpha}$ and $\bar{t_\beta}$ can generate attracting infinite paths and we have proved (ii).

Now suppose that i_β generates an attracting infinite path X. Then

$$X = \lim_{n\to\infty} f^n(i_\beta) = \lim_{n\to\infty} f^n(\beta)$$

as β is an r-legal path. But then

$$\rho X = \lim_{n\to\infty} f^n(\rho\beta) = \lim_{n\to\infty} f^n(\alpha)$$

and so if either i_α or i_β generates an infinite attracting infinite path then they both do, but these paths are in the same equivalence class. This finishes the proof.

A simple euler characteristic argument almost proves the index theorem. If we have a fixed vertex v such that $\pi_1(f, v)$ is an automorphism of positive index, U is the set of vertices which are Nielsen equivalent to v and K is the number of INP's with endpoints at U then certainly,

$$a(\pi_1(f,v)) - 2 \leq \sum_{u\in U} E(u) - K \leq \sum_{u\in U} -|U| + 1 \tag{5.1}$$

These inequalities follow because if we have an INP of height r then either: (i)$\mathcal{H}_r$ is level and so an attracting fixed infinite path of height r is equivalent to one of lower height or, (ii) $\mathcal{H}_r$ is growing and the distinct initial edges of an INP do not give rise to inequivalent attracting fixed infinite paths. (Note that $K = |U| - 1 + K - rank(\pi_1(\mathcal{G}, v : \text{Fix}\pi_1(f,v).)$

To prove the index theorem we need to make full use of the graph of complexes discussed in chapter 1, which are used in [CT94] to prove the free product analogue of the Scott conjecture.

Recall that a graph of complexes Σ is constructed with a map $p : \Sigma \to \mathcal{G}$ so that the image of every vertex is a vertex fixed by f and the image of every edge is an INP. The notation Σ_r is used to denote the sub-graph of complexes of Σ generated by the preimages of INP's of height at most r, and Σ_r^v the component of Σ_r containing the vertex v. Recall that if $\mathcal{Z}$ is a connected graph of complexes, then

$$\text{K-rank}(\mathcal{Z}) = \text{K-rank}(\pi_1(\mathcal{Z}).$$

If $\mathcal{Z}$ has non contractible components, $\mathcal{Z}_1, \ldots, \mathcal{Z}_p$ then,

$$\tilde{\mathrm{K}}(\mathcal{Z}) = 1 + \sum_{i=1}^{p} (\text{K-rank}(\mathcal{Z}_i) - 1).$$

Suppose that $\mathcal{Z}_1 \subseteq \mathcal{Z}_2$ then,

Lemma 5.2.15 (3.10 [CT94]) $\tilde{\mathrm{K}}(\mathcal{Z}_1) \leq \tilde{\mathrm{K}}(\mathcal{Z}_2)$.

The inequality is strict if and only if either

(i) there is a path in $\mathcal{Z}_2$ meeting $\mathcal{Z}_2 - \mathcal{Z}_1$ whose endpoints lie in non contractible components of $\mathcal{Z}_1$ or,

(ii) there is a complex in $\mathcal{Z}_2 - \mathcal{Z}_1$ which is attached at a vertex of a non contractible component of $\mathcal{Z}_1$.

In [CT94] it is then shown that p induces an isomorphism, $p|\Sigma^v : \Sigma_v \to \mathrm{Fix}\pi_1(f, v)$ and that

$$\begin{aligned} \tilde{\mathrm{K}}(\Sigma_0) &\leq \tilde{\mathrm{K}}(\mathcal{G}_0) \\ \tilde{\mathrm{K}}(\Sigma_r) - \tilde{\mathrm{K}}(\Sigma_{r-1}) &\leq \tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1}). \end{aligned} \tag{5.2}$$

Recall that we have assigned a height to each equivalence class of attracting fixed inifinite paths (the least among the heights in the class) and that each such corresponds to an edge. (For growing strata, they are generated by a single edge and for level strata there is exactly one edge corresponding to the attracting fixed infinite path of least height.) We define

$$\mathrm{i}(\Sigma_r^v) = \text{K-rank}\Sigma_r^v + a(\Sigma_r^v)/2,$$

where $a(\Sigma_r^v)$ is the number of equivalence classes of attracting fixed inifinite paths of height at most r and which correspond to an edge whose initial vertex is a vertex in Σ_r^v. Then,

$$\tilde{\mathrm{i}}(\Sigma_r) = 1 + \sum \max(0, \mathrm{i}(\Sigma_r^v) - 1),$$

where the sum ranges over the components of Σ_r.

Note: Throughout this $f : \mathcal{G} \to \mathcal{G}$ is a relative train track map as described in chapter 1. However, in the following theorem we make a counting argument that will be simplified if we make some restrictions.

Suppose that $\rho = \alpha\bar{\beta}$ is an INP of height r, where $\mathcal{H}_r$ is a growing stratum. So $f(\alpha) = \alpha\tau$ and $f(\beta) = \beta\tau$ and α, β are r-legal paths. Let w be the vertex at which the illegal turn of $\mathcal{H}_r$ occurs (the endpoint of α and β) and let t_α, t_β be the first edges of $\bar{\alpha}$ and $\bar{\beta}$ respectively. By 5.2.14, the first edge of $f(t_\alpha)$ is the same as the first edge of $f(t_\beta)$ which is just the first edge of $\bar{\tau}$. We may thus subdivide these edges, $t_\alpha = t''_\alpha t'_\alpha$ and $t_\beta = t''_\beta t'_\beta$ so that $f(t''_\alpha) = f(t''_\beta)$ and then identify t''_α and t''_β to form a new edge t''. This is an operation called folding (as in the irreducible case) which induces a new topological map which is actually another relative train track map as discussed in [CT94] and [BH92]. For the purposes of the next theorem we assume that w, the vertex at which the illegal turn of the INP occurs, is not a vertex of $\mathcal{G}_{r-1}$, is incident to at least 3 edges of $\mathcal{H}_r$ and that one of these is the initial edge of τ. Since we may perform the above folding operation we know that every outer automorphism has a relative train track representative with the above property. Thus, although we only prove the following theorem for such relative train tracks, it will be sufficient for our purposes.

Theorem 5.2.16

$$\tilde{\mathrm{i}}(\Sigma_r) \leq \tilde{\mathrm{K}}(\mathcal{G}_r).$$

Proof: We shall actually prove, for each r, an equation as in 5.2,

$$\tilde{\mathrm{i}}(\Sigma_r) - \tilde{\mathrm{i}}(\Sigma_{r-1}) \leq \tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1}).$$

As $\mathcal{G}_0$ consists of factor complexes and their stems there can be no attracting fixed infinite paths of height 0, so $\tilde{\mathrm{i}}(\Sigma_r) = \tilde{\mathrm{K}}(\Sigma_r)$ and the result follows by 5.2. If $\mathcal{H}_r$ is either a descending stratum or a level stratum with an INP then there can be no equivalence classes of attracting fixed infinite words of height r and so again the result follows by that of 5.2. Suppose then that $\mathcal{H}_r$ is a level stratum with no INP, then

$$\tilde{\mathrm{i}}(\Sigma_r) - \tilde{\mathrm{i}}(\Sigma_{r-1}) = 0 \text{ or } 1/2.$$

In the former case there is nothing to prove. In the latter case, $\mathcal{H}_r$ consists of a single edge E such that $f(E) = E\tau$ for some $\tau \in \mathcal{G}_{r-1}$. Let v be the initial vertex of E. By the equation above, we must have that $\mathrm{i}(\Sigma^v_{r-1}) \geq 1$ and hence there is either a closed fixed path at v contained in $\mathcal{G}_{r-1}$ or there is an attracting fixed infinite word starting at v and contained in $\mathcal{G}_{r-1}$. In either case, v must lie in a non contractible component of $\mathcal{G}_{r-1}$. Also, as E occurs as the first edge of an attracting fixed inifinite path but by the discussion before 5.2.9 E only occurs once in this path. Hence the terminal vertex of E lies in a non contractible component of $\mathcal{G}_{r-1}$. Thus in this case,

$$\tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1}) \geq 1.$$

The only situation left to consider is where $\mathcal{H}_r$ is a growing stratum. This is by far the most difficult and we shall need to define various quantities. (Recall that $\mathcal{H}_r$ is a collection of real edges.)

For every *oriented* edge $E \in \mathcal{H}_r$ define $r(E)$ as follows:

(i) If there is a loop (reduced) in $\mathcal{G}_r$, with E as its first edge and meeting no non contractible component of $\mathcal{G}_{r-1}$ then $r(E) = 1$.

(ii) If there is a path (reduced) with E as its first edge and ending in a non contractible component of $\mathcal{G}_{r-1}$ then, $r(E) = 1$.

(iii) $r(E) = 0$ otherwise.

Then for each vertex, v, of $\mathcal{H}_r$ let,

$r(v) = \sum_{i(E)=v} r(E)$, if v lies in some non contractible component of $\mathcal{G}_{r-1}$ and

$r(v) = \max(0, (\sum_{i(E)=v} r(E)) - 2)$ otherwise.

Note that $i(E)$ denotes the initial vertex of E and that the sums above are taken over edges of $\mathcal{H}_r$. (In fact $r(E)$ is only defined if $E \in \mathcal{H}_r$.)

Then we define,

$$r(\mathcal{H}_r) = \sum_{v \in \mathcal{H}_r} r(v).$$

Claim: $r(\mathcal{H}_r)/2 \leq \tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1})$.

Proof of Claim: If no vertex of $\mathcal{H}_r$ lies in $\mathcal{G}_{r-1}$ then this is clear as we have a type of euler characteristic. Similarly, if no vertex of $\mathcal{H}_r$ lies in a non contractible component of $\mathcal{G}_{r-1}$ or if every vertex, v, which does has $r(v) = 0$, then the claim is clear. (One can reduce these cases to the previous one by using 5.2.15.)

So we suppose that there is a vertex, v, of $calH_r$ which lies in a non contractible component of $\mathcal{G}_{r-1}$ and which is the initial vertex of the edge $E \in \mathcal{H}_r$ and $r(E) = 1$. Form the graph $\mathcal{G}'_r$ from $\mathcal{G}_r$ by removing the edge E and $\mathcal{H}'_r$ similarly.

To prove the claim it will be enough to show that $r(\mathcal{H}_r) \leq r(\mathcal{H}'_r) + 2$ since by the definition of E it is clear that,

$$\tilde{\mathrm{K}}(\mathcal{G}'_r) - \tilde{\mathrm{K}}(\mathcal{G}'_{r-1}) = \tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1}) - 1.$$

Thus by repeatedly removing edges we will have shown the claim. We consider $\mathcal{G}'_r$ a sub-graph of complexes of $\mathcal{G}_r$ and we use subscripts $r_{\mathcal{H}_r}$ and $r_{\mathcal{H}'_r}$ to distinguish these quantities when applied to the same edge (or vertex).

Now, by our hypotheses, there is a vertex, v, which is the initial vertex of the edge $E \in \mathcal{H}_r$ with $r_{\mathcal{H}_r}(E) = 1$. Hence, either there is a reduced loop in $\mathcal{G}_r$ whose first edge is E or there is a reduced path whose first edge is E and which ends in a non contractible component of $\mathcal{G}_{r-1}$. In both cases call the path p, and let w be the first vertex of $\mathcal{H}_r$ crossed by p with $r_{\mathcal{H}_r}(w) > 0$. (In the case where p is a loop it could be that $v = w$.) By definition of w there is an edge e with initial vertex w, such that $r_{\mathcal{H}_r}(e) = 1$ and e is the edge pointing back along the path p toward v. Potentially, $r_{\mathcal{H}'_r}(e) = 0$, although this is not neccessarily the case. (It is clear that e is distinct from E as an *oriented* edge.)

Now consider an edge f of $\mathcal{H}_r$, which is distinct from E and e with initial vertex u. Suppose that $r_{\mathcal{H}_r}(u) > 0$ and that $r_{\mathcal{H}_r}(f) = 1$. We shall show that $r_{\mathcal{H}'_r}(f) = 1$. Let p_f be the path associated with f. So p_f is either a loop whose first edge is f or a path whose first edge is f and which ends in a non contractible component of $\mathcal{G}_r$. The only way that $r_{\mathcal{H}'_r}(f) \neq 1$ is if p_f is a path containing the edge $\bar{E}$, and without loss we may assume that the last edge of p_f is $\bar{E}$ and that otherwise every edge of p_f lies in $\mathcal{H}'_r$. Now we know that there is a path from v to w starting with the edge E and ending with the edge $\bar{e}$ and otherwise involving no occurence of either (or their inverses). This is just a subpath of p, given above. Neither f nor $\bar{f}$ may occur in this last path since we know that every vertex $u' \in \mathcal{H}_r$ which is crossed on this path has $r_{\mathcal{H}_r}(u') = 0$ and by hypothesis, $u = i(f)$ and $r_{\mathcal{H}_r}(u) > 0$. Thus there is a reduced path, $p'_f \subset \mathcal{G}'_r$, whose first edge is f which ends at w. If $u = w$ or w lies in a non contractible component of $\mathcal{G}_{r-1}$ then we are done. Otherwise call the last edge of p'_f, e'. As $r_{\mathcal{H}_r}(w) > 0$, there is an edge $e_1 \neq e, \bar{e'}$ with initial vertex w. (Also $e_1 \neq E$ as in this case w does not lie in a non contractible component of $\mathcal{G}_{r-1}$ and so $w \neq v$.) Thus $p'_f p'_{e_1}$ is a reduced path which lies in $\mathcal{G}'_r$ and hence $p'_f p'_{e_1} \bar{p'_f}$ is a path in $\mathcal{G}'_r$ which (after reduction) is a loop whose first edge is f. In every case, we get that $r_{\mathcal{H}'_r}(f) = 1$. Hence we have shown that

$$r(\mathcal{H}_r) \leq r(\mathcal{H}'_r) + 2,$$

which implies that,

$$r(\mathcal{H}_r)/2 \leq \tilde{\mathrm{K}}(\mathcal{G}_r) - \tilde{\mathrm{K}}(\mathcal{G}_{r-1}).$$

This completes the proof of the claim.

To complete the argument we now show that,

$$\tilde{\mathrm{i}}(\Sigma_r) - \tilde{\mathrm{i}}(\Sigma_{r-1}) \leq r(\mathcal{H}_r)/2. \tag{5.3}$$

By 5.2.10, if we are given an equivalence class of attracting fixed infinite paths of height r, then there is a fixed infinite path in this class which is generated by a single edge $E \in \mathcal{H}_r$, with $E \leq f(E)$. Furthermore, as E is the first edge of an infinite path, $r(E) = 1$. Thus we associate to each equivalence class of attracting fixed infinite words a unique edge, $E \in \mathcal{H}_r$, with $r(E) = 1$. Suppose that u is a vertex of $\mathcal{H}_r$ which is not the endpoint of the INP of height r. If $\mathrm{i}(\Sigma^u_{r-1}) > 0$ then u must lie in a non contractible component of $\mathcal{G}_{r-1}$, since either there is a closed Nielsen path at u or there is an attracting fixed infinite path at u (of height $r-1$). Then,

$$\mathrm{i}(\Sigma^u_r) - \mathrm{i}(\Sigma^u_{r-1}) \leq \sum_{i(E)=u} r(E)/2 \leq r(u)/2.$$

If $\mathrm{i}(\Sigma^u_{r-1}) = 0$, then

$$\max(0, \mathrm{i}(\Sigma^u_r) - 1) \leq (\sum_{i(E)=u} r(E)/2) - 1 \leq r(u)/2.$$

If there is no INP of height r, we may then add these inequalities to deduce 5.3. Suppose then that there is an INP, $\rho = \alpha\bar{\beta}$ of height r connecting vertices v_1 and v_2 and let w be the vertex at which the illegal turn of ρ occurs. Let E_α and E_β be the first edges of α and β. Now, E_α is the first edge in $f(E_\alpha)$ and this is an r-legal path. Thus $f^n(E_\alpha)$ is arbitrarily long and always has E_α as its first edge. Thus $r(E_\alpha) = 1$ and similarly $r(E_\beta) = 1$. Recall now the note before the statement of this theorem. This guarantees that we have an edge $t \in \mathcal{H}_r$, which has initial vertex w and is the first edge of τ. By the same argument, $r(t) = 1$. Let t_α, t_β be the first edges of $\bar{\alpha}$ and $\bar{\beta}$. We know that one of these edges does not generate an attracting fixed infinite path. Also, if $E_\alpha \neq E' \in \mathcal{H}_r$ has initial vertex v_1 and $r(E') = 1$ then clearly $r(t_\alpha) = 1$, by concatenating paths and a similar condition holds for v_2. We are now in a position to prove, 5.3, in the case where there is an INP of height r. Suppose that, $\mathrm{i}(\Sigma^{v_1}_{r-1}) = \mathrm{i}(\Sigma^{v_2}_{r-1}) = 0$ and without loss, that $r(v_1) \geq r(v_2)$. If E_β is the only edge with initial vertex, v_2 with $r(E_\beta) = 1$, then as E_α and E_β cannot generate inequivalent attracting fixed infinite paths (5.2.14),

$$\max(0, \mathrm{i}(\Sigma^{v_1}_r) - 1) \leq (\sum_{i(E)=v_1} r(E)/2) - 1 \leq r(v_1)/2.$$

If however, there is an edge $E' \neq E_\beta$ with initial vertex v_2 and $r(E') = 1$ then as there is a similar edge for v_1 (or we are back in the previous case) $r(w) > 0$. But as t does not generate an attracting fixed infinite path,

$$\begin{aligned}\max(0, \mathrm{i}(\Sigma_r^w) - 1) &\leq (\textstyle\sum_{i(E)=w} r(E)/2) - 1 - 1/2 \\ &\leq r(w)/2 - 1/2,\end{aligned}$$

and as before (E_α and E_β do not generate inequivalent attracting fixed inifinite paths),

$$\begin{aligned}\max(0, \mathrm{i}(\Sigma_r^{v_1}) - 1) &\leq (\textstyle\sum_{i(E)=v_1} r(E)/2) + (\textstyle\sum_{i(E)=v_2} r(E)/2) - 1 - 1/2 \\ &\leq r(v_1)/2 + r(v_2)/2 + 1/2.\end{aligned}$$

Hence in either of these cases, we deduce 5.3. The case where $\mathrm{i}(\Sigma_{r-1}^{v_1}) \neq \mathrm{i}(\Sigma_{r-1}^{v_2}) = 0$ is similar to the above, so we are left with considering the situation where, $\mathrm{i}(\Sigma_{r-1}^{v_1}), \mathrm{i}(\Sigma_{r-1}^{v_2}) > 0$. As before, this means that both v_1 and v_2 must lie in noncontractible components of $\mathcal{G}_{r-1}$ and so $r(t_\alpha) = r(t_\beta) = 1$. Thus we repeat the calculation above,

$$\begin{aligned}\max(0, \mathrm{i}(\Sigma_r^w) - 1) &\leq (\textstyle\sum_{i(E)=w} r(E)/2) - 1 - 1/2 \\ &\leq r(w)/2 - 1/2.\end{aligned}$$

Also,

$$\begin{aligned}\mathrm{i}(\Sigma_r^{v_1}) - \mathrm{i}(\Sigma_{r-1}^{v_1}) - \mathrm{i}(\Sigma_{r-1}^{v_2}) + 1 &\leq 1 - 1/2 + \textstyle\sum_{i(E)=v_1} r(E)/2 + \textstyle\sum_{i(E)=v_1} r(E)/2 \\ &\leq r(v_1) + r(v_2) + 1/2.\end{aligned}$$

where the above sum takes into account the effect of the INP and the fact that E_α and E_β cannot generate inequivalent attracting fixed infinite words. If there is a Nielsen path connecting v_1 to v_2 in $\mathcal{G}_{r-1}$, we may replace the first expression by, $\mathrm{i}(\Sigma_r^{v_1}) - \mathrm{i}(\Sigma_{r-1}^{v_1})$.

Adding these inequalities we get 5.3. We note that we have tacitly assumed that the vertices v_1, v_2, w are distinct, but that the arguments are changed only superficially if we drop this assumption. Thus we have proved the theorem.

Theorem 5.2.17 (Index Theorem) *If Φ is an outer automorphism of a finite free product G then,*

$$\mathrm{i}(\Phi) \leq \text{K-rank}(G).$$

Proof: Recall that

$$\mathrm{i}(\Phi) = 1 + \sum \text{K-rank}(G : \mathrm{Fix}\phi) + a(\phi)/2 - 1$$

where the sum is taken over a set of representatives of the positive similarity classes of Φ. There is a relative train track representative, $f : \mathcal{G} \to \mathcal{G}$, of Φ, satisfying the conclusion of the above theorem. Any positive similarity class is point represented by f and distinct similarity classes are represented on Nielsen inequivalent vertices (the proof of this is as in 2.1.10). If v is a vertex fixed by f then the number of equivalence classes of attracting fixed infinite paths starting at a vertex Nielsen equivalent to v is equal to the number of equivalence classes of attracting fixed infinite words for $\pi_1(f, v)$ so,

$$\mathrm{i}(\pi_1(f, v)) = \mathrm{i}(\Sigma^v)$$

and thus,

$$\mathrm{i}(\Phi) = \mathrm{i}(\pi_1(f)) = \tilde{\mathrm{i}}(\Sigma) \leq \tilde{\mathrm{i}}(\Sigma) \leq \tilde{\mathrm{i}}(\mathcal{G}) = \text{K-rank}(G).$$

Chapter 6

Irreducible Automorphisms

6.1 Train Tracks and Maximal Index

We now turn our attention and the proof of the previous section to the case of free groups and in particular irreducible automorphisms of maximal index. We start by representing our irreducible automorphism, ϕ,as a stable train track map f on a graph G. Actually this only defines an outer automorphism so we also have a vertex v which is fixed by f (by ~~lemma~~ 5.2.7) and a commuting diagram:

$$\begin{array}{ccc} F_n & \xrightarrow{\cong} & \pi_1(G,v) \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \pi_1(f,v)} \\ F_n & \xrightarrow{\cong} & \pi_1(G,v) \end{array}$$

We shall show that under the assumption that ϕ is an irreducible automorphism of maximal index very specific information can be obtained about the train track which represents it. Throughout the rest of the section ϕ is such an automorphism with a train track representative and a diagram as above.

Lemma 6.1.1 *f must have an INP, ρ, which has one endpoint at v.*

Proof: Suppose that there is no INP with an endpoint at v, then we know that $Fix\phi = 1$ and hence that $a(\phi) = 2n$. We also know from the previous section that attractive fixed words for ϕ correspond to edges E which start at v and such that $E \leq f(E)$, by 5.2.10 . Thus the valence of v is $2n$ and so G must be a rose (a wedge of n circles) with v as the principal vertex. Additionally, for every (oriented) edge E in G, $E \leq f(E)$. However in this case f does not define a homotopy equivalence and thus we get a contradiction.

We wish to simplify the situation even further, by reducing our analysis to the situation where the only vertices of G are the endpoints or endpoint of the INP ρ according to whether or not ρ is closed. This is the content of the next lemma.

Lemma 6.1.2 *Let ϕ be an irreducible automorphism of maximal index of F_n. Then ϕ has a stable train track representative f on a graph G with an INP ρ such that the only vertices of G are the endpoints of ρ.*

(We note that thus G has 1 vertex precisely when $Fix\phi \neq \{1\}$ but we do not preclude the possibility that if G has 2 vertices then one of these has valence 2. However the endpoints of INP's are the only vertices that we allow to have valence 2.)

Proof: We note that lemma 6.1.1 gives us a stable train track representative of ϕ and an INP ρ with one endpoint at v. Let us first consider the case where ρ is not closed and joins v to a vertex w, where these have valencies n_v and n_w respectively. Now by assumption $a(\phi) = 2n$ since Fix$\phi = \{1\}$. (Stable train tracks have at most one INP.) By 5.1 we get that

$$n_v + n_w - 1 \geq a(\phi) = 2n$$

However note that if the illegal turn of ρ occurs at v or w then we get the stricter inequality (using 5.2.14)

$$n_v + n_w - 2 \geq a(\phi) = 2n$$

In the latter case we have that, by an euler characteristic argument, v and w are the only vertices of G and hence we are done. In the former case there must be another vertex y of G with valence $n_y \geq 3$ at which the illegal turn of ρ occurs. Another euler characteristic argument leads us to conclude that v,w and y are the only vertices of G and that $n_y = 3$.

In the case where ρ is closed we get that $a(\phi) = 2n - 2$ and hence that

$$n_v - 1 \geq a(\phi) = 2n - 2$$

As before if the illegal turn of ρ occurs at v then

$$n_v - 2 \geq a(\phi) = 2n - 2$$

In the latter case the graph G has exactly one vertex, namely v. In the former case there G has two vertices v and y where the illegal turn of ρ occurs at y and $n_y = 3$. For both the problem situations, that is when we have a vertex y of valence 3, we fold the illegal turn of ρ which occurs at the vertex y and follow this by a valence 2 homotopy at the vertex produced by the folding. Note that as the train track is stable the fold must be full and hence we remove the vertex y and are left with a stable train track whose only vertices are at the endpoints of ρ.

Let us consider the case where the graph G, given by lemma 6.1.1, is a rose. This can happen in one of two ways: either ρ is a closed INP and G has exactly one vertex or ρ is not closed, G has two vertices and without loss of generality the vertex w (which is an endpoint of ρ) has valence two and the illegal turn of ρ must occur at v.

In either case it is possible to obtain a basis of F_n and a very neat description of the images of the basis under ϕ. Since G is a rose label the edges of G, $E_1, ...E_n$. In the case where G has a valence two vertex we have also subdivided one of these edges say $E_j = E'E''$. In this case we assume that ρ ends at the valence two vertex and that the last edge of ρ is E'. By 5.2.14 and 5.2.10, we have that for all the oriented edges of G except one, $E \leq f(E)$ and that the exception to this is part of the illegal turn of ρ. Since the illegal turn of ρ must occur at a vertex of valence at least three we assume without loss that ${E_n}^{-1} \not\leq f({E_n}^{-1})$. Now if we let $x_i \in F_n$ correspond to E_i and $g \in F_n$ correspond to ρ if ρ is closed and to $\rho E''$ otherwise, we get:

If ϕ is a maximal index automorphism represented on the rose then there is a basis $x_1, x_2, ..., x_n$ and an element $g \in F_n$ such that

(i) $\phi x_i = x_i ... x_i$ for $1 \leq i \leq n-1$

(ii) $\phi x_n = x_n ... x_1$, and

(iii) The words $\phi^k {x_i}^{\pm 1}, 1 \leq i \leq n, k \geq 0$ are all 'naturally' reduced. (6.1)

That is, applying ϕ to each letter of $\phi^{k-1} {x_i}^{\pm 1}$ and then concatenating, results in a reduced word.

Then we have that each of $\lim_{k \to \infty} \phi^k {x_i}^{\pm 1}$ for $1 \leq i \leq n-1$ and $\lim_{k \to \infty} \phi^k x_n$ is an attracting fixed infinite word.

The role of the element g is different in the two cases and derives from the INP ρ. If $\text{Fix}\phi = \{1\}$, then $\lim_{k \to \infty} \phi^k g$ is an attracting infinite fixed word and two attracting infinite words are equivalent only if they are equal (the attracting infinite words listed above are a complete set). This is the case where ρ is not closed and $a(\phi) = 2n$.

If $\text{Fix}\phi \neq 1$ then $\text{Fix}\phi = \{g\}$, ρ is closed and for some i and j $\lim_{k \to \infty} \phi^k x_i = g \lim_{k \to \infty} \phi^k x_j$.

The other case given by lemma 6.1.1 is where the graph G has exactly two vertices each of valence at least three and there is an INP connecting them. Using our standard arguments we know that for every oriented edge E bar one, $E \leq f(E)$. (Using 5.2.14 and 5.2.10)

For some unique edge e we have that $f(e) = e...e'$ where $e' \neq e$ and that $\{e^{-1}, e'^{-1}\}$

is the unique illegal turn of G. Note that since f is a stable train track map this means that all folds are full and hence that $e' \neq e^{-1}$. We assume that the illegal turn of ρ occurs at the vertex w and hence that both e and e' end at w.

We now wish to describe the automorphism $\pi_1(f, v)$ with respect to some 'nice' basis. This is quite elementary, however the statement has to allow for several similar but distinct cases.

We label the edges of G as follows: the edges which start and end at the vertex v are labelled A_i. (Note that there is a choice of orientation involved.) The edges which start and end at the vertex w are labelled B_j and the edges which start at v and end at w are labelled C_k.

We then choose one of the connecting edges C_0 to be our maximal tree and we arrive at a basis of $\pi_1(G, v)$ which is $\{x_i\} \cup \{y_j\} \cup \{z_k\}$ where

$$x_i \simeq A_i$$

$$y_j \simeq C_0 B_j {C_0}^{-1}$$

$$z_k \simeq C_k {C_0}^{-1} \text{ for } C_k \neq C_0$$

If we define the element $h \in \pi_1(G, v)$ so that $h \simeq f(C_0){C_0}^{-1}$ then we get that with the exception of one of the basis elements.

$$x_i \mapsto x_i ... x_i$$

$$y_j \mapsto h y_j ... y_j h^{-1}$$

$$z_k \mapsto z_k ... z_k h^{-1}$$

where each of these products is reduced as written. Also note that since $C_0 \leq f(C_0)$ then h cannot begin with ${x_i}^{\pm 1}$ nor can it begin with z_k, although it can begin with ${z_k}^{-1}$ or ${y_j}^{\pm 1}$. The exception in the above statement occurs due to our edge e since $e^{-1} \not\leq f(e^{-1})$, and to the illegal turn $\{e^{-1}, {e'}^{-1}\}$.

First we suppose that $e, e' \neq C_0$. We then get that $e = C_k$ for some $C_k \neq C_0$ or that for some j, $e = B_j^{\pm 1}$.

If $e = C_k$ then we have

$$z_k \mapsto z_k...ah^{-1}$$

where the letter a corresponds to the edge e' and hence either $a = z_{k'}$ for some $k' \neq k$ or $a = y_j^{\pm 1}$ for some j.

Otherwise (still assuming that $e, e' \neq C_0$) we have that $e = B_j^{\pm 1}$ for some j and so

$$y_j \mapsto hy_j...ah^{-1} \text{ or } ha...y_jh^{-1}$$

where again, a is the letter corresponding to the edge e' and either $a = z_k$ for some k or $a = y_{j'}^{\pm 1}$ for some $j' \neq j$.

The case where $C_0 = e$ or e' is slightly more complicated. First we suppose that the illegal turn is $\{C_0{}^{-1}, C_1{}^{-1}\}$ and we use the fact that f is a stable train track map so that all folds are full. Hence either $f(C_0{}^{-1})$ is an initial subpath of $f(C_1{}^{-1})$ or vice versa. In the former case,

$$z_1 \mapsto z_1...a$$

where the letter a is either $z_k{}^{-1}$ or $x_i{}^{\pm 1}$.

In the latter case we can write $h = h'w$ where w begins with z_1 and then,

$$z_1 \mapsto h'^{-1}.$$

We note that if we instead took C_1 to be our maximal tree, we turn this case into the previous one, avoiding this part of the analysis.

The final cases occur where the illegal turn is $\{C_0{}^{-1}, B_j{}^{\pm 1}\}$. We shall just deal with the case where $B_j{}^{-1}$ forms part of the illegal turn and simply state the results for the other case since it is very similar.

As before, we use the fact that f is a stable train track to deduce that either $f(C_0)$ has $f(B_j)$ as a terminal subpath or vice versa. In the former case we write

$$f(C_0) = pf(B_j)^m$$

where $m \geq 1$ is chosen maximally so that p does not have $f(B_j)$ as a terminal subword. Additionally we know that p begins with C_0. Thus we can write $h = h'w^m$ where w begins with y_j and so,

$$y_j \simeq C_0 B_j {C_0}^{-1} \mapsto pf(B_j)p^{-1} \simeq h'wh'^{-1}$$

If this product is reduced as written then we are done. Otherwise a little further analysis is needed and we subdivide C_0 into $m+1$ paths, $C, {C_0}^{(1)}, {C_0}^{(2)}, ..., {C_0}^{(m)}$. The subdivision is such that $f(C) = p$ and $f({C_0}^{(i)} = f(B_j)$. We then perform a sequence of m foldings, in effect folding each ${C_0}^{(i)})$ with B_j. A rewriting occurs at each stage but it is clear that B_j survives and it's image may be folded. We finally arrive at a stable train track map f' on a graph G' whose topology is the same as that of G. We shall denote the image of an edge $e \in G$, in G' by e^*. By assumption we have that the path $f(C)f(B_j)f(C^{-1})$ is not reduced and hence that either $\{C^{*-1}, {B_j}^{*-1}\}$ or $\{C^{*-1}, {B_j}^*\}$ is the unique illegal turn in G'. However we may then take preimages of these turns in G to deduce that $f(C)$ is a terminal subpath of either $f(B_j)$ or $f({B_j}^{-1})$. (Remember that by definition $f(B_j)$ and $f({B_j}^{-1})$ cannot be terminal subpaths of $f(C)$.) Hence for some path q, either

$$f(B_j) = qf(C) \text{ or } \; f(B_j) = f(C^{-1})q$$

We can easily see that only the first of these can occur since $f(C_0) = f(C)f(B_j)^m$ and is reduced. We recall that we may write $h = h'w^m$ where the first letter of w is y_j. In addition we get that $w = w'h'$ and so that

$$y_j \mapsto h'w'$$

Note that had we assumed that B_j formed part of the illegal turn then we would have that $h = h'z$ where z begins with ${y_j}^{-1}$ and that ${y_j}^{-1} \mapsto h'zh'^{-1}$ and hence if we put $w = z^{-1}$ then w ends with y_j and

$$y_j \mapsto h'wh'^{-1}$$

and if this is not reduced we may write $z = z'h'$ and so if $w' = z'^{-1}$ then

$$y_j \mapsto w'h'^{-1}.$$

Note that the folding operations can be applied in both the cases considered above. Hence if we are only interested in choosing a good basis we can skip the case where $f({B_j}^{\pm 1}$ is a terminal subpath of $f(C_0)$ turning this into one of the other cases.

The last case to consider is where $f(C_0)$ is a terminal subpath of $f(B_j)$ so we may write

$$f(B_j) = pf(C_0)$$

where p is a path whose first letter is B_j.

As always this path is reduced as written. Now as $f(C_0)$ and $f(B_j)$ end with the same edge, this edge must be either C_0 or B_j as there exists exactly one oriented edge E with $E \not\leq f(E)$. Using this fact agian tells us that $f(C_0)$ begins with C_0 and $f(B_j)$ begins with B_j from which we deduce that p begins with B_j and ends with ${C_k}^{-1}$ for some $k \neq 0$ or ${A_i}^{\pm 1}$ for some i. Thus the path $f(C_0)p$ is reduced. If we then let z be the group element corresponding to the closed path C_0p, then z begins with y_j and ends with ${z_k}^{-1}$ for some k or ${x_i}^{\pm 1}$ for some i and

$$y_j \mapsto hz.$$

Had we assumed in this case that $f(C_0)$ were a terminal subword of $f({B_j}^{-1})$ then we would have found a group element z^{-1} whose first letter was ${y_j}^{-1}$ and whose last letter was ${z_k}^{-1}$ for some k or ${x_i}^{\pm 1}$ for some i and then

$$y_j \mapsto zh^{-1}.$$

An important observation to be made after this analysis is that if we regard our basis elements as paths in the graph G then the image of each of these paths is always legal even if the paths themselves are not necessarily.

We now collect these results in the following proposition. We note that the products given are always reduced as written.

Proposition 6.1.3 *Let ϕ be a maximal index irreducible automorphism of F_n. Then either ϕ has a stable train track representative on a rose and has the form described in 6.1 or F_n has a basis $\{x_i\} \cup \{y_j\} \cup \{z_k\}$, which has a special element $e \in \{y_j\} \cup \{z_k\}$, together with elements $g, h \in F_n$ where h does not begin with ${x_i}^{\pm 1}$ or z_k and where $g^{-1}\phi(g) = h^{-1}$. Then for except for the special element e*

$$x_i \mapsto x_i ... x_i$$

$$y_j \mapsto h y_j ... y_j h^{-1}$$

$$z_k \mapsto z_k ... z_k h^{-1}$$

For the special element one of the following occurs:

(i) $e \mapsto e...ah^{-1}$ *where* $a \in \{z_k\} \cup \{{y_j}^{\pm 1}\} - \{e^{\pm 1}\}$

(ii) $e \mapsto he...ah^{-1}$ *where* $a \in \{z_k\} \cup \{{y_j}^{\pm 1}\} - \{e^{\pm 1}\}$

(iii) $e \mapsto e...a$ *where* $a \in \{{x_i}^{\pm 1}\} \cup \{{z_k}^{-1}\} - \{e^{\pm 1}\}$

(iv) $e \mapsto he...a$ *where* $a \in \{{z_k}^{-1}\} \cup \{{x_i}^{\pm 1}\} - \{e^{\pm 1}\}$.

Furthermore for each basis element a, $\phi^k(a)$ is 'naturally reduced' for all $k \geq 1$. That is to say that if we apply ϕ to each of the letters of $\phi^{k-1}a$ in reduced form and then concatenate the result is reduced as written.

Notes on proof: The proof for this is actually given above. The only points to note are that the element g is taken from the INP ρ. Actually we have that $g \simeq \rho {C_0}^{-1}$. The statement that each $\phi^k(a)$ is naturally reduced is just a statement about train tracks and that each image of a basis path is legal. We have reduced the number of cases outlined above by either changin the maximal tree, repeated folding or by relabelling a y_j to ${y_j}^{-1}$. However we could have gone futher and got rid of the fourth case by folding and switching the role of the vertices v and w. It is clear that the above automorphism has $2n$ attracting fixed infinite words. Each basis element except for e defines exactly two attracting infinite words. (This is clear for the x_i and z_k. For the y_j we need to premultiply by g. In other words $\lim_{k\to\infty} \phi^k(gy_j)$ is a fixed attracting infinite word. This will work for ${y_j}^{-1}$ and ${z_k}^{-1}$ similarly.) In the same way e will get us one fixed attracting infinite word. The last one comes from the edge C_0 and is seen on the group level as $\lim_{k\to\infty} \phi^k h$.

We note that we have the following commuting diagram:

$$\begin{array}{ccc} F_n & \xrightarrow{\gamma_g} & F_n \\ \big\downarrow \phi & & \big\downarrow \gamma_h \circ \phi \\ F_n & \xrightarrow{\gamma_g} & F_n \end{array}$$

where by γ_g we mean conjugation by g. Our choice of orientation has meant that we have discarded the basis arrived at by conjugating by g. This basis would have a similar form to the one above, but with the role of the x_i's and y_j's swapped round as well as the z_k's going in the other 'direction'.

One last thing to note is that if an automorphism is given as in the proposition then it is clear that we may construct a train track map which represents it. We simply construct a graph with two vertices, v and w. We place loops labelled A_i at v so that the number corresponds to the number of x_i, and similarly loops B_j at w which correspond to the y_j and edges from v to w corresponding to the z_k and labelled by C_k. Pick a connecting edge C_0, and we construct a train track map f on this graph by the identification

$$A_i \simeq x_i$$

$$C_0 B_j {C_0}^{-1} \simeq y_j$$

$$C_k {C_0}^{-1} \simeq z_k$$

$$f(C_0){C_0}^{-1} \simeq h$$

These conditions are enough to construct f and the statement of the proposition ensures that f is a train track map.

6.2 Rank Two Automorphisms

We now wish to comment on the rank 2 case which is always more simple and show that maximal ~~rank~~ index outer automorphisms are particularly easy to describe.

Recall that a reducible outer automorphism permutes some non trivial set of free factors of the group. In rank 2 an outer automorphism will be reducible if it fixes the conjugacy class of some primitive element (up to inversion) or $F_2 = A * B$ for some non trivial free factors A and B which are interchanged by the automorphism.

Thus consider an outer automorphism, Φ, of F_2 of maximal index. It is easy to check that in the reducible case the only possibilities are that Φ contains an automorphism ϕ of maximal rank. In the irreducible case, an finite order automorphism cannot have maximal index. This is because in a stable train track representative it permutes all the edges, so cannot have a fixed infinite word or a Nielsen path. Hence we are left with the case where Φ is irreducible, and it has a stable train track representative whose transition matrix has Perron Frobenius eigenvalue greater than one. We wish to show that Φ has a stable train track representative on the rose. Suppose first that a stable train track representative for Φ, f, G, does not support an INP. From the index formula Φ has at most 2 similarity classes of positive index and we know that each of these will be point represented in the stable train track.

Let v be a vertex of G fixed by f. By 5.2.10, the number of equivalence classes of attracting infinite words $a(\pi_1(f,v))$ is bounded by the the number of edges E with initial vertex v and with $E \leq f(E)$. Since there is no INP, there must be two fixed vertices in the graph each the initial vertex of 3 such oriented edges. (In rank 2 a graph cannot have a vertex of valence more than 3 unless it is the rose.) However, this accounts for all the oriented edges of the graph and so f cannot be a homotopy equivalence. Thus if Φ has a stable train track representative, (f, G), where G is not a rose, then f must support an INP. Each vertex of the graph has valence 3, and the illegal turn of the INP, must occur at a vertex of valence at least 3. Hence the illegal turn occurs at a vertex of exactly 3 and as in the previous section, we may fold the INP at the illegal turn and then perform a valence 2 homotopy at the newly created vertex. We end up with a graph with one less vertex. In rank 2 there are only three graphs to consider, and reducing the number of vertices always gets us a rose. So Φ has a stable train track representative on a rose, (f, R_2). Again we argue that f must support an INP. If not, then we know that an automorphism point representative an a vertex which is not the principle vertex cannot have positive index (it has valence 2). Hence the automorphism induced at the principle vertex has maximal index and we use the previous section to concluded that there must be an INP. The unique illegal turn then occurs at the principal vertex. Label the edges of R_2, A, B. Without loss we may say that the illegal turn of f is $\{\bar{A}, \bar{B}\}$. Thus $f(A)$ connot contain any occurence of the subwords $A\bar{B}$ or $B\bar{A}$ and similarly for $f(B)$. Hence, after possibly reorienting the edges, we see that $f(A)$ and $f(B)$ are positive.

We claim that there is a basis $\{a, b\}$ of F_2 and a $\phi \in \Phi$ (the automorphism point represented on the principal vertx of the rose) such that

(i) ϕ is a positive automorphism.

(ii) ϕa begins and ends with a, ϕb begins with b and ends with a.

The first statement follows from the statement that ϕ is point represented at the principle vertex of a rose. If the INP of f, has an endpoint at the principle vertex, then ϕ has maximal index and (ii) follows from the previous section. Otherwise, we know that an automorphism point represented at a vertex of valence 2, has index at most 3/2, and this is only if the INP has an endpoint at that vertex. (Recall that the distinct first and last edges of the INP do not generate inequivalent attracting fixed infinite paths.) Hence for Φ to have maximal index, ϕ would have index of 3/2, and (ii) follows in this case as well, by 5.2.10.

An easy inductive argument shows that any automorphism satisfying (i) and (ii) can be written as a product of the two elementary Nielsen transformations

$$\begin{array}{ll} a \mapsto a & \text{and} \quad a \mapsto ab \\ b \mapsto ba & \qquad\;\; b \mapsto b \end{array}$$

Another easy induction will show that any automorphism which can be written as a product of the above two Nielsen transformations not only has the form described above, but also fixes the commutator $aba^{-1}b^{-1}$. Hence we have shown the following:

Corollary 6.2.1 *Let Φ be a maximal index irreducible outer automorphism of F_2 then there is a $\phi \in \Phi$ of maximal index and either ϕ has maximal rank or there is a basis $\{a, b\}$ of F_2 such that*

(i) ϕ is positive with respect to this basis.

(i) ϕ fixes $aba^{1-}b^{-1}$

(iii) $\phi a = a...a$ and $\phi b = b...a$.

Note: It is also clear that an automorphism as above has maximal index since each of $\lim_{k\to\infty} \phi^k a^{\pm 1}$, $\lim_{k\to\infty} \phi^k b$ is an infinite attracting infinite word. Also, if ϕ is an automorphism of maximal rank, then there is a basis $\{a, b\}$ of F_2, such that,

$$\begin{aligned} \phi a &= a \\ \phi b &= ba^n \text{ , for some } n \end{aligned}$$

and hence may also be written as a product of the above Nielsen transformations.

6.3 Normal Forms

We use in this section a solution of the conjugacy problem, given in [Los96] to determine how unique the form in proposition 6.1.3 is. One could in fact just state the algorithm given there from which one could determine whether two irreducible automorphisms of maximal index are conjugate, however we wish to give a more direct formulation for this special case.

We start by defining a few terms. In [Los96] a forward evolution path is defined to be a sequence of marked topological representatives for some irreducible outer automorphism Φ;

$$(f_0, G_0) \to (f_1, G_1)... \to (f_i, G_i) \to ...(f_m, G_m)$$

where each transformation $(f_i, G_i) \to (f_{i+1}, G_{i+1})$ is either the collapsing of an invariant or pretrivial forest or a folding or a quasi folding operation as defined in [Los96]. The folding and quasi-folding operations are the same as the folding operation of [BH92] except that the collapsing of pretrivial and invariant forests is *not* part of the operation and if the fold occurs at a valence 3 vertex, it is immediately followed by a valence 2 homotopy at the newly created vertex. To avoid confusion we shall use the term *quasi-folding* for the foldings defined in [Los96] and folding for foldings in the sense of [BH92], as given in chapter 1. Another point to note is that in [Los96] the vertices of graphs are assumed to have a valence of at least 3. We shall assume that this is the case from now on, but in fact this is not a real restriction since we may apply valence 1 and valence 2 homotopies to any topological map.

Then a backward evolution path is defined to be a forward evolution path in reverse and an evolution path is a sequence of marked topological representatives where each transformation is either a forward or a backwards evolution path. An elementary evolution path is an evolution path in which every transformation in the sequnce is either a quasi folding or a collapsing.

In [Los96, Theorem 3.3.1] it is proved that if Φ is an irreducible outer automorphism of F_n then;

Theorem 6.3.1 *Let (f, G) and (f', G') be topological representatives of Φ then there is an evolution path connecting them.*

This is actually improved upon there. If (f, G) is a topological representative of an irrducible outer automorphism, then $\lambda(f)$ is the largest eigenvalue of the transition matrix for f. It is then shown that,

Theorem 6.3.2 *Let (f, G) and (f', G') be topological representatives of Φ then there is an elementary evolution path,*

$$(f, G) = (f_0, G_0) \rightarrow (f_1, G_1) \rightarrow \ldots \rightarrow (f_t, G_t) = (f', G'),$$

where for all i, $\lambda(f_i) \leq \max(\lambda(f), \lambda(f'))$.

Note that an evolution path is said to connect two marked topological representatives if the said representatives occur at each end of the sequence.

For our purposes we would like a slightly different formulation of this theorem.

Definition 6.3.3 A *special fold* is a fold as in [BH92], followed by a valence 2 homotopy if the fold creates a velence 2 vertex.

The point of the above definition is that it combines the two slightly different notions of foldings in [BH92] and [Los96] to enable us to state the results of the former for train track maps.

We define a special forward evolution path to be a sequence of marked topological representatives where each transformation is a special folding. The terms special backward evolution path and special evolution path are defined analogously.

Now, for our irreducible outer automorphism Φ we let $T(\Phi)$ denote the set of projective equivalence classes of train track representatives of Φ and let $ST(\Phi)$ denote the subset defined by the stable train track representatives of Φ. Recall that two graphs are projectively equivalent if they have the same combinatorial structure and it is possible to relabel the edges of one graphg to turn it into the other. Two topological maps, (f, G) and (f', G'), are said to be projectivel equivalent if there exists such a relabelling $h : G \to G'$ making the following diagram commute:

$$\begin{array}{ccc} G & \xrightarrow{f} & G \\ \downarrow{\scriptstyle h} & & \downarrow{\scriptstyle h} \\ G' & \xrightarrow{f'} & G' \end{array}$$

In [Los96] train track maps are defined to be topological maps (without valence 1 or 2 vertices) representing an irreducible outer automorphism where the maximum eigenvalue of the transition matrix is minimal amongst all such. In [BH92] it is further required that the toplogical map have no invariant forests and this latter is the definition we use. We need the following lemma (contained in [Los96, 3.4.5]) to compensate for this difference.

Lemma 6.3.4 *The operation of collapsing an invariant forest commutes with that of quasi-folding.*

We then get an easy corollary to 6.3.2:

Corollary 6.3.5 *Let (f, G), $(f', G') \in T(\Phi)$ (resp. $ST(\Phi)$), then there is a special evolution path connecting them where each term of the sequence is in $T(\Phi)$ (resp. $ST(\Phi)$).*

Proof: This is clear as train tracks do not have invariant or pretrivial forests and the fact that special folds preserve the property of being a train track or a stable train track.

Also it is possible to apply the proof of [Los96, Theorem 4.2] to our situation and get,

Proposition 6.3.6 *Given two irreducible outer automorphisms, Φ and Ψ, of F_n the following are equivalent:*
(i) $T(\Phi) = T(\Psi)$ (resp. $ST(\Phi) = ST(\Psi)$),
(ii) $T(\Phi) \cap T(\Psi)$ is non empty ($ST(\Phi) \cap ST(\Psi)$ is non empty),
(iii) Φ and Ψ are conjugate in $\mathrm{Out}F_n$.

We now endow $ST(\Phi)$ with an oriented graph structure. The vertices are the elements of $ST(\Phi)$ and there is an oriented edge from (f, G, τ) to (f', G', τ') if there is a special folding operation taking (f, G, τ) to (f', G', τ'). Then 6.3.5 tells us that $ST(\Phi)$ is connected. Now if we consider those outer automorphisms whose stable train track representatives support an INP then, as these maps have exactly one illegal turn (which is the illegal turn of the INP), each vertex of $ST(\Phi)$ has exactly one outgoing edge. Hence the fundamental group of $ST(\Phi)$ as a graph is infinite cyclic. We know that this situation occurs if Φ is the outer automorphism of some maximal index automorphism and so this gives us a particularly easy algorithm for determining whether two marked stable train tracks give rise to the same maximal index automorphism. (We note that given two normal forms as in 6.1.3, it is relatively easy to construct train track maps representing them and to make the train track maps stable. Thus the algorithm given below is very easy to apply.)

We recall that $ST(\Phi)$ is finite as there are only finitely many matrices of a given Perron-Frobenius eigenvalue and size, and so as every vertex of $ST(\Phi)$ has exactly one outward edge if we follow the directions of edges from any vertex we must end up circling the topological core of $ST(\Phi)$.

We are now in a position to state the algorithm which determines whether two irreducible maximal index automorphisms are conjugate. Our starting point will be a stable train track representative.

Algorithm (Train Track)

Let ϕ, ψ be irreducible automorphisms of F_n of maximal index and Φ and Ψ be the corresponding outer automorphisms with marked stable train track representatives. Repeatedly fold the unique illegal turn of stable train track maps (that is, fold the INP) until we get a repetition (remembering to perform a valence 2 homotopy if we create a valence 2 vertex). As in the remarks after 6.3.5 we are then able to list the cores of $ST(\Phi)$ and $ST(\Psi)$ and futhermore we know that Φ is conjugate in $Out(F_n)$ to Ψ if and only if $ST(\Phi)$ and $ST(\Psi)$ are the same (this is due to 6.3.6) and hence they are conjugate if and only if the cores of these graphs are equivalent and in fact it is enough for the cores to coincide. Thus having listed the elements in the core we just check to see if an element of one is equivalent to an element of the other and this establishes a conjugacy between Φ and Ψ. By the index theorem 5.2.17 this also gives us a conjugacy between ϕ and ψ as index is preserved by isomorphism.

We now give an example of how to find the core for $ST(\Phi)$ where Φ is a maximal index irreducible outer automorphism. We start with the following stable train track map.

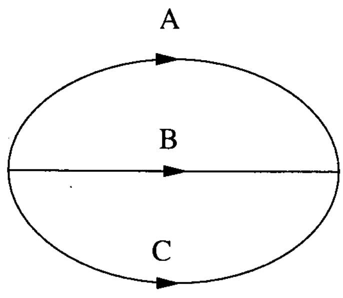

figure 1

$$\begin{aligned} A &\longmapsto AB^{-1}CB^{-1}A \\ B &\longmapsto BA^{-1}BC^{-1}B \\ C &\longmapsto CB^{-1}AB^{-1}AB^{-1}CB^{-1}A \end{aligned} \tag{6.2}$$

This is clearly a train track map whose only illegal turn is $\{A^{-1}, C^{-1}\}$. It is actually quite easy to check that this is a stable train track map and in fact the following calculations demonstrate that there are finitely many projective equivalence classes of this train track map. We note that the path $AB^{-1}CA^{-1}BC^{-1}$ is an INP.

If we fold the INP, by first subdividing $C = C_1C_2$, where $C_2 \mapsto AB^{-1}CB^{-1}A$ and then identifying C_2 with A we get the following stable train track map:

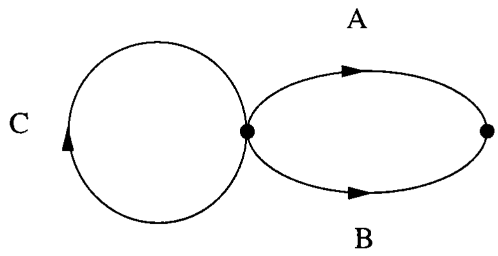

figure 2

$$A \longmapsto AB^{-1}C_1AB^{-1}A$$
$$B \longmapsto BA^{-1}BA^{-1}C_1{}^{-1}B$$
$$C_1 \longmapsto C_1AB^{-1}AB^{-1}$$

A quasi-folding operation at a valence three vertex incorporates a valence two homotopy across, in our case, A. However, in the case of train track maps this is equivalent to a reverse subdivision at the new valence 2 vertex. This, while clear, may also be checked by inspection. We remove the valence 2 vertex and we have an edge E corresponding to AB^{-1} and an edge D corresponding to C_1 and the following stable train track map:

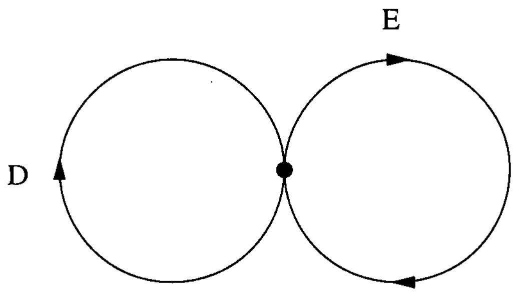

figure 3

$$D \longmapsto DE^2$$
$$E \longmapsto EDE^2DE^2$$

We note that we have an INP $DED^{-1}E^{-1}$ which we proceed to fold. As before we subdivide the edge E into edges E_1 and E_2 so that E_1 maps to EDE^2 and then fold E_2 with D and we get the resultant train track map:

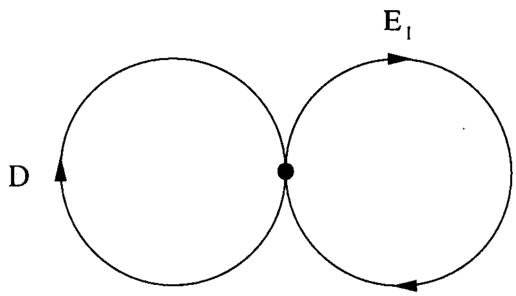

figure 4

$$D \longmapsto DE_1DE_1D$$
$$E_1 \longmapsto E_1D^2E_1DE_1D$$

We now relabel the edge D by the letter F and we relabel E_1 by G, so that,

$$F \mapsto FGFGF$$
$$G \mapsto GF^2GFGF$$

The INP is then $FGF^{-1}G^{-1}$ which we fold by subdividing G into G_1 and G_2, where G_1 maps to GF and then fold G_2 with F. Again as this is a full fold we end up with a rose where the edges are labelled F and G_1 and the map is given by:

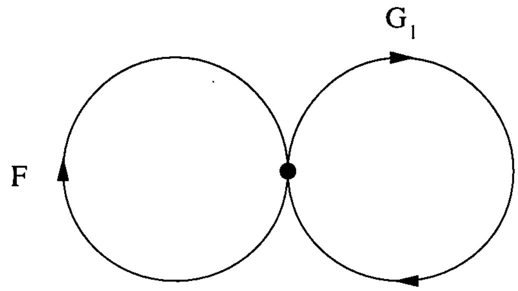

figure 5

$$F \longmapsto FG_1F^2G_1F^2$$
$$G \longmapsto G_1F^2$$

It is easy to see that the train track map given in figure 5 is equivalent to the one given in figure 3 where the equvalence is given by

$$F \mapsto E, G_1 \mapsto D$$

and hence we have found the core of $ST(\Phi)$. Establishing whether this automorphism is conjugate to another automorphism Ψ just involves following the same procedure, finding the core of $ST(\Psi)$ and seeing whether there is an equivalence between elements of the core. We note that this algorithm works for any irreducible outer automorphism whose stable train track maps have an INP, as then any stable train track representative will have a unique illegal turn. Furthermore we note that the results of section 2 can be extended to the case of outer automorphisms, but that then the final statements become very complicated.

The Core of $ST(\Phi)$

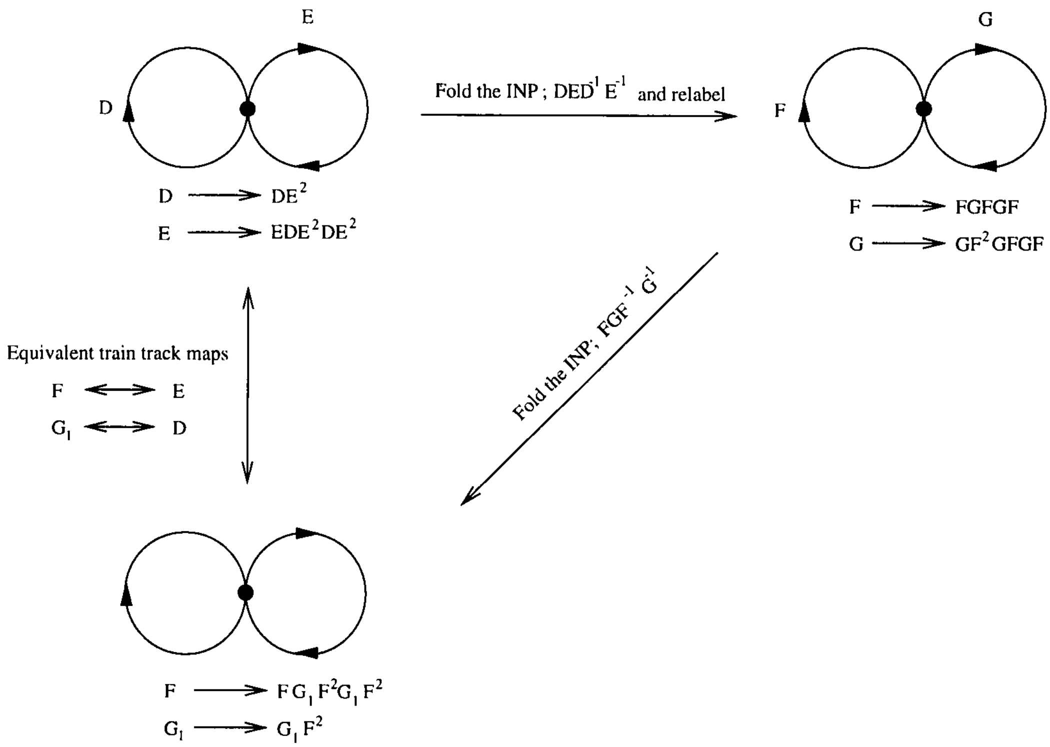
E
D
D ⟶ DE2
E ⟶ EDE2DE2
Fold the INP ; DED-1 E-1 and relabel
G
F
F ⟶ FGFGF
G ⟶ GF2GFGF
Fold the INP; FGF-1 G-1
Equivalent train track maps
F ⟷ E
G1 ⟷ D
F ⟶ FG1F2G1F2
G1 ⟶ G1F2

Chapter 7

Hyperbolic Groups

In this chapter our goal is to apply the result of Theorem 5.2.17 to the situation where our free product G is also a hyperbolic group and to try to get information about the action of an automorphism ϕ on the boundary of G.

Given a finitely generated group G and a generating set A of G we form the Cayley graph, $\Gamma_A(G)$ in the usual manner. We abbreviate this to Γ where no confusion will arise. A metric is defined on Γ by letting each edge have length one (in fact each edge is realised by a closed unit interval) and we define a metric d on Γ, the path metric by setting $d(x,y)$ to be the length of the shortest path joining x and y. For $g \in G$ we have that $|g|_A = d(g,1)$. Also if $g,h \in G$ then $d(g,h) = |gh^{-1}|_A$. We note that we are now dealing with the geometric realisation of Γ

Given a basepoint $w \in \Gamma$ there is an inner product

$$(x.y)_w = 1/2(d(x,w) + d(y,w) - d(x,y)).$$

We say that G is $\delta-$hyperbolic or just hyperbolic, if there is a constant $\delta \geq 0$ such that

$$\forall x,y,z \in \Gamma,\ (x.y)_w \geq \min\{(x.z)_w, (z.y)_w\} - \delta \tag{7.1}$$

For our purposes it is sufficient to note that this definition is indepedent of the choice of both basepoint and generating set. (Although these result in different values of δ.)

A sequence of points $\{x(i)\} \subset \Gamma$ is said to converge to infinity if

$$\lim_{i,j\to\infty} (x(i).x(j)) = \infty.$$

This is clearly independent of the basepoint chosen since, $|(x.y)_w - (x.y)_{w'}| \leq d(w, w')$.

Then define an equivalence relation on sequences which converge to infinity by $\{x(i)\} \sim \{y(i)\}$ iff

$$\lim_{i\to\infty} (x(i).y(i)) = \infty.$$

This relation is clearly reflexive and symmetric and the hyperbolicity ensures that it is transitive. The boundary ∂G of G is the set of equivalence classes of sequences which tend to infinity. Note that this is *not* the same as the set of infinite reduced words defined in chapter 5 - it is in general much bigger.

Now if $\{x(i)\}$ is a sequence which tends to infinity, then any subsequence also tends to infinity and moreover is equivalent to the original. Also $\{x(i)\}$ is clearly equivalent to a sequence all of whose elements are vertices of Γ or in other words elements of G. This is due to the fact that we can find a sequence $\{x'(i)\}$ of vertices with $d(x(i), x'(i)) \leq 1/2$ for all i. Given these comments from now on we assume that our chosen basepoint is the identity element of the group. Additionally, whenever we take a point of ∂G and a sequence $\{x(i)\}$ in it's equivalence class we will assume that each $x(i)$ is a vertex in the Cayley graph and in particular, each $x(i)$ will be a group element. If $x \in \partial G$ then a sequence $\{x_i\}$ is said to tend to x, $x_i \to x$, if $\{x_i\} \in x$. If $x \in \Gamma$, then $x_i \to x$ if this holds with the usual topology. The inner product can then be extended to the boundary,

Definition 7.0.7 Let $x, y \in \Gamma \cup \partial G$, then

$$(x.y)_S = inf\{\liminf_i (x_i.y_i)\},$$

where the infimum is taken over all pairs of sequences $x_i \to x$ and $y_i \to y$.

Some properties of the inner product,

Lemma 7.0.8 (4.5,4.6 [ABC+91]) *(1) If $x \in \Gamma$ and $y \in \Gamma \cup \partial G$ then $(x.y)_S = inf\{\liminf_i(x.y_i)\}$ where the infimum is taken over all sequences $y_i \to y$. (ie. it is enough to consider the constant sequence at x).*

(2) For $x, y \in \Gamma$, $(x.y)_S = (x.y)$

(3) For $x, y \in \Gamma \cup \partial G$, $(x.y)_S = \infty$ implies that $x = y \in \partial G$.

(4) iF $x, y \in \partial G$ and $x_i \to x, y_i \to y$ then

$$(x.y)_S \leq \liminf_i(x_i.y_i) \leq (x.y)_S + 2\delta.$$

As the extended inner product has all these properties (and more) the subscript is dropped.

Now there is a topology on $\Gamma \cup \partial G$ which has the following as a basis (see [ABC+91, Proposition 4.8]).

(1) $B_r(x) = \{y \in \Gamma | d(y,x) < r\}$, for each $x \in \Gamma$,

(2) $N_{x,k} = \{y \in \Gamma \cup \partial G | (x.y) > k\}$ for each $x \in \partial G$.

Clearly this extends the topology on Γ which is a dense subspace of $\Gamma \cup \partial G$. Now if $\alpha \in \mathrm{Aut}G$ then α induces a pseudo isometry on Γ so is defined on $\Gamma \cup \partial G$.

For a fuller treatment of the above material we refer the reader to [Gro88], [CDP90] and in particular [ABC+91] which account we have followed most closely.

We now restrict our attention to the situation where $G = G_1 * G_2 * \ldots G_n$ is our hyperbolic group. We fix a generating set A_i for each factor, and the union of these is the generating for G, which we call A. We now have two concepts of length in G. We shall denote by $|g|$ the length of the element g with respect to the generating set A. In other words $d(g,1) = |g|$. We also have the free product length which we denote $|g|_*$.

Now we have analysed to some extent infinite reduced words in G. Given such an infinite reduced word X, we can form a subsequence X_i of group elements, where X_i is as before the initial subword of X of length i (with respect to $|.|_*$).

The next lemma tells us that we may regard infinite reduced words as elements of ∂G. Note that for $w \in G$, w_k denotes the initial subword of w of free product length k and that this is taken to be w if $k \geq |w|_*$. Also as in Chapter 5 we take $w \wedge w'$ to be the longest (in free product length) initial common subword of w and w'.

Lemma 7.0.9 *Let $\{y(i)\}$ be a sequence which tends to infinity and X an infinite reduced word then,*

(i) $|y(i)_k| \leq (y(i).y(j)) \leq |y(i)_{k+1}|$ where $k = |y(i) \wedge y(j)|_*$

(ii) $\{X_i\}$ tends to infinity.

(iii) $\{y(i)\} \sim \{X_i\}$ if and only if $|X \wedge y|_*$ is unbounded.

Proof: (i) We may write $y(i) = y(i)_k az$ and $y(j) = y(i)_k bz'$ where a and b each belong to, perhaps the same, single factor $b^{-1}a \neq 1$ and the products are reduced as written. This follows from the definition of $y(i) \wedge y(j)$. Then, remembering that we always take our basepoint to be 1,

$$\begin{aligned} 2(y(i).y(j)) &= |y(i)_k| + |a| + |z| + |y(i)_k| + |b| + |z'| - |z'^{-1}b^{-1}az| \\ &= 2|y(i)_k| + |a| + |b| + |z| + |z'| - |z'| - |z| - |b^{-1}a| \\ &= 2|y(i)_k| + |a| + |b| - |b^{-1}a| \end{aligned}$$

Since $|b| - |a| \leq |b^{-1}a| \leq |a| + |b|$ we have proved (i).

To prove (ii) note that X_i is a subword of X_j if $i \leq j$ and so by using (i) we get that

$$(X_i.X_j) \geq \min\{|X_i|, |X_j|\}.$$

In fact by repeating the argument of the first part we see that this is actually an equality. Thus we have proved (ii).

To prove (iii) we consider the first implication. So suppose that $\{y(i)\} \sim \{X_i\}$. By (i) we know that

$$(X_i.y(i)) \leq |X_{k+1}| \text{ where } k = |y(i) \wedge X_i|_*.$$

and thus $|y(i) \wedge X|_*$ is not bounded since $(X_i.y(i)) \to \infty$. In fact this also shows that $|y(i) \wedge X|_* \to \infty$ as $i \to \infty$.

Conversely, suppose that $|y(i) \wedge X|_*$ is unbounded. Then,

$$\begin{aligned}(X_i.y(i)) &\geq \min\{(y(i).y(j)), (X_i.y(j))\} - \delta \\ &\geq \min\{(y(i).y(j)), |X_k|\} - \delta, \\ &\quad \text{where } k = |X_i \wedge y(j)|_* = \min\{|X \wedge y(j)|_*, i\}\end{aligned}$$

and the first inequality is due to 7.1, the second is simply part (i) of this lemma. Given $K \geq 0$ we may find an $N \geq K + \delta$ such that $(y(i).y(j)) \geq K + \delta$. (Since $\{y(i)\}$ tends to infinity.) We may also find a $j_0 \geq N$ such that $|X \wedge y(j_o)| \geq K + \delta$. (Since $|X \wedge y(j)|$ is unbounded.) Thus by the above inequality we have that for all $i \geq N$, $(y(i).X_i) \geq K$. Hence $\{y(i)\} \sim \{X_i\}$.

By the above lemma we may consider infinite reduced words of G a subset of ∂G which we shall denote G_∞. (By part (iii), two infinite reduced words cannot have the same image in ∂G.)

Now if $\alpha \in \text{Aut}G$ and $x \in \Gamma \cup \partial G$ then αx is defined. x is said to be attracting for α if there is an open set U (in the topology given above) containing x and such that for every $y \in U$, $\lim_{n\to\infty} \phi^n y = x$. We then get the following proposition.

Proposition 7.0.10 *Let X be an infinite reduced word fixed by $\alpha \in \text{Aut}G$. Then X is attracting in the sense of 5.1.14 if and only if it is attracting in $\Gamma \cup \partial G$.*

Proof: If $\{y(i)\}$ is a sequence in G then by 7.0.9,

$$|X_{k_n+1}| \geq \liminf_i (\alpha^n y(i).X_i) \geq |X_{k_n}|.$$

where $k_n = \liminf_i |\alpha^n y(i) \wedge X_i|_* = \liminf_i |\alpha^n y(i) \wedge X|_*$. By 7.0.8 this implies that if X is attracting for α in one topology then it is attracting in the other and vice versa. (We note that we use the fact that if $x \in \Gamma$ then there is an $x' \in G$ such that $|(x.X_i) - (x'.X_i)| \leq 1$, to apply the above inequality to the whole of Γ.)

With the above proposition we may apply the results of 5, in particular Theorem 5.1.14 to the situation where G is a hyperbolic group. However not every point in the boundary of G is an infinite reduced word. The obvious exceptions arise from the factor groups (recall that a free product is hyperbolic if and only if its factor groups are also). In general then we would expect that the boundary of G contain not only the infinite reduced words but also copies of the boundaries of its factor groups and all their conjugates. We say that $y \in \partial(G_j{}^w)$ if there is a sequence $\{y(i)\} \in y$ such that each $y(i) \in G_j{}^w$. The content of the next theorem is that these are all.

Theorem 7.0.11 *The boundary of G consists of the infinite reduced words and those points in $\partial(G_i{}^w)$, where G_i is a factor of G and $w \in G$.*

Proof: Let $\{y(i)\}$ be a sequnce of elements of G which tends to infinity. We will show that either this sequence is equivalent to one which lies completely in $G_i{}^w$ for some factor G_i and some $w \in G_i$ or there is an infinite reduced word X such that $|y(i) \wedge X|_* \to \infty$. For each $g \in G$ let,

$$n_g = \liminf_i |g \wedge y(i)|_*.$$

Suppose that the set $\{n_g | g \in G\}$ is bounded. Then let $n = \max\{n_g | g \in G\}$. Then we may find a $g_0 \in G$ such that $n_{g_0} = n$. By choosing such a g_0 of minimal free product length we also get that g_0 is a subword of all but finitely many $y(i)$. Suppose that there is a g_1 which is also subword of infinitely many $y(i)$ and that $|g_1|_* > |g_0|_*$, which in particular means that g_0 is a proper subword of g_1. Since $n_{g_0} = \max\{n_g | g \in G\}$ we must also have that there are infinitely many $y(i)$ which do not have g_1 as a subword.

Hence we may find infinitely many pairs of integers i, j such that $y(i) \wedge y(j) = g_0 \wedge g_1 = g_0$. Thus for these pairs we get, by 7.0.9 that

$$(y(i).y(j)) \leq |g_1|.$$

However, this is in contradiction to the fact that $\{y(i)\}$ is a sequence which converges to infinity. Thus by passing to a subsequence if neccessary we may write,

$$y(i) = g_0 a(i) z(i)$$

where the product is reduced as written and each $a(i)$ lies in a single factor and they are all distinct. We note that if $a(i)$ and $a(j)$ lie in distinct factors then

$$(y(i), y(j)) = |g_0|.$$

Thus by passing to a further subsequence we may assume that all the $a(i)$ lie in the same factor, G_l. Hence,

$$(y(i).y(j)) = |g_0| + (a(i).a(j)) \leq |g_0 a(i)| \leq |g_0 a(i) {g_0}^{-1}|.$$

It is then clear that $\{g_0 a(i) {g_0}^{-1}\}$ is a sequence which converges to infinity and is equivalent to $\{y(i)\}$, since by 7.0.9, $(g_0 a(i) {g_0}^{-1}.y(i)) \geq |g_0 a(i)|$. Thus if $\{n_g | g \in G\}$ is bounded we see that $\{y(i)\}$ is equivalent to a sequence which lies entirely in G_l^w for some l, w.

On the other hand we will show that if $\{n_g | g \in G\}$ is not bounded then $\{y(i)\}$ will be an infinite reduced word. (This is to say that there is an infinite reduced word X such that $\{y(i)\}$ is equivalent to $\{X_i\}$.) By the hypotheses we may find a sequence of elements $w(i) \in G$ such that $\{n_{w(i)}\}$ is a strictly increasing sequence of positive integers. By choosing, as before, the elements $w(i)$ to be of minimal length we may assume that each $w(i)$ is a subword of all but finitely many of the $\{y(i)\}$. Since we also have that $|w(i+1)|_* > |w(i)|_*$, we get that each $w(i)$ is a strict subword of $w(i+1)$. Hence the sequence $w(i)$ defines an infinite reduced word X. In fact each $w(i)$ is a subword of X and hence $|y(i) \wedge X|_*$ is unbounded. Thus by 7.0.9 we are done.

We now consider the action of an automorphism α on ∂G. For those points in G_∞, the infinite reduced words, we may apply 5.1.14 to show that these must all be either attracting, repelling or the limit of a sequence of fixed elements. By the above theorem other points of ∂G lie in $\partial G_i{}^w$ for some integer i and some $w \in G$. Clearly, in this case G_i must be a non infinite cyclic factor of G, since otherwise we are in fact looking at an infinite reduced word. Also, if $w^{-1}w' \in G_i$, then $\partial G_i{}^w = \partial G_i{}^{w'}$. So we assume, when considering points in $\partial G_i{}^w$ that w does not begin with an element from G_i.

Now we recall that if G_j is a non infinite cyclic factor of G and α an automorphism of G then G_j is sent by α to a conjugate of possibly a different non cyclic factor. Thus if $a \in G_j$, we write $\alpha a = u_j{}^{-1}\alpha_j(a)u_j$ where this product is reduced as written. We then get the following,

Proposition 7.0.12 *If α fixes a point of $\partial G_j{}^w$, then $\alpha w = wxu_j$ (not neccessarily reduced), where $|x|$ is bounded independently of w.*

(We note that a consequence of this proposition is that $\alpha(G_j{}^w) = G_j{}^w$.)

Proof: Consider the sequence $\{a_i{}^w\} \subset G_j{}^w$ and suppose that $\{a_i{}^w\} \sim \{\alpha a_i{}^w\}$. We write $\alpha a_i{}^w = b_i{}^v$, where b_i is in the same factor as $\alpha_j(a_i)$ and the product is reduced as written. (Note that if we anticipate the result we see that b_i may not be $\alpha_j(a_i)$ but an internal conjugate of it.) First we claim that $v = w$, since otherwise we apply 7.0.9 to get,

$$(a_i{}^w . b_i{}^v) \leq \max\{|v|, |w|\}$$

which contradicts our assumptions. We get the same inequality if a_i and b_i lie in different non cyclic factors. Combining these we see that $w^{-1}\alpha(w)u_j{}^{-1} \in G_j$. Thus

$\alpha w = wxu_j$ where $x \in G_j$. If this is reduced we apply 5.1.3 to get the result. Otherwise since we asssume that w does not end with an element from G_j and u_j does not begin with one, we must have that $x = 1$.

Suppose that we attempted to count the number of fixed points of α on the boundary. It is fairly natural to consider the equivalence of fixed points under the action of the fixed subgroup. By the above proposition, under this equivalence, we only need to consider finitely many of the 'factor boundaries' $\partial {G_j}^w$. This is because, if we have a fixed point in $\partial {G_j}^w$ then $\alpha w = wxu_j$, and there are finitely many possibilities for x. If we also have a w' such that $\alpha w' = w'xu_j$ then $w'w^{-1} \in \mathrm{Fix}\alpha$. This reduces the problem of studying the fixed points of α on the boundary of a free product, to that of studying the possible actions on the boundaries of indecomposable groups. It seems likely that this problem is amenable, especially in the torsion free case as here the indecomposable groups will be one-ended and hence have connected and locally connected boundary

We finally give an easy corollary of the above results for some virtually free groups. In the finitely generated case they are actually hyperbolic since they contain a hyperbolic subgroup of finite index. We consider finitely generated groups which can be written as the free product of finite and infinite cyclic groups - in other words, virtually free groups which are free products. We note that it is clear that the boundary of a finite group is empty. We thus have the following.

Theorem 3.7 *Let G be a finitely generated free product of finite and infinite cyclic groups and $\alpha \in \mathrm{Aut}G$. Then regarding G as a hyperbolic group we have that the points on the boundary fixed by α are either:*

(i) Attracting

(ii) Repelling

(iii) The limit of fixed elements of G.

Suppose $G = G_1 * G_2 * \ldots G_k$ where each G_i is either finite or infinite cyclic. Set $r(G_i)$ to be the maximal number of generators for any subgroup of G_i and $r(G)$ to be the maximum of these. Set $a(\alpha)$ to be the number of equivalence classes of attracting fixed points on the boundary of G for α. The equivalence is given by the action of the fixed subgroup. We then have the following:

Theorem 5.8 $1/r(G)$ *(# generators of* $\mathrm{Fix}\alpha$*)* $+\ a(\alpha)/2 \leq k$.

Proof: This is merely a corollary of 5.2.17.

Bibliography

[ABC+91] J. M. Alonso, T. Brady, D. Cooper, T.Delzant, V.Ferlini, M.Lustig, M.Mihalik, M.Shapiro, and H.Short, *Notes on word hyperbolic groups*, Group Theory from a Geometrical Viewpoint (H.Short, ed.), World Scientific, 1991, pp. 3–63.

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

[BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math.(2) **135** (1992), 1–51.

[BL36] R. Baer and F. Levi, *Freie produkte und ihre untergruppen*, Compositio Math. **3** (1936), 391–398.

[CDP90] M. Coornaert, T. Delzant, and A. Papdopoulos, *: les groupes hyperbolique de gromov*, Lecture notes in mathematics, no. 1441, Springer-Verlag, 1990.

[CL] M.M. Cohen and M. Lustig, *Dehn twist automorphisms*, preprint.

[CL95] M.M. Cohen and M. Lustig, *Very small actions on $\mathbb{R}$-trees and dehn twist automorphisms*, Topology **34** (1995), no. 3, 575–617.

[CM87] M. Culler and J. Morgan, *Group actions on $\mathbb{R}$-trees*, Proc. London Math. Soc. **55** (1987), 571–604.

[CM89] M.M. Cohen and M.Lustig, *On the dynamics and the fixed subgroup of a free group automorphism*, Inv. Math. **96** (1989), 613–638.

[Coh89] D. E. Cohen, *Combinatorial group theory: a topological approach*, London Math. Soc. Stud. Texts, no. 14, CUP, 1989.

[Coo87] D. Cooper, *Automorphisms of free groups have finitely generated fixed point sets*, Jour. Alg. **111** (1987), 453–456.

[CT88] D. J. Collins and E. C. Turner, *Free product fixed points*, J. London Math. Soc. (2) **38** (1988), 67–76.

[CT93] D. J. Collins and E. C. Turner, *An automorphism of a free group of finite rank with maximal rank fixed point subgroup fixes a primitive element*, Journal of Pure and Applied Algebra **88** (1993), no. 1–3, 43–49.

[CT94] D. J. Collins and E. C. Turner, *Efficient representatives for automorphisms of free products*, Michigan Math. Journal **41** (1994), 443–464.

[CT96] D. J. Collins and E. C. Turner, *All automorphisms of free groups with maximal rank fixed subgroups*, Math. Proc. Camb. Phil. Soc. **119** (1996), no. 4, 615–630.

[CV86] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 91–119.

[DS75] J.L. Dyer and G.P. Scott, *Periodic automorphisms of free groups*, Comm. Algebra **3** (1975), 195–201.

[GJLL] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, preprint.

[Gol90] R. Goldstein, *Bounded cancellation of automorphisms of free products*, Topology and Combinatorial Group Theory, Lecture notes in mathematics, Springer, 1990, pp. 81–89.

[Gro88] M. Gromov, *Hyperbolic groups*, Essays in group theory (S.M. Gersten, ed.), MSRI, vol. 8, Springer-Verlag, 1988, pp. 75–263.

[KPS73] A. Karrass, A. Pietrowski, and D. Solitar, *Finite and infinite cyclic extensions of free groups*, J. Austr. Math. Soc. **16** (1973).

[Kur34] A. G. Kuroš, *Die untergruppen der freien produkte von belieben gruppen*, Math. Ann. **109** (1934), 647–660.

[Lev95] D. Gaboriau G. Levitt, *The rank of actions on* $\mathbb{R}$*-trees*, Ann. Sci. Ecole. Norm.Sup. (4) **28** (1995), no. 5, 549–570.

[Los96] J.E. Los, *On the conjugacy problem for automorphisms of free groups*, Topology **35** (1996), no. 3, 779–806.

[Sko] R. Skora, *Deformations of length functions in groups*, preprint.