Automorphisms of free groups and train tracks

Gilbert Levitt

Laboratoire de Mathématiques Nicolas Oresme
Université de Caen

September 2008

F_n is the free (non-abelian) group of rank n.
For topologists: the fundamental group of a finite graph, or a compact surface with boundary.

In $F_3 = F(a, b, c)$:

Product: $(a^2b^{-1}a^{-3}c^2)\,(c^{-2}acb^2) = a^2b^{-1}a^{-2}cb^2$

Inverse: $(a^2b^{-1}a^{-3}c^2)^{-1} = c^{-2}a^3ba^{-2}$

We want to understand automorphisms of free groups.

- Individual automorphisms
- The global structure of $Aut(F_n)$ and $Out(F_n) = Aut(F_n)/Inn(F_n)$

A remark for topologists: the natural object associated to a homeomorphism $f : X \to X$ is $f_* \in Out(\pi_1(X))$.

Dehn-Nielsen: $MCG(\Sigma_g) \simeq Out^+(\pi_1(\Sigma_g))$, with Σ_g the closed orientable surface of genus g.

Two famous results:

- the Scott conjecture (Bestvina-Handel)
- the Tits alternative (Bestvina-Feighn-Handel)

Train tracks (not quite the same as Thurston's) are the basic tool used in the proofs.

The Scott conjecture

$\alpha \in Aut(F_n)$
Its fixed subgroup is $Fix(\alpha) = \{g \in F_n \mid \alpha(g) = g\}$.

Theorem (Bestvina-Handel 1992)

$Fix(\alpha)$ has rank at most n.

In case you're wondering...

Fact

A subgroup of F_n is free (Nielsen-Schreier), but its rank may be $> n$ (even infinite).

Example

The commutator subgroup $[F_2, F_2] = \ker(F_2 \to \mathbb{Z}^2)$ has infinite rank.

The Scott conjecture

$\alpha \in Aut(F_n)$
Its fixed subgroup is $Fix(\alpha) = \{g \in F_n \mid \alpha(g) = g\}$.

Theorem (Bestvina-Handel 1992)

$Fix(\alpha)$ has rank at most n.

Example

$a \mapsto a$
$b \mapsto ba$

$Fix(\alpha) = \langle a, bab^{-1} \rangle$
Check: $bab^{-1} \mapsto ba\ a\ a^{-1}b^{-1} = bab^{-1}$

Similar results:

- Free abelian groups
  If $\alpha \in GL(n, \mathbb{Z}) = Aut(\mathbb{Z}^n)$, then $Fix(\alpha) \simeq \mathbb{Z}^p$ with $p \leq n$.
- Surface groups
  Let $G = \pi_1(\Sigma_g)$, and $\alpha \in Aut(G)$ be induced by a homeomorphism $f : \Sigma_g \to \Sigma_g$.

  Nielsen: $Fix(\alpha)$ is cyclic or is the fundamental group of an incompressible subsurface.
  In particular, $Fix(\alpha)$ is free of rank $\leq 2g$ if $\alpha \neq id$.

  Dehn-Nielsen: every automorphism of G is induced by a homeomorphism of Σ_g.

- Nielsen's theorem about $Fix(\alpha)$ also applies to homeomorphisms of a compact surface Σ with boundary (Jaco-Shalen 1977).

  In this case $\pi_1(\Sigma) \simeq F_n$ with $n = 1 - \chi(\Sigma)$, and $Fix(\alpha)$ has rank $\leq n$.

  Example: Dehn twist on a punctured torus

The Tits alternative

Theorem (Tits 1972)

If $H \subset GL(n, \mathbb{C})$, then either H contains F_2 or H is virtually solvable.

virtually solvable = has a solvable subgroup of finite index

Theorem (Ivanov, Birman-Lubotzky-McCarthy 1983)

If $H \subset MCG(\Sigma_g)$, then either H contains F_2 or H is virtually abelian.

The proof uses Nielsen-Thurston theory: the classification of individual elements of $MCG(\Sigma_g)$.

Theorem (Bestvina-Feighn-Handel 2000)

If $H \subset Out(F_n)$, then either H contains F_2 or H is virtually abelian.

- The proof is muuuch harder, because there is no Jordan form or Nielsen-Thurston classification for automorphisms of F_n.
- The best replacement (so far) is Train tracks (relative, improved, completely split, β-train tracks)
- Much of the work on automorphisms of free groups is motivated by analogies with linear groups and mapping class groups.

Geometric automorphisms

Definition

$\Phi \in Out(F_n)$ is *geometric* if $\Phi = f_*$, where Σ is a compact surface with $\pi_1(\Sigma) \simeq F_n$ and $f : \Sigma \to \Sigma$ is a homeomorphism.

Example (n=2)

Every $\Phi \in Out(F_2)$ is geometric, realized on a punctured torus.
$Out(F_2) = Out(\mathbb{Z}^2) = GL(2, \mathbb{Z}) = \widehat{MCG}(T^2 \setminus \{*\})$.

Up to isotopy, f is

- of finite order
- (a root of) a Dehn twist
- pseudo-Anosov.

A simple pseudo-Anosov map

$A = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$ has eigenvalues $\frac{1}{\lambda} < 1 < \lambda$. It acts on $\mathbb{R}^2$ with two invariant foliations (by lines parallel to the eigenvectors).

If K is compact, then $A^N(K)$ looks like a very thin segment of length $\sim \lambda^N$ for N large.

A induces a homeomorphism $\bar{A}$ of the torus $T^2 = \mathbb{R}^2/\mathbb{Z}^2$. Images by $\bar{A}^N$ are almost dense in T^2.

Removing the fixed point and adding a boundary yields a pseudo-Anosov homeomorphism on a punctured torus.

Nielsen-Thurston theory

Nielsen-Thurston in general: a power of $f : \Sigma \to \Sigma$ is isotopic to a homeomorphism f' with the following property:
it preserves a decomposition of Σ into subsurfaces such that on each subsurface the restriction of f' is

- the identity
- a Dehn twist on an annulus
- pseudo-Anosov.

For $n \geq 3$, not every automorphism of F_n is geometric.

Example

$$a \mapsto a \quad b \mapsto ba \quad c \mapsto cb$$

is not geometric.

Why?

Iterate c:
$c \mapsto cb \mapsto cbba \mapsto cbbaba^2 \mapsto cbbaba^2ba^3 \mapsto \dots$ grows quadratically.

If $\Phi = f_*$, the growth of (the conjugacy class of) c is the growth of a closed curve under iteration of f, hence linear or exponential.

Automorphisms of finite order

If Γ is a finite graph with $\pi_1(\Gamma) \simeq F_n$, any symmetry $f : \Gamma \to \Gamma$ induces $f_* \in Out(F_n)$ of finite order.

Example

If Γ is a "Theta graph" Θ, then $Sym(\Gamma)$ is a subgroup of order 12 of $Out(F_2) \simeq GL(2, \mathbb{Z})$ (which does not lift to $Aut(F_2)$).

Theorem (Zimmermann, Culler, Khramtsov)

Every finite subgroup $F \subset Out(F_n)$ is realized on a finite graph Γ.

Corollary

For a given n, there are finitely many conjugacy classes of finite subgroups of $Out(F_n)$.

Corollary (Dyer-Scott)

If $\alpha \in Aut(F_n)$ has finite order, then $Fix(\alpha)$ is cyclic or is a free factor of F_n.

In particular, $Fix(\alpha)$ has rank $\leq n$.

Idea of proof: $Fix(\alpha)$ is a free factor because it is the fundamental group of a subgraph of Γ.

Towards train tracks

Every $\Phi \in Out(F_n)$ is represented by a map f (a homotopy equivalence, but not a homeomorphism) on a rose (wedge of circles).
The petals are labelled by the generators of F_n.

In all previous examples, the automorphism α was positive: there were no a^{-1}, b^{-1}, c^{-1} in images of generators.
Consequence: no cancellation when computing $\alpha^p(a)$, $\alpha^p(b)$, $\alpha^p(c)$.
Geometrically: no cancellation in $f^p(edge)$.

But cancellation occurs in iterated images of other words.

Example

$a \mapsto aba$
$b \mapsto ba$
$aba^{-1}b^{-1} \mapsto ab\ a\ ba\ a^{-1}b^{-1}a^{-1}\ a^{-1}b^{-1} = aba^{-1}b^{-1}$

This is unavoidable. But we would like to have the non-cancellation property on generators (edges) for general automorphisms.
This is the basic idea behind train tracks: represent Φ by a map on a graph whose cancellation is as controlled as possible.

Applying this idea to surface homeomorphisms leads to a quick algorithmic proof of Nielsen-Thurston (and an invariant train track à la Thurston in the pseudo-Anosov case).

An example

Example

$a \mapsto c$
$b \mapsto c^{-1}a$
$c \mapsto c^{-1}b \mapsto b^{-1}c\ c^{-1}\ a$

$c^{-1}b$ is an illegal turn.
The previous example had plenty of illegal turns, for instance ab^{-1}.
The problem is that $c^{-1}b$ appears in the image of c.

Answer to the problem: fold.

- We get a map f_1 on a graph Γ_1 with 4 edges a, b_1, c_1, d
  $a \mapsto dc_1$
  $b_1 \mapsto a$
  $c_1 \mapsto db_1$
  $d \mapsto c^{-1}d^{-1}$
- Γ_1 has 2 vertices, no preferred basepoint.
  f_1 represents Φ via $\pi_1(\Gamma_1) \simeq F_3$.
- Are we done? Or is there cancellation in iterated images of edges?
- The image of a contains the illegal turn dc_1, so we fold again.

- We get a map f_2 on a graph Γ_2 with 5 edges a, b_1, c_2, d_2, e.
  $a \mapsto d_2c_2$
  $b_1 \mapsto a$
  $c_2 \mapsto b_1$
  $d_2 \mapsto c_2^{-1}e^{-1}$
  $e \mapsto d_2e^{-1}$
- After collapsing e, we get a train track map f'.
  $a' \mapsto d'c'$
  $b' \mapsto d'^{-1}a'$
  $c' \mapsto b'$
  $d' \mapsto c'^{-1}$

Train track maps

Definition

$f : \Gamma \to \Gamma$ is a train track map if:

1. Γ is a finite graph (with no vertex of valence 1)
2. f maps a vertex to a vertex, an edge to a non-trivial edge-path
3. for $p \geq 1$, there is no cancellation in $f^p(edge)$

f represents $\Phi \in Out(F_n)$ if there is an isomorphism $j : \pi_1(\Gamma) \simeq F_n$ such that $j \circ f_* = \Phi \circ j$ (everything is up to inner automorphisms).

Example

f' as above, the obvious map on the rose for positive automorphisms.

The important condition is 3 (no cancellation).
Maps satisfying 1 and 2 are called topological representatives.

f, f_1, f_2 were topological representatives.
Every Φ has a topological representative (on a rose).

They don't always exist.
What always exists is a relative train track.

One tries to construct a train track by folding and cleaning up a topological representative (as before). Why should it stop?

Key ideas (Bestvina-Handel):

- A topological representative has a transition matrix, with spectral radius $\lambda \geq 1$.
- λ cannot increase during the process.
- λ belongs to a discrete set, so can only take finitely many values.
- When λ is minimal, one gets a train track map or a reduction (in a sense to be defined).

In particular:

Theorem

Any irreducible automorphism has a train track representative.

Reducible automorphisms have relative train track representatives.

Surface homeomorphisms

These ideas apply to surface homeomorphisms (creating a puncture and using a spine).

- A reduction gives a periodic subsurface. The map is reducible (in Thurston's sense).
- In the irreducible case, the map has finite order (if $\lambda = 1$), is pseudo-Anosov (if $\lambda > 1$).

The transition matrix

A topological representative f has a transition matrix M, which records how many times the image of the i-th edge crosses the j-th edge (in either direction).

$$f : \begin{cases} a \mapsto c \\ b \mapsto c^{-1}a \\ c \mapsto c^{-1}b \end{cases} \qquad M = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

$$f' : \begin{cases} a' \mapsto d'c' \\ b' \mapsto d'^{-1}a' \\ c' \mapsto b' \\ d' \mapsto c'^{-1} \end{cases} \qquad M' = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \end{pmatrix}$$

The ℓ^1-norm of a column is the length of the image of the corresponding edge.

Perron-Frobenius theory

Let M be a non-negative integral square matrix.

- M has an eigenvalue $\lambda \in \{0\} \cup [1, \infty)$, such that every eigenvalue μ satisfies $|\mu| \leq \lambda$.
- If v is a positive vector, then $\|M^p(v)\|^{1/p} \to \lambda$ as $p \to \infty$ (exponential growth rate).

λ is called the Perron-Frobenius eigenvalue, or largest eigenvalue, or spectral radius.

Given d, the set of Perron-Frobenius eigenvalues of matrices of size d is discrete

because the coefficients of the characteristic polynomial may be bounded in terms of λ and d.

λ does not increase

$$f : \begin{cases} a \mapsto c \\ b \mapsto c^{-1}a \\ c \mapsto c^{-1}b \end{cases} \qquad M = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix} \qquad \lambda$$

$$f' : \begin{cases} a' \mapsto d'c' \\ b' \mapsto d'^{-1}a' \\ c' \mapsto b' \\ d' \mapsto c'^{-1} \end{cases} \qquad M' = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix} \qquad \lambda'$$

Claim: λ does not increase during the "fold and clean up" process. In particular, $\lambda' \leq \lambda$.

Train tracks minimize λ

Suppose that $f : \Gamma \to \Gamma$ and $f' : \Gamma' \to \Gamma'$ both represent Φ, with f' a train track map. Then $\lambda' \leq \lambda$.
In particular, any two train track maps give rise to the same λ.

Idea (on the example).

- The loops a and a' represent the same element of F_3, so have similar growth under iteration of f (resp. f').
- λ represents the growth rate of a if we don't do the cancellations, so is an upper bound for the actual growth rate.
- λ' represents the actual growth rate of a' because there is no cancellation. So $\lambda' \leq \lambda$.

In general:

- Let e_i be the i-th edge of Γ. The ℓ^1-norm of the i-th column of M^p is the length of $f^p(e_i)$, before cancellation.
- It has exponential growth rate at most λ, and there is an edge e_{i_0} for which it equals λ.
- In particular, the exponential growth rate of any $|f_\sharp^p(e_i)|$ (after cancellation) is at most λ. But it may be smaller for every i, even i_0, because of cancellation...
  unless f is a train track map,

Gilbert Levitt Automorphisms of free groups and train tracks

- Since f' is a train track map, there is an edge $e' \subset \Gamma'$ such that the actual length $|f'^{p}_{\sharp}(e')|$(after cancellation) has growth rate λ'.
- In fact, there is a loop γ' such that $|f'^{p}_{\sharp}(\gamma')|$ has growth rate λ'. (In the case of f_3, we can take γ' to be the loop $a \subset \Gamma_3$).
- γ' represents an element (or conjugacy class) g of F_3 whose growth rate under Φ is that of γ' under f', namely λ'.
- But g is represented by a loop $\gamma \subset \Gamma$, and the growth rate of g is bounded by the growth rate of γ under f, hence by λ. Thus $\lambda' \leq \lambda$. □

Back to the "fold and clean up" process.

If f is arbitrary and f' is a train track map, $\lambda - \lambda'$ measures the amount of cancellation that takes place in f. Since the process is designed to reduce cancellation, λ cannot increase.

By discreteness, we reach a point where λ cannot change anymore.

But this does not mean that we have a train track map.

Example

$$a \mapsto aba$$
$$b \mapsto ba$$
$$c \mapsto caba^{-1}b^{-1}$$

View this as a map f on the rose. The matrix $M = \left(\begin{array}{cc|c} 2 & 1 & 2 \\ 1 & 1 & 2 \\ \hline 0 & 0 & 1 \end{array}\right)$ is block-triangular. It has the same Perron-Frobenius eigenvalue as $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$.

This eigenvalue is the exponential growth rate of a, so it cannot be decreased. There is no train track representative: the cancellation in $f^2(c)$ won't go away.

The subgraph consisting of the first two petals is f-invariant. Algebraically, we get an invariant free factor $\langle a, b\rangle$.

Definition

An automorphism is reducible if there is a periodic free factor (different from F_n and the trivial group).

The process yields a map with λ minimal. Such a map is a train track map or shows a reduction (a non-trivial periodic subgraph).

We get:

Theorem

Any irreducible automorphism admits a train track representative.

These ideas apply to surface homeomorphisms (creating a puncture and using a spine).

- A reduction gives a periodic subsurface. The complementary subsurface is also periodic, and the map is reducible (in Thurston's sense).
- In the irreducible case, the map has finite order (if $\lambda = 1$), is pseudo-Anosov with dilation factor λ (if $\lambda > 1$).

Reducible automorphisms

$$\begin{aligned} a &\mapsto aba \\ b &\mapsto ba \\ c &\mapsto caba^{-1}b^{-1} \end{aligned}$$ is reducible.

$\langle a, b \rangle$ is an invariant free factor. But there is no invariant complementary free factor: one cannot write $F_3 = G_1 * G_2$ with G_1 and G_2 invariant.

This is a key difference with the surface case: if a surface homeomorphism is reducible, Σ is the union of two invariant subsurfaces.
Unlike in the surface case, it is not true that every automorphism can be built out of irreducible ones. This is a major problem.

Compare representation theory: a *unitary* representation is a direct sum of irreducible representations (if V is invariant, so is $V^{\perp}$). But this is not true in general.

Towards relative train tracks

In other words, we cannot stop once we've found a reduction, so let's keep going.

If there is a reduction, the matrix becomes block-triangular after raising f to a power and reordering edges.

Do this so as to get as many blocks as possible. For instance, replace $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ by its square $\left(\begin{array}{c|c} 1 & 0 \\ \hline 0 & 1 \end{array}\right)$.

Strata

$$a \mapsto aba \qquad b \mapsto ba \qquad c \mapsto caba^{-1}b^{-1}$$

$$\left(\begin{array}{cc|c} 2 & 1 & 2 \\ 1 & 1 & 2 \\ \hline 0 & 0 & 1 \end{array}\right)$$

- The block-triangular structure partitions the set of edges into strata (one stratum per block). In the example, there are two strata: one contains a and b, the other c.
- The set of strata is linearly ordered: the c stratum is higher than the (a, b)-stratum. The image of an edge cannot contain edges from a higher stratum.

$$a \mapsto a \qquad b \mapsto ba \qquad c \mapsto cb$$

has 3 strata

What about cancellation?

In the first example, there is no cancellation in the iterated images of a and b.

There is some in the images of c:
$c \mapsto caba^{-1}b^{-1} \mapsto caba^{-1}b^{-1}ababaa^{-1}b^{-1}a^{-1}a^{-1}b^{-1}$.

But all cancellation takes place in the stratum below c.

Relative train track maps (slightly simplified definition)

Definition

$f : \Gamma \to \Gamma$ is a relative train track map if:

1. Γ is a finite graph, with edges grouped into strata $H_1, H_2, \ldots, H_r$
2. f maps a vertex to a vertex, an edge to a non-trivial edge-path
3. the image of an edge in H_j contains no edge from a higher stratum H_k with $k > j$
4. given an edge $e \subset H_j$ and $p \geq 1$, all cancellation in $f^p(e)$ takes place in lower strata H_k with $k < j$

Theorem (Bestvina-Handel)

Given $\Phi \in Out(F_n)$, some power of Φ has a relative train track representative.

The proof uses moves as in the absolute case, but one has to decrease all Perron-Frobenius eigenvalues...

Having a relative train track map makes it possible to prove results by induction, from the lowest stratum up.

An application

Train tracks are a basic tool in the study of automorphisms of free groups. Here is a baby application.

Theorem

Given $\alpha \in Aut(F_n)$ and $g \in F_n$, the length of $\alpha^p(g)$ grows like $\lambda^p p^m$ for some $\lambda \geq 1$ and $m \in \mathbb{N}$, as $p \to \infty$.

"Grows like" means the ratio is bounded away from 0 and ∞.

The theorem is better stated in terms of conjugacy classes.

I do not know how to prove the theorem without improved relative train tracks.

(Improved train tracks were introduced by Bestvina-Feighn-Handel in the first of their Tits papers.)

The problem, of course, is cancellation: if there is no cancellation, $\alpha^p(g)$ grows like entries of M^p, as if the group were abelian.

The theorem is true in free abelian groups. What can be said in a general finitely generated (finitely presented) group?

Examples

We have seen that the growth is linear or exponential if α is geometric (represented by a surface homeomorphism). We have seen quadratic growth. Here is an example with growth $p\lambda^p$.

Start with
$a \mapsto aba$
$b \mapsto ba$
$c \mapsto cdc$
$d \mapsto dc$

There are two strata, whose order is arbitrary. All four generators have the same growth λ^p under iteration of the map.

$a \mapsto aba$
$b \mapsto ba$
$c \mapsto cdc\,a$
$d \mapsto dc$

c and d grow like $p\lambda^p$.
More generally, one may achieve growth $p^k\lambda^p$ in F_{2k+2}.

It is important for this example that both strata have the same eigenvalue $\lambda_{ab} = \lambda_{cd}$.

Otherwise c grows like λ_{ab}^p if $\lambda_{ab} > \lambda_{cd}$, like λ_{cd}^p if $\lambda_{ab} < \lambda_{cd}$.

Proof of the growth theorem

If f is a relative train track map, the transition matrix M is block-triangular. Each diagonal block M_j corresponds to a stratum H_j.
Since we have created as many blocks as possible, there are only two possibilities (three, actually):

- M_j is a positive matrix of size ≥ 2 with eigenvalue $\lambda > 1$: exponential stratum
- M_j is a single 1: polynomial (NEG, parabolic) stratum. It consists of a single edge.

Splitting of paths

Definition

$\gamma = \gamma_1.\gamma_2$ is a splitting if $f_\sharp^p(\gamma) = f_\sharp^p(\gamma_1).f_\sharp^p(\gamma_2)$ for every $p \geq 1$.

$f_\sharp$ denotes the tightened image (after cancellations).

In words: γ_1 and γ_2 don't interact when we apply powers of f.

In particular, if γ_1 or γ_2 becomes long, so does γ. If we know the growth of γ_1 and γ_2, we know that of γ.

Selected properties of improved relative train track maps

Exponential strata

- If e belongs to an exponential stratum H_j, then $f(e)$ splits into edges of H_j and lower paths δ_i. (lower = contained in strata H_k with $k < j$)
- No $f_\sharp^p(\delta_i)$ is trivial.

It follows that $f_\sharp^p(e)$ splits into edges of H_j and lower paths $f_\sharp^q(\delta_i)$ with $q < p$.

The δ_i's appear in the image by f of an edge of H_j, so they belong to a finite set.

Polynomial strata

- If e is a polynomial stratum, then $f(e)$ splits as $f(e) = e.u$ with u a loop in strata lower than e.
- If u is non-trivial, no $f_\sharp^p(u)$ is trivial.

It follows that $f_\sharp^2(e)$ splits as $e.u.f_\sharp(u)$, etc.

This implies splitting for images of arbitrary paths.

Splitting lemma

Given an edge-path γ of height j, there exists p_0 such that $f_\sharp^{p_0}(\gamma)$ splits into:

- Edges of H_j.
- Paths lower than H_j.
- Paths whose growth is at most linear (Nielsen paths, exceptional paths).

(The height of γ is the index of the highest stratum that it meets.)

Nielsen path

$$a \mapsto aba$$
$$b \mapsto ba$$

$aba^{-1}b^{-1}$ is a Nielsen path

Exceptional path

$$a \mapsto a$$
$$b \mapsto ba^2$$
$$c \mapsto ca$$

bc^{-1} is an exceptional path

These paths don't split, but we don't care since we know their growth.

To show that elements (or conjugacy classes) of F_n grow like some $\lambda^p p^m$ under α, we show that any path (loop) γ has this type of growth under f. The proof is by induction on the height of γ.

Since we understand the growth of lower paths by induction, the splitting lemma implies that we only have to understand the growth of an edge e. There are two cases.

- If e is a polynomial stratum H_j, then $f(e) = e.u$ with u of height $< j$, and $f_\sharp^p(e) = e.u.f_\sharp(u).f_\sharp^2(u)\dots f_\sharp^{p-1}(u)$. By induction, we know how u grows.
  - If u grows exponentially, e has the same growth.
  - If u grows as a polynomial of degree m, then e grows as a polynomial of degree $m+1$.

- If e is in an exponential stratum H_j, then $f_\sharp^p(e)$ consists of
  - edges of H_j
  - lower paths $f_\sharp^q(\delta_i)$, where $q < p$ and δ_i is a lower path which appears in $f(e')$ for some e' in H_j.

There are only finitely many types of δ_i's, and by induction we know how they grow.

When we apply powers of f to e, two things happen.

- The number of edges in H_j grows.
- Lower paths appear and grow.

In $f_\sharp^p(e)$, the number of edges in H_j is $\sim \lambda_j^p$, where λ_j is the eigenvalue of H_j.
Passing from $f_\sharp^q(e)$ to $f_\sharp^{q+1}(e)$ creates $\sim \lambda_j^q$ lower paths δ_i.

The growth of e depends on comparing the growth of the δ_i's to λ_j^p.

- If every δ_i has growth $< \lambda_j^p$, then e grows like λ_j^p (as if there were no δ_i's).
- If some δ_i has exponential growth rate $> \lambda_j$, then e grows like the fastest-growing δ_i (what matters is how the δ_i's grow, not how fast they are created).
- If the fastest growth of a δ_i is $\lambda_j^p p^m$, then e grows like $\lambda_j^p p^{m+1}$ (the rate of growth and the rate of creation both matter). □

Attracting laminations

Given $\Phi \in Out(F_n)$, (relative) train track representatives are not unique. But they have certain features in common.

In particular, the number of exponential strata and their λ's only depend on Φ. In fact, each exponential stratum gives rise to an attracting lamination, a set of unordered pairs of distinct elements of $\partial F_n \times \partial F_n$. These laminations are intrinsic.

The interpretation in the surface case is the following. Let f be a pseudo-Anosov homeomorphism on Σ_g. It has an attracting (stable) lamination. Lift it to the universal covering of Σ, the Poincaré disc. Every geodesic defines a pair of points on the circle at infinity, which is the boundary of $\pi_1(\Sigma_g)$.

Numerical bounds

There are bounds for certain numerical invariants attached to Φ.

We have seen that elements of F_n grow like $\lambda^p p^m$. The degree m is bounded by $\frac{n}{2} - 1$ if $\lambda > 1$, as in the example given above. If the growth is polynomial ($\lambda = 1$), the degree is bounded by n for elements of F_n, by $n - 1$ for conjugacy classes.

The number of exponential strata is at most $\left[\frac{3n-2}{4}\right]$ (L.-Lustig). This bound is achieved by geometric automorphisms induced by homeomorphisms whose Nielsen-Thurston reduction consists of pseudo-Anosov maps on once-punctured tori and four-punctured spheres.

References

- The original Bestvina-Handel paper: Annals of Maths, 1992
  Application to surfaces: BH, Topology, 1995
- Bogopolski's book: Introduction to group theory, EMS textbooks in mathematics 2008
- Surveys about automorphisms of free groups:
  Vogtmann: Automorphisms of free groups and outer space
  Bridson-Vogtmann: Automorphism groups of free groups, surface groups and free abelian groups
- Catherine Pfaff's notes from a short course I gave at MSRI, available from http://www.math.cornell.edu/~vogtmann/
- The growth theorem is proved in the appendix of arXiv:0801.4844.