A basis of the fixed point subgroup of an automorphism of a free group

Oleg Bogopolski and Olga Maslakova

GAGTA-6, Düsseldorf, 1.08.12
Workshop "Aut. of free groups", Barcelona, 10.11.12
Webinar "GT", NY, 6.12.12

Outline

1. Main Theorem
2. Names
3. A relative train track for α
4. Graph D_f for the relative train track $f : \Gamma \to \Gamma$
5. A procedure for construction of $CoRe(D_f)$
6. How to convert this procedure into an algorithm?
7. Cancelations in f-iterates of paths of Γ
8. μ-subgraphs in details

Scott Problem

Let F_n be the free group of finite rank n and let $\alpha \in \mathrm{Aut}(F_n)$. Define

$$\mathrm{Fix}(\alpha) = \{x \in F_n \mid \alpha(x) = x\}.$$

Rang problem of P. Scott (1978): $\mathrm{rk}(\mathrm{Fix}(\alpha)) \leqslant n$

M. Bestvina and M. Handel (1992): Yes

Main Theorem

Basis problem. Find an algorithm for computing a basis of $\mathrm{Fix}(\alpha)$.

It has been solved in three special cases:
– for positive automorphisms (Cohen and Lustig)
– for special irreducible automorphisms (Turner)
– for all automorphisms of F_2 (Bogopolski).

Theorem (O. Bogopolski, O. Maslakova, 2004-2012).
A basis of $\mathrm{Fix}(\alpha)$ is computable.
(see http://de.arxiv.org/abs/1204.6728)

Names

Dyer
Scott
Gersten
Goldstein
Turner
Cooper
Paulin
Thomas
Stallings
Bestvina
Handel
Gaboriau
Levitt
Cohen
Lustig
Sela
Dicks
Ventura
Brinkmann

Names

Dyer
Scott
Gersten

Goldstein
Turner

Cooper
Paulin
Thomas
Stallings

Bestvina
Handel

Gaboriau
Levitt

Cohen
Lustig

Sela
Dicks
Ventura

Brinkmann

Relative train tracks

Let Γ be a finite connected graph
and $f : \Gamma \to \Gamma$ be a homotopy equivalence s.t.
f maps vertices to vertices and edges to reduced edge-paths.

The map f is called a *relative train track* if ...

Relative train tracks

Let Γ be a finite connected graph
and $f : \Gamma \to \Gamma$ be a homotopy equivalence s.t.
f maps vertices to vertices and edges to reduced edge-paths.

The map f is called a *relative train track* if ...

To define this, we first need to define

- Turns in Γ (illegal and legal)
- Transition matrix
- Filtrations
- Stratums (exponential, polynomial, zero)

Turns

Let Γ be a finite connected graph
and $f : \Gamma \to \Gamma$ be a homotopy equivalence s.t.
f maps vertices to vertices and edges to reduced edge-paths.

A *turn:*

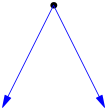

A *degenerate turn:*

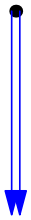

Turns

Let Γ be a finite connected graph
and $f : \Gamma \to \Gamma$ be a homotopy equivalence s.t.
f maps vertices to vertices and edges to reduced edge-paths.

A *turn:*

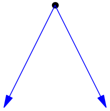

A *degenerate turn:*

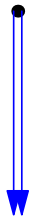

Differential of f.

$Df : \Gamma^1 \to \Gamma^1$, $\quad (Df)(E) =$ the first edge of $f(E)$.

Tf : Turns $\to$ Turns, $\quad (Tf)(E_1, E_2) = ((Df)(E_1), (Df)(E_2))$.

An illegal turn

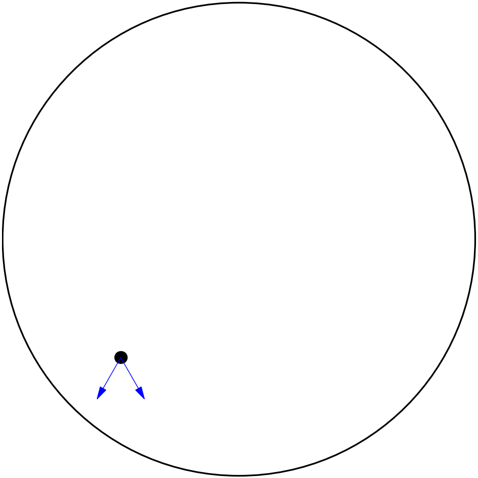

An illegal turn

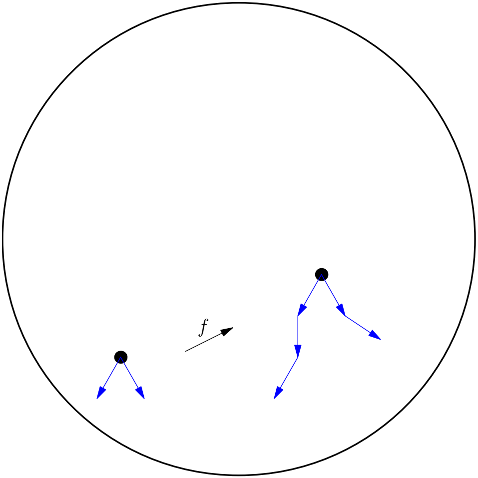

An illegal turn

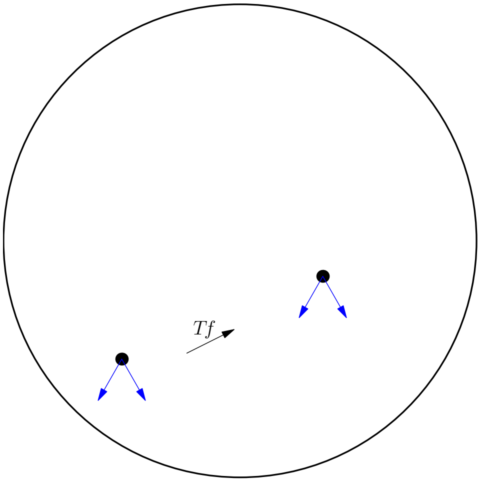

An illegal turn

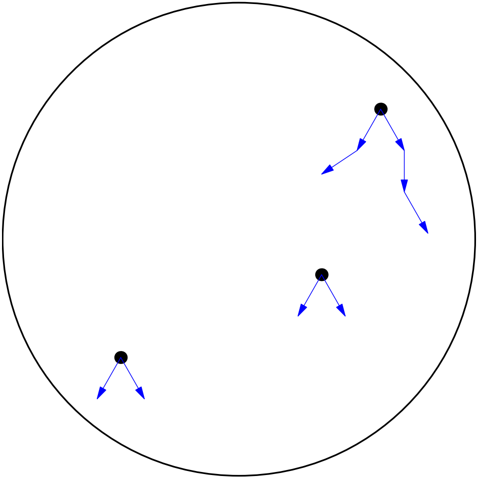

An illegal turn

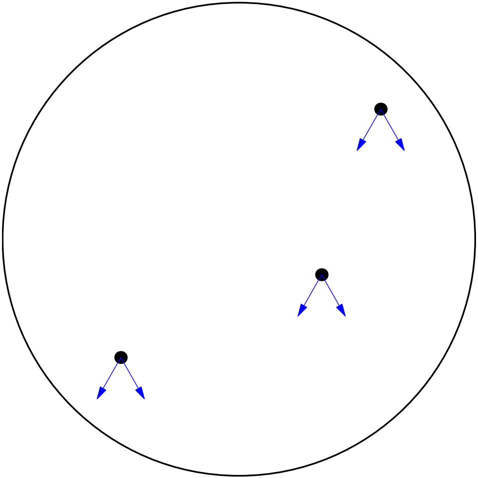

An illegal turn

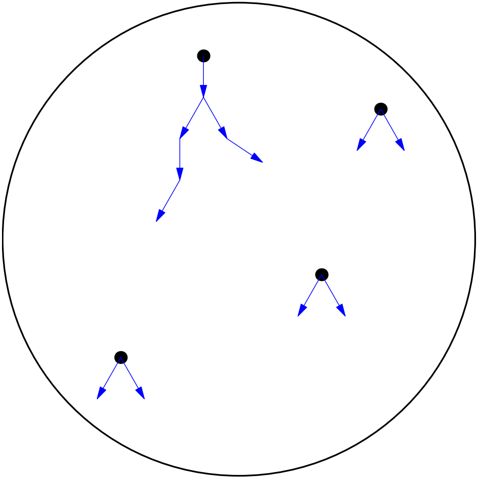

An illegal turn

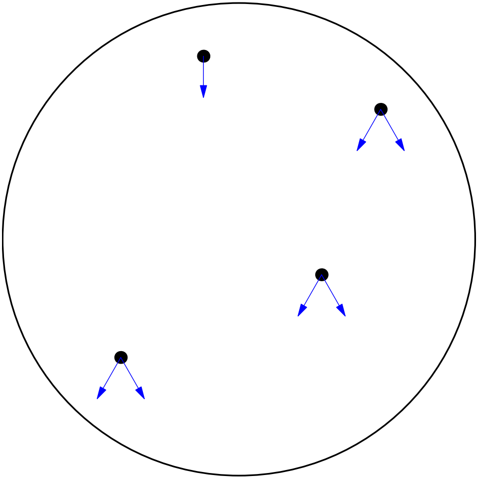

A turn (E_1, E_2) is called illegal
if $\exists n \geqslant 0$ such that the turn $(Tf)^n(E_1, E_2)$ is degenerate.

Legal turns and paths

A turn (E_1, E_2) is called legal
if $\forall n \geqslant 0$ the turn $(Tf)^n(E_1, E_2)$ is nondegenerate.

An edge-path p in Γ is called legal if each turn of p is legal.
Legal paths are reduced.

Legal turns and paths

A turn (E_1, E_2) is called legal
if $\forall n \geqslant 0$ the turn $(Tf)^n(E_1, E_2)$ is nondegenerate.

An edge-path p in Γ is called legal if each turn of p is legal.
Legal paths are reduced.

Claim. Suppose that $f(E)$ is legal for each edge E in Γ.
Then, for every legal path p in Γ, the path $f^k(p)$ is legal $\forall k \geqslant 1$.

Transition matrix of the map $f : \Gamma \to \Gamma$

From each pair of mutually inverse edges of Γ we choose one edge. Let $\{E_1, \dots, E_k\}$ be the set of chosen edges.

The *transition matrix* of the map $f : \Gamma \to \Gamma$ is the matrix $M(f)$ of size $k \times k$ such that the ij^{th} entry of $M(f)$ is equal to the total number of occurrences of E_i and $\overline{E_i}$ in the path $f(E_j)$.

Transition matrix of the map $f : \Gamma \to \Gamma$

From each pair of mutually inverse edges of Γ we choose one edge. Let $\{E_1, \ldots, E_k\}$ be the set of chosen edges.

The *transition matrix* of the map $f : \Gamma \to \Gamma$ is the matrix $M(f)$ of size $k \times k$ such that the ij^{th} entry of $M(f)$ is equal to the total number of occurrences of E_i and $\overline{E_i}$ in the path $f(E_j)$.

Ex.:

$$\begin{aligned} E_1 &\to E_1\overline{E}_2 \\ E_2 &\to E_2 \end{aligned}$$

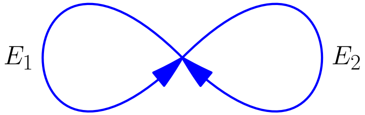

$$M(f) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

Filtration

$\emptyset = \Gamma_0 \subset \Gamma_1 \subset \cdots \subset \Gamma_N = \Gamma$, where $f(\Gamma_i) \subset \Gamma_i$

$H_i := cl(\Gamma_i \setminus \Gamma_{i-1})$ is called the i-th *stratum*.

$M(f) =$ 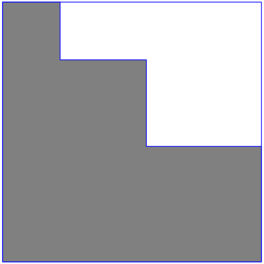

Filtration

$\emptyset = \Gamma_0 \subset \Gamma_1 \subset \cdots \subset \Gamma_N = \Gamma$, where $f(\Gamma_i) \subset \Gamma_i$

$H_i := cl(\Gamma_i \setminus \Gamma_{i-1})$ is called the i-th *stratum*.
If the filtration is maximal, then the matrices $M_1, \ldots, M_N$ are irreducible.

$M(f) =$

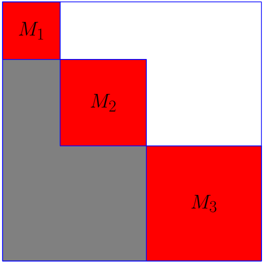

Strata

Frobenius: If $M \geqslant 0$ is a nonzero irreducible integer matrix, then
$\exists\, \vec{v} > 0$ and $\lambda \geqslant 1$ such that $M\vec{v} = \lambda\vec{v}$.
If $\lambda = 1$, then M is a permutation matrix.
v is unique up to a positive factor.
$\lambda =$ max of absolute values of eigenvalues of M.

Strata

Frobenius: If $M \geqslant 0$ is a nonzero irreducible integer matrix, then $\exists\, \vec{v} > 0$ and $\lambda \geqslant 1$ such that $M\vec{v} = \lambda\vec{v}$.
If $\lambda = 1$, then M is a permutation matrix.
v is unique up to a positive factor.
$\lambda =$ max of absolute values of eigenvalues of M.

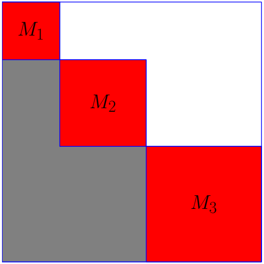

A stratum $H_i := cl(\Gamma_i \setminus \Gamma_{i-1})$ is called

exponential if $M_i \neq 0$ and $\lambda_i > 1$

polynomial if $M_i \neq 0$ and $\lambda_i = 1$

zero if $M_i = 0$

A metric for an exponential stratum

Let $H_r = cl(\Gamma_r \setminus \Gamma_{r-1})$ be an exponential stratum and let $E_{\ell+1}, \ldots, E_{\ell+s}$ be the edges of H_r.

A metric for an exponential stratum

Let $H_r = cl(\Gamma_r \setminus \Gamma_{r-1})$ be an exponential stratum and let $E_{\ell+1}, \ldots, E_{\ell+s}$ be the edges of H_r.

We have $vM_r = \lambda_r v$ for some $v = (v_1, \ldots, v_s) > 0$ and $\lambda_r > 1$.

A metric for an exponential stratum

Let $H_r = cl(\Gamma_r \setminus \Gamma_{r-1})$ be an exponential stratum and let $E_{\ell+1}, \ldots, E_{\ell+s}$ be the edges of H_r.

We have $vM_r = \lambda_r v$ for some $v = (v_1, \ldots, v_s) > 0$ and $\lambda_r > 1$.

We set $L_r(E_{\ell+i}) = v_i$ for edges $E_{\ell+i}$ in H_r
and $L_r(E) = 0$ for edges E in Γ_{r-1},
and extend L_r to paths in Γ_r.

A metric for an exponential stratum

Let $H_r = cl(\Gamma_r \setminus \Gamma_{r-1})$ be an exponential stratum and let $E_{\ell+1}, \dots, E_{\ell+s}$ be the edges of H_r.

We have $vM_r = \lambda_r v$ for some $v = (v_1, \dots, v_s) > 0$ and $\lambda_r > 1$.

We set $L_r(E_{\ell+i}) = v_i$ for edges $E_{\ell+i}$ in H_r
and $L_r(E) = 0$ for edges E in Γ_{r-1},
and extend L_r to paths in Γ_r.

Claim. For any path $p \subset \Gamma_r$ holds $L_r(f^k(p)) = \lambda_r^k(L_r(p))$.

Relative train track

Let $f : \Gamma \to \Gamma$ be a homotopy equivalence such that $f(\Gamma^0) \subseteq \Gamma^0$ and f maps edges to reduced paths.
The map f is called a *relative train track* if there exists a maximal filtration in Γ such that each exponential stratum H_r of this filtration satisfies the following conditions:

Relative train track

Let $f : \Gamma \to \Gamma$ be a homotopy equivalence such that $f(\Gamma^0) \subseteq \Gamma^0$ and f maps edges to reduced paths.
The map f is called a *relative train track* if there exists a maximal filtration in Γ such that each exponential stratum H_r of this filtration satisfies the following conditions:

(RTT-i) Df maps the set of oriented edges of H_r to itself; in particular all mixed turns in (G_r, G_{r-1}) are legal;

(RTT-ii) If $\rho \subset G_{r-1}$ is a nontrivial edge-path with endpoints in $H_r \cap G_{r-1}$, then $[f(\rho)]$ is a nontrivial path with endpoints in $H_r \cap G_{r-1}$;

(RTT-iii) For each legal edge-path $\rho \subset H_r$, the subpaths of $f(\rho)$ which lie in H_r are legal.

Relative train track

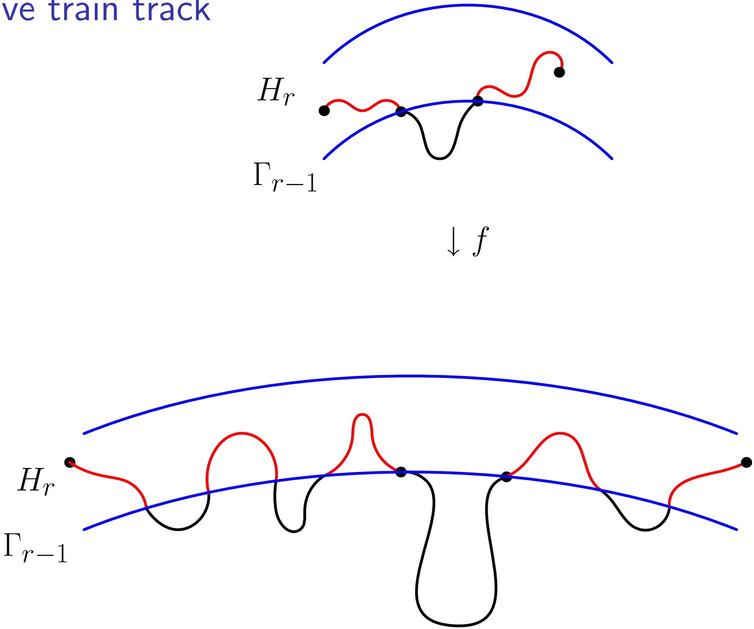

A useful fact

A path $p \subset \Gamma_r$ is called *r-legal*
if the pieces of p lying in H_r are legal.

Claim. For any r-legal reduced path $p \subset \Gamma_r$ holds

$$L_r([f^k(p)]) = \lambda_r^k(L_r(p)).$$

Theorem of Bestvina and Handel (1992)

Theorem [BH] Let F be a free group of finite rank. For every automorphism $\alpha : F \to F$, one can algorithmically construct a relative train track $f : \Gamma \to \Gamma$ which realizes the outer class of α.

Theorem of Bestvina and Handel (1992)

Theorem [BH] Let F be a free group of finite rank. For any automorphism α of F one can algorithmically

- construct a relative train track $f : \Gamma \to \Gamma$
- indicate a vertex $v \in \Gamma^0$ and path p in Γ from v to $f(v)$
- indicate an isomorphism $i : F \to \pi_1(\Gamma, v)$

such that the automorphism $i^{-1}\alpha i$ of the group $\pi_1(\Gamma, v)$ coincides with the map given by the rule

$$[x] \mapsto [p \cdot f(x) \cdot \bar{p}],$$

where $[x] \in \pi_1(\Gamma, v)$.

First improvement

Theorem [BH] Let F be a free group of finite rank. For any automorphism α of F one can algorithmically

- construct a relative train track $f : \Gamma \to \Gamma$
- indicate a vertex $v \in \Gamma^0$ and path p in Γ from v to $f(v)$
- indicate an isomorphism $i : F \to \pi_1(\Gamma, v)$
- compute a natural number n,

such that the automorphism $i^{-1}\alpha^n i$ of the group $\pi_1(\Gamma, v)$ coincides with the map given by the rule

$$[x] \mapsto [p \cdot f(x) \cdot \bar{p}],$$

where $[x] \in \pi_1(\Gamma, v)$.

(Pol) Every polynomial stratum H_r consists of only two mutually inverse edges, say E and $\overline{E}$. Moreover, $f(E) \equiv E \cdot a$, where a is a path in Γ_{r-1}.

Second improvement

Theorem Let F be a free group of finite rank. For any automorphism α of F one can algorithmically

- construct a relative train track $f_1 : \Gamma_1 \to \Gamma_1$
- indicate a vertex $v_1 \in \Gamma_1^0$ fixed by f_1
- indicate an isomorphism $i : F \to \pi_1(\Gamma_1, v_1)$
- compute a natural number n,

such that

$$i^{-1}\alpha^n i = (f_1)_*$$

and

(Pol) Every polynomial stratum H_r consists of only two mutually inverse edges, say E and $\overline{E}$. Moreover, $f_1(E) \equiv E \cdot a$, where a is a path in Γ_{r-1}.

Setting

Claim. Let α be an automorphism of a free group F of finite rank. If we know a basis of $\mathrm{Fix}(\alpha^n)$, we can compute a basis of $\mathrm{Fix}(\alpha)$.

Proof. $H = \mathrm{Fix}(\alpha)$ is a subgroup of $G = \mathrm{Fix}(\alpha^n)$.
The restriction $\alpha|_G$ is an automorphism of finite order of G.
Let

$$\overline{G} = G \rtimes \langle \alpha|_G \rangle.$$

Kalajdzevski: one can compute a finite generator set of $C_{\overline{G}}(\alpha|_G)$.
Reidemeister-Schreier: one can compute a finite generator set of $H = C_{\overline{G}}(\alpha|_G) \cap G$.

Setting

Passing from α to appropriate α^n, we can

- construct a relative train track $f : (\Gamma, v) \to (\Gamma, v)$
- indicate an isomorphism $i : F \to \pi_1(\Gamma, v)$

such that

$$i^{-1}\alpha i = f_*$$

and

(Pol) Every polynomial stratum H_r consists of only two mutually inverse edges, say E and $\overline{E}$. Moreover, $f(E) \equiv E \cdot a$, where a is a path in G_{r-1}.

Setting

Passing from α to appropriate α^n, we can

- construct a relative train track $f : (\Gamma, v) \to (\Gamma, v)$
- indicate an isomorphism $i : F \to \pi_1(\Gamma, v)$

such that

$$i^{-1} \alpha i = f_*$$

and

(Pol) Every polynomial stratum H_r consists of only two mutually inverse edges, say E and $\overline{E}$. Moreover, $f(E) \equiv E \cdot a$, where a is a path in G_{r-1}.

Claim. To construct a basis of $\mathrm{Fix}(\alpha)$, it suffices to construct a basis of

$$\overline{\mathrm{Fix}}(f) = \{[p] \in \pi_1(\Gamma, v) \,|\, f(p) = p\}.$$

Graph D_f for the relative train track $f : \Gamma \to \Gamma$

1. Definition of f-paths in Γ
2. Definition of D_f
3. Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f) \;\; \cong \mathrm{Fix}(\alpha)$
4. Preferable directions in D_f
5. Repelling edges, dead vertices in D_f
6. A procedure to construct a core of D_f
7. How to convert this procedure into an algorithm

1. f-paths in Γ

An edge-path μ in Γ is called an *f-path* if $\omega(\mu) = \alpha(f(\mu))$:

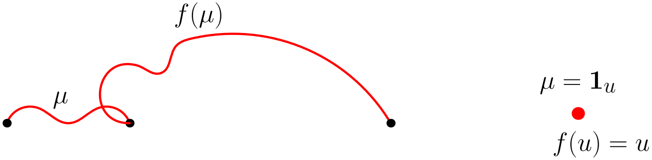

1. f-paths in Γ

An edge-path μ in Γ is called an *f-path* if $\omega(\mu) = \alpha(f(\mu))$:

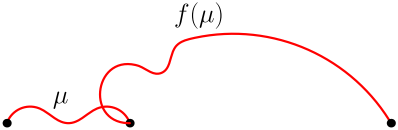

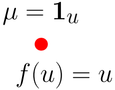

If μ is an f-path and E is an edge in Γ such that $\alpha(E) = \alpha(\mu)$, then $\overline{E}\mu f(E)$ is also an f-path:

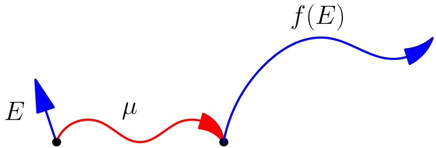

Definition of D_f

Vertices of D_f are reduced f-paths in Γ.

Two vertices μ and τ in D_f are connected by an edge with label E if E is an edge in Γ satisfying $\alpha(E) = \alpha(\mu)$ and $\tau = [\overline{E}\mu f(E)]$.

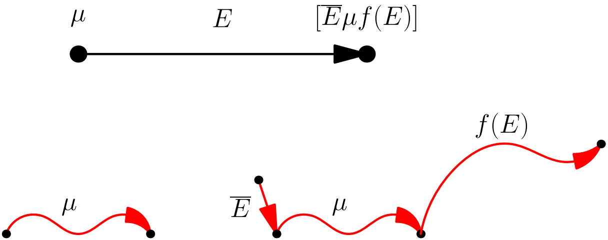

Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f)$

D_f :

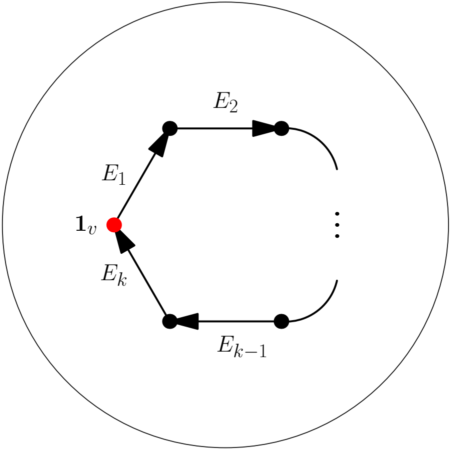

$\mathbf{1}_\mathrm{v}$

Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f)$

D_f :

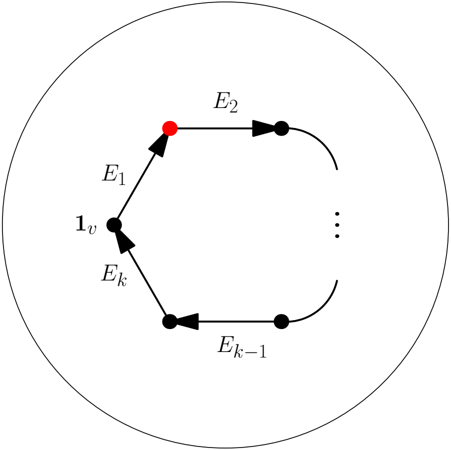

$[\overline{E}_1 \mathbf{1}_{\mathrm{v}} f(E_1)]$

Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f)$

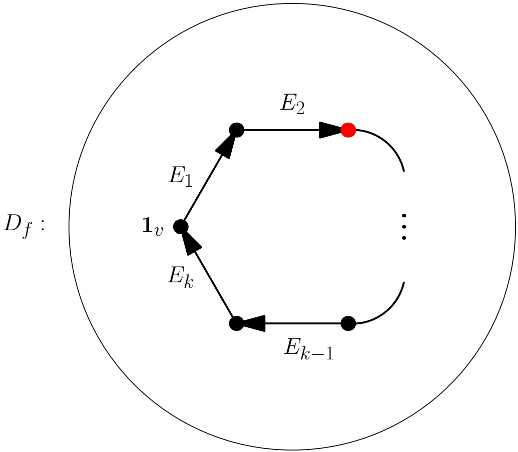

$$[\overline{E}_2[\overline{E}_1\mathbf{1}_\mathrm{v}f(E_1)]f(E_2)]$$

Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f)$

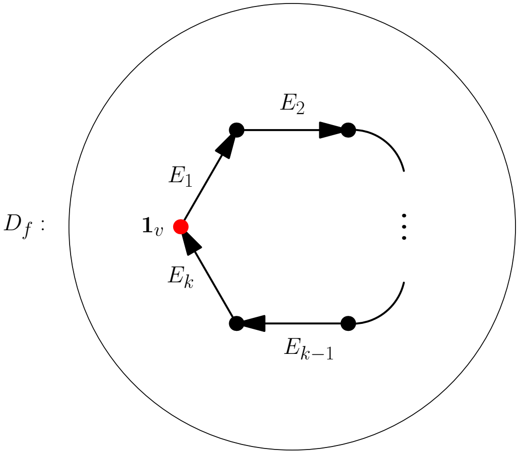

$$[\overline{E}_k \dots [\overline{E}_2[\overline{E}_1 \mathbf{1}_\mathrm{v} f(E_1)] f(E_2)] \dots f(E_k)] = \mathbf{1}_\mathrm{v}$$

Proof that $\pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \overline{\mathrm{Fix}}(f)$

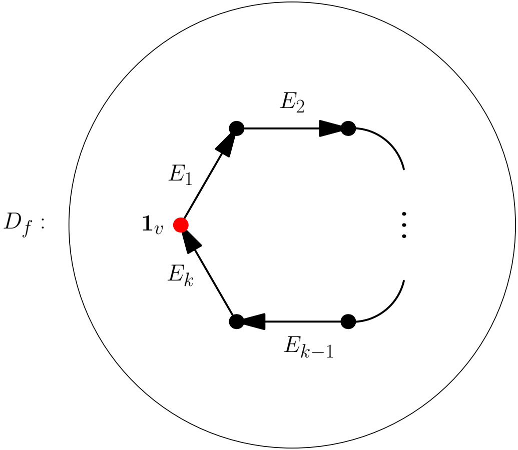

$$[\overline{E}_k \ldots [\overline{E}_2[\overline{E}_1 \mathbf{1}_\mathrm{v} f(E_1)]f(E_2)] \ldots f(E_k)] = \mathbf{1}_\mathrm{v}$$

$$[E_1 E_2 \ldots E_k] = [f(E_1 E_2 \ldots f(E_k))] \in \overline{\mathrm{Fix}}(f)$$

Preferable directions in D_f

Let μ be an f-path in Γ.
Suppose $E_1, \dots, E_k$ are all edges outgoing from $\alpha(\mu)$.
Then the vertex μ is connected with the vertices $[\overline{E}_i \mu f(E_i)]$ of D_f.
We set $\widehat{f}(\mu) := [\overline{E} \mu f(E)]$ if E is the first edge of the f-path μ.

in Γ:

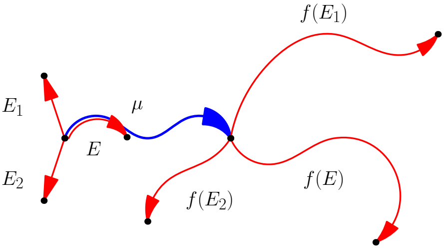

in D_f :

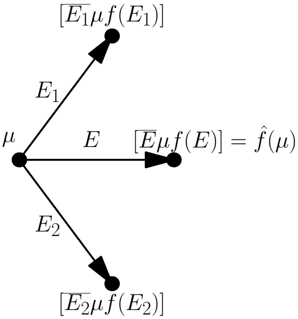

Preferable directions in D_f

Let μ be an f-path in Γ.
Suppose $E_1, \ldots, E_k$ are all edges outgoing from $\alpha(\mu)$.
Then the vertex μ is connected with the vertices $[\overline{E}_i \mu f(E_i)]$ of D_f.
We set $\widehat{f}(\mu) := [\overline{E} \mu f(E)]$ if E is the first edge of the f-path μ.

in Γ:

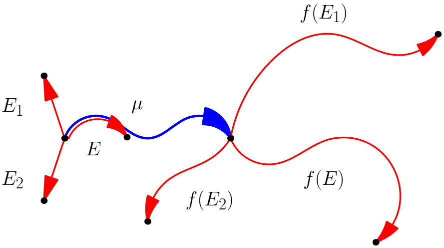

in D_f :

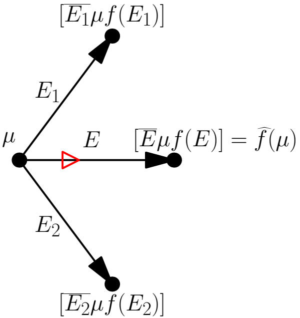

The preferable direction at the vertex $\mu \in D_f$ is the direction of the edge from μ to $\widehat{f}(\mu)$ with label E.

Graph D_f: example

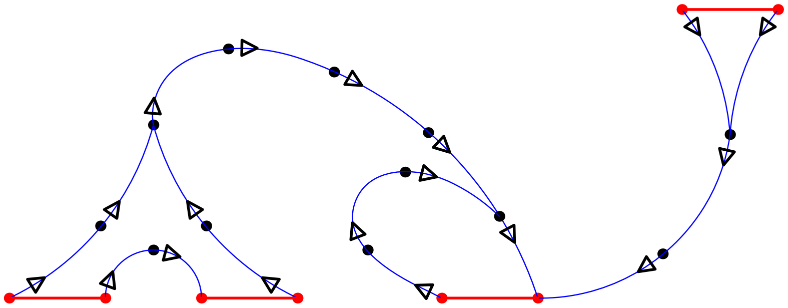

Graph D_f: example

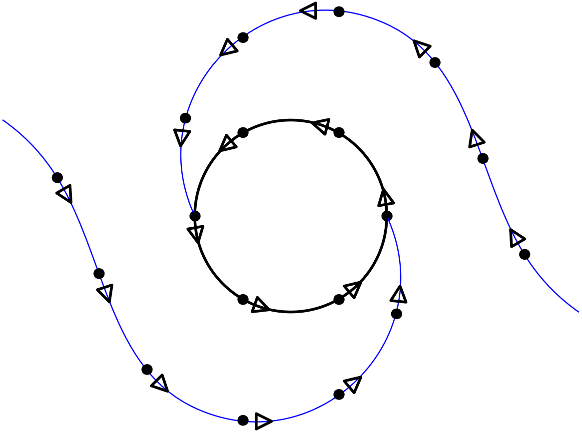

Graph D_f: example

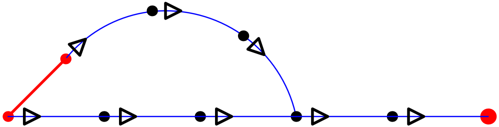

Graph D_f: example

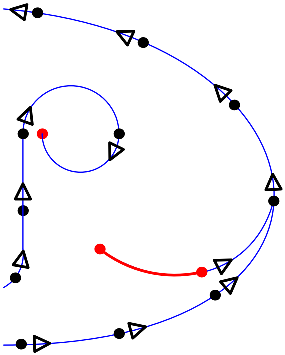

Definition of repelling edges in D_f

repelling edges — not repelling edges

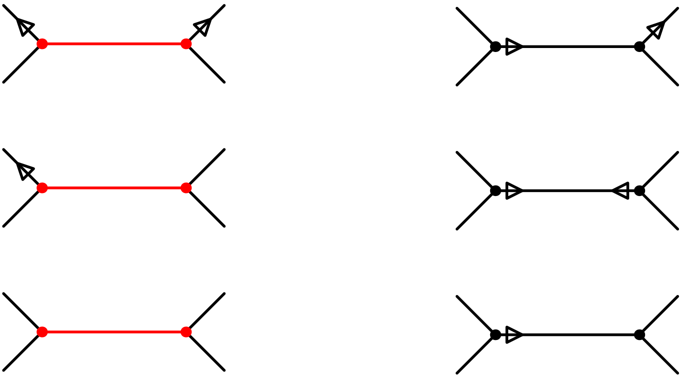

Let e be an edge of D_f with $\alpha(e) = u$, $\omega(e) = v$, and $Lab(e) = E$. The edge e is called *repelling* in D_f if E is not the first edge of the f-path u in Γ and $\overline{E}$ is not the first edge of the f-path v in Γ.

How to find repelling edges

Proposition (Cohen, Lustig). The repelling edges of D_f are in 1-1 correspondence with the occurrences of edges E in $f(E)$, where $E \in \Gamma^1$. More precisely, there exists a bijection of the type:

$f(E) \equiv \overline{u} \cdot E \cdot v \iff$

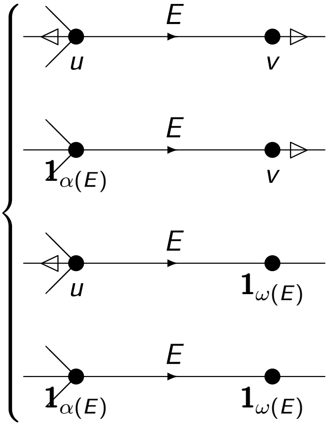

- if u and v are nonempty,
- if u is empty and v not,
- if v is empty and u not,
- if u and v are empty.

There is only finitely many repelling edges and they can be algorithmically found.

μ-subgraphs in D_f

Recall that if $\mu = E_1 E_2 \dots E_m$ is a vertex in D_f with $m \geqslant 1$, then

$$\widehat{f}(\mu) = [E_2 \dots E_m f(E_1)].$$

We define $\mu_1 := \mu$ and $\mu_{i+1} := \widehat{f}(\mu_i)$ if μ_i is nondegenerate.
The μ-subgraph consists of the vertices $\mu_1, \mu_2, \dots$ and the edges which connect μ_i with μ_{i+1} and carry the preferable direction at μ_i.

μ-subgraphs in D_f

Recall that if $\mu = E_1 E_2 \ldots E_m$ is a vertex in D_f with $m \geqslant 1$, then

$$\widehat{f}(\mu) = [E_2 \ldots E_m f(E_1)].$$

We define $\mu_1 := \mu$ and $\mu_{i+1} := \widehat{f}(\mu_i)$ if μ_i is nondegenerate.
The μ-subgraph consists of the vertices $\mu_1, \mu_2, \ldots$ and the edges which connect μ_i with μ_{i+1} and carry the preferable direction at μ_i.

Types of μ-subgraphs:

a segment ending at a dead vertex

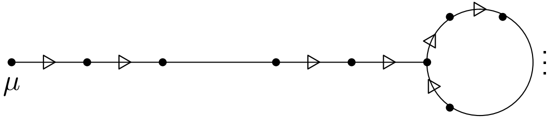

a segment with a cycle

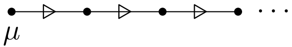

a ray

An important claim

Claim. If $\mathbf{1}_v$ lies in a non-contractible component C of D_f, then C contains a repelling vertex μ such that $\mathbf{1}_v$ belongs to the μ-subgraph.

Inverse preferred direction

Let f be a homotopy equivalence $\Gamma \to \Gamma$ s.t. f maps vertices to vertices and edges to reduced edge-paths.

We have algorithmically defined preferred directions at almost all vertices of D_f. There exists finitely many repelling edges in D_f and they can be algorithmically found.

Inverse preferred direction

Let f be a homotopy equivalence $\Gamma \to \Gamma$ s.t. f maps vertices to vertices and edges to reduced edge-paths.

We have algorithmically defined preferred directions at almost all vertices of D_f. There exists finitely many repelling edges in D_f and they can be algorithmically found.

Turner: One can algorithmically define the so called *inverse preferred direction* at almost all vertices of D_f. It has the following properties.
1) There exists finitely many inv-repelling edges in D_f and they can be algorithmically found.

Inverse preferred direction

2) Suppose that R is a μ-ray in D_f. Then the preferred direction on all but finitely many edges in R is opposite to the inverse preferred direction.

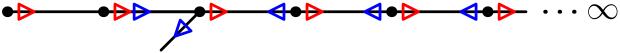

In particular R contains a normal vertex, i.e. a vertex where the red and the blue directions exist and different.

Inverse preferred direction

3) Let R_1 be a μ_1-ray and R_2 be a μ_2-ray, both don't contain inv-repelling edges and suppose that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one is contained in the other.

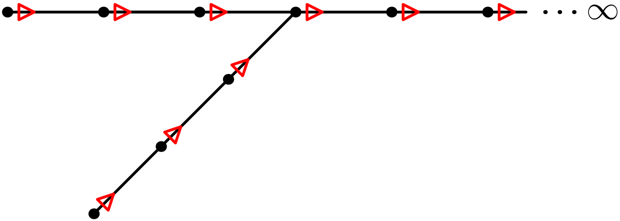

Inverse preferred direction

3) Let R_1 be a μ_1-ray and R_2 be a μ_2-ray, both don't contain inv-repelling edges and suppose that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one is contained in the other.

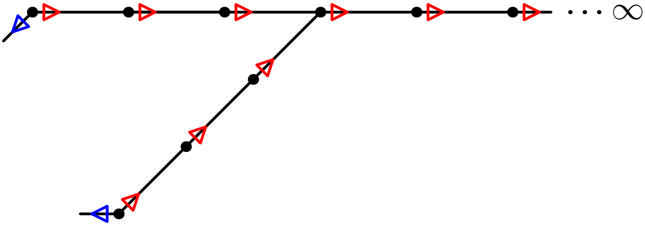

Inverse preferred direction

3) Let R_1 be a μ_1-ray and R_2 be a μ_2-ray, both don't contain inv-repelling edges and suppose that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one is contained in the other.

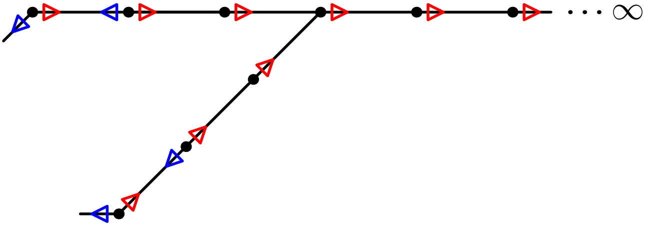

Inverse preferred direction

3) Let R_1 be a μ_1-ray and R_2 be a μ_2-ray, both don't contain inv-repelling edges and suppose that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one is contained in the other.

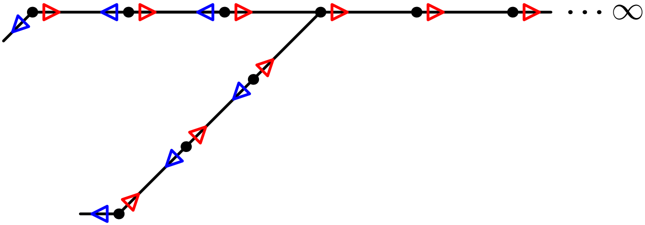

Inverse preferred direction

3) Let R_1 be a μ_1-ray and R_2 be a μ_2-ray, both don't contain inv-repelling edges and suppose that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one is contained in the other.

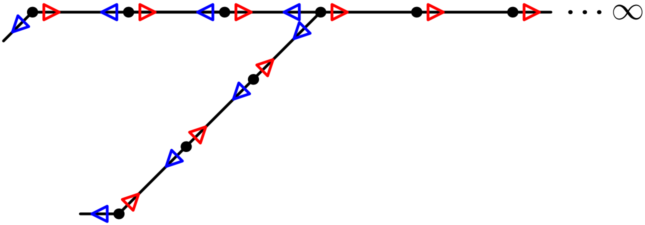

A procedure for construction of $CoRe(D_f)$

(1) Compute repelling edges.

(2) For each repelling vertex μ determine, whether the μ-subgraph is finite or not.

(3) Compute all elements of all finite μ-subgraphs from (2).

(4) For each two repelling vertices μ and τ with infinite μ-and τ-subgraphs determine, whether these subgraphs intersect.

(5) If the μ-subgraph and the τ-subgraph from (4) intersect, find their first intersection point and compute their initial segments up to this point.

How to convert this procedure into an algorithm?

It suffices to solve the following problems:

Problem 1. Given a vertex μ of the graph D_f, determine whether the μ-subgraph is finite or not.

Problem 2. Given two vertices μ and τ of the graph D_f, verify whether τ is contained in the μ-subgraph.

How to convert this procedure into an algorithm?

It suffices to solve the following problems:

Problem 1. Given a vertex μ of the graph D_f, determine whether the μ-subgraph is finite or not.

Problem 2. Given two vertices μ and τ of the graph D_f, verify whether τ is contained in the μ-subgraph.

We solve these problems in:

http://de.arxiv.org/abs/1204.6728

r-cancelation points in paths

A path $\mu \subset \Gamma$ has height r if $\mu \subset \Gamma_r$ and μ has at least one edge in H_r.

r-cancelation points in paths

A path $\mu \subset \Gamma$ has height r if $\mu \subset \Gamma_r$ and μ has at least one edge in H_r.

Let $\mu \subset \Gamma$ be a path of height r, where H_r is exponential.

A vertex v in μ is called an r-cancelation point in μ
if the turn (A, B) at v is an illegal r-turn:

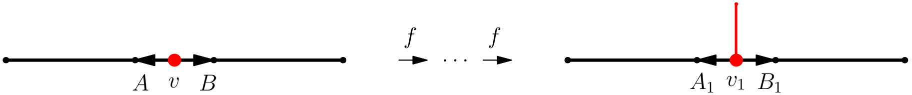

Non-deletable r-cancelation points

Let $\mu \subset \Gamma$ be a path of height r, where H_r is an exponential stratum.

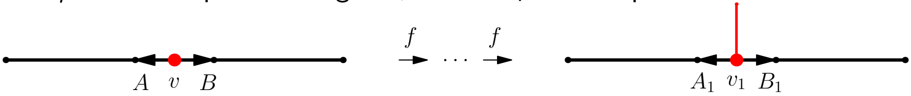

Suppose

- v divides μ into two r-legal subpaths
- v is an r-cancelation point in μ

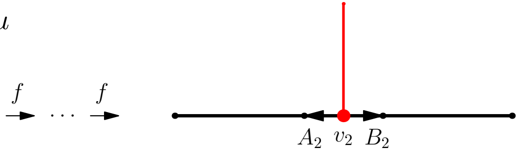

Then

- v is called a nondeletable r-cancelation point in μ if $\exists\ \infty$ illegal r-turns (A_k, B_k).

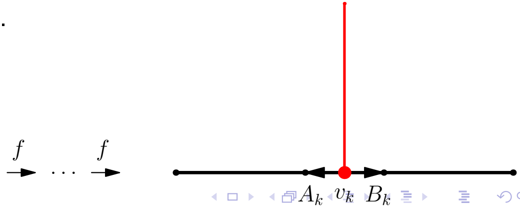

Nondeletability of r-cancelation points in paths is verifiable

Theorem. Let $f : \Gamma \to \Gamma$ be a relative train track. Let μ be a path in Γ of height r, where H_r is exponential. Suppose that a vertex v divides μ into two r-legal paths and v is an r-cancelation point.

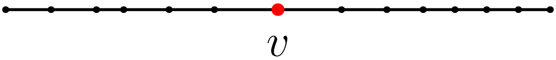

Nondeletability of r-cancelation points in paths is verifiable

Theorem. Let $f : \Gamma \to \Gamma$ be a relative train track. Let μ be a path in Γ of height r, where H_r is exponential. Suppose that a vertex v divides μ into two r-legal paths and v is an r-cancelation point.

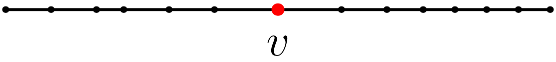

Then:
1) One can (effectively and uniformly) decide, whether v is deletable in μ or not.
2) If v is non-deletable in μ, one can compute the so called cancelation area $A(v, \mu)$ and the cancelation radius $a(v, \mu)$.

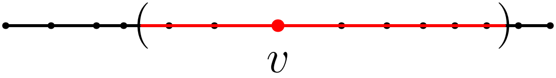

$$a(v, \mu) = L_r(A_{left}(v, \mu)) = L_r(A_{right}(v, \mu)).$$

r-cancelation areas in iterates of μ

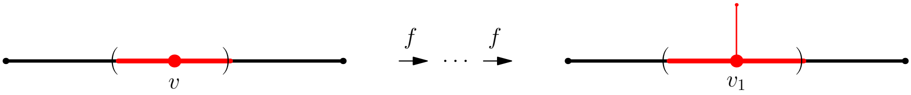

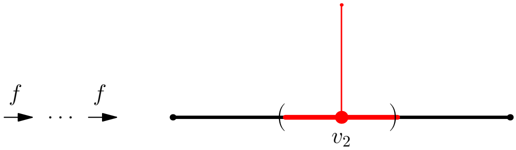

Let

- H_r be exp
- $\mathrm{Height}(\mu) = r$
- μ is not r-legal
- v divides μ into two r-legal subpaths
- v is a nondeletable r-cancelation point in μ

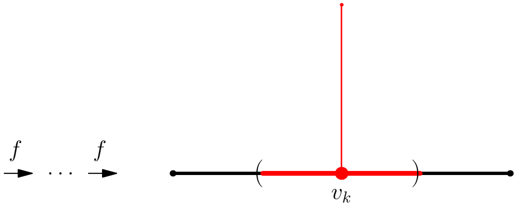

Different r-cancelation areas can interact

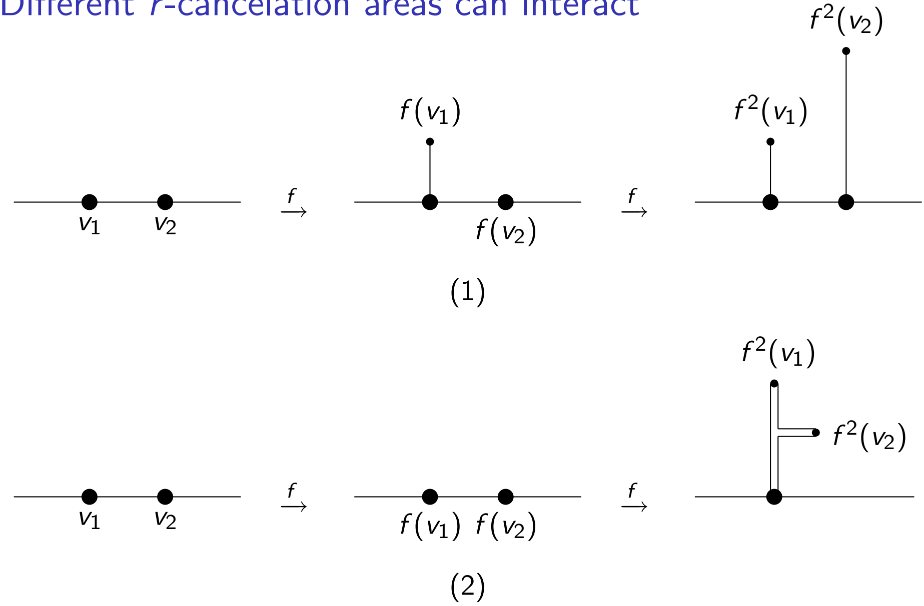

r-stability of paths

Def. Let $\mu \subset \Gamma_r$ be a path of height r, where H_r is exponential. μ is called *r-stable* if the number of r-cancelation points in

$$\mu, [f(\mu)], [f^2(\mu)], \ldots$$

is the same. Hence these points are non-deletable.

Several r-cancelation points in one path

Let μ be a path in Γ of height r, where H_r is exponential. Suppose:

- vertices $v_1, \ldots, v_n$ divide μ into r-legal paths $\mu_0, \ldots, \mu_n$.
- v_i is a nondeletable r-cancelation point in $\mu_{i-1}\mu_i$ for all i.

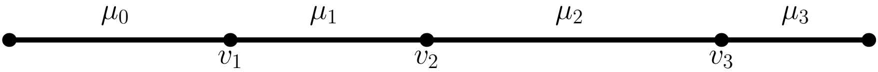

Let $a(v_i)$ be the cancelation radius of v_i in $\mu_{i-1}\mu_i$.
Theorem. μ is *stable* iff $a(v_i) + a(v_{i+1}) \geqslant L_r(\mu_i)$ for all i.

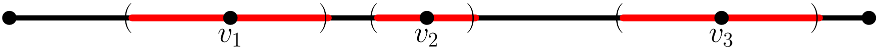

Stability theorem

Theorem. One can check, whether μ is r-stable.
If μ is not r-stable, one can compute n such that $[f^n(\mu)]$ is r-stable.

Finiteness and computability of the r-cancelation areas

Theorem.

1) There exists only finitely many r-cancelation areas in the infinite set of paths of height r. All r-cancelation areas $A_1, \ldots, A_k$ can be computed.

2) After appropriate subdivision of $f : \Gamma \to \Gamma$ the following holds: One can compute a natural $P = P(f)$ such that for every exponential stratum H_r and every r-cancelation area A, the r-cancelation area $[f^P(A)]$ is an edge-path.

μ-subgraphs in details (no cancelations)

Let $\mu = E_1 E_2 \dots E_n$ be an f-path.
Below is an ideal situation (no cancelations):

$$
\begin{aligned}
\mu &\equiv E_1 E_2 \dots E_n & , \\
\widehat{f}(\mu) &\equiv E_2 E_3 \dots E_n \cdot f(E_1) & , \\
\widehat{f}^{\,2}(\mu) &\equiv E_3 E_4 \dots E_n \cdot f(E_1) \cdot f(E_2) & , \\
&\vdots \\
\widehat{f}^{\,n}(\mu) &\equiv f(E_1) \cdot f(E_2) \cdot \ldots \cdot f(E_n), \\
&\vdots
\end{aligned}
$$

μ-subgraphs in details (no cancelations)

Let $\mu = E_1 E_2 \dots E_n$ be an f-path.
Below is an ideal situation (no cancelations):

$$
\begin{array}{lll}
\mu & \equiv E_1 E_2 \dots E_n & , \\
\widehat{f}(\mu) & \equiv E_2 E_3 \dots E_n \cdot f(E_1) & , \\
\widehat{f}^{\,2}(\mu) & \equiv E_3 E_4 \dots E_n \cdot f(E_1) \cdot f(E_2) & , \\
\vdots & & \\
\widehat{f}^{\,n}(\mu) & \equiv f(E_1) \cdot f(E_2) \cdot \dots \cdot f(E_n), & \\
\vdots & &
\end{array}
$$

Then Problems 1 and 2 can be reduced to:
Problem 1'. Do there exist $p > q$ such that $f^p(\mu) \equiv f^q(\mu)$?
Problem 2'. Does there exist p such that $f^p(\mu) \equiv \tau$?

μ-subgraphs in details (no cancelations)

Let $\mu = E_1 E_2 \dots E_n$ be an f-path.
Below is an ideal situation (no cancelations):

$$\begin{array}{lll}
\mu & \equiv E_1 E_2 \dots E_n & , \\
\widehat{f}(\mu) & \equiv E_2 E_3 \dots E_n \cdot f(E_1) & , \\
\widehat{f}^{\,2}(\mu) & \equiv E_3 E_4 \dots E_n \cdot f(E_1) \cdot f(E_2) & , \\
\vdots & & \\
\widehat{f}^{\,n}(\mu) & \equiv f(E_1) \cdot f(E_2) \cdot \dots \cdot f(E_n), & \\
\vdots & &
\end{array}$$

Then Problems 1 and 2 can be reduced to:
Problem 1'. Do there exist $p > q$ such that $f^p(\mu) \equiv f^q(\mu)$?
Problem 2'. Does there exist p such that $f^p(\mu) \equiv \tau$?

Solution. In this special case we have $\ell(\widehat{f}^{\,i+1}(\mu)) \geqslant \ell(\widehat{f}^{\,i}(\mu))$.

μ-subgraphs in details (there are cancelations)

We define 3 types of *perfect* f-paths:

- r-perfect
- A-perfect
- E-perfect

Definition of an r-perfect path

Let H_r be an exponential stratum. An edge-path $\mu \subset \Gamma_r$ is called r-perfect if the following conditions are satisfied:

- μ is a reduced f-path and its first edge belongs to H_r,
- μ is r-legal,
- $[\mu f(\mu)] \equiv \mu \cdot [f(\mu)]$ and the turn of this path at the point between μ and $[f(\mu)]$ is legal.

Definition of an A-perfect path

Let H_r be an exponential stratum. A reduced f-path $\mu \subset \Gamma_r$ containing edges from H_r is called A-perfect if

- all r-cancelation points in μ are non-deletable, the corresponding r-cancelation areas are edge-paths,
- the A-decomposition of μ starts on an A-area, i.e. it has the form $\mu \equiv A_1 b_1 \dots A_k b_k$,
- $[\mu f(\mu)] \equiv \mu \cdot [f(\mu)]$ and the turn at the point between μ and $[f(\mu)]$ is legal.

Definition of an E-perfect path

We may assume that $f : \Gamma \to \Gamma$ satisfies the condition (Pol): Each polynomial stratum H_r has a the unique (up to inversion) edge E and $f(E) \equiv E \cdot \sigma$, where σ is a path in Γ_{r-1}.

Let μ be an f-path of height r, where H_r is a polynomial stratum. μ is called E-perfect if

- the first edge of μ is E or $\overline{E}$,
- every path $\widehat{f}^{\,i}(\mu)$, $i \geqslant 1$ contains the same number of E-edges as μ.

μ-subgraphs in details (there are cancelations)

We define 3 types of *perfect* f-paths:

- r-perfect
- A-perfect
- E-perfect

Property. If σ is an r-perfect or A-perfect f-path, then there is no cancelation in passing from σ to $\widehat{f}(\sigma)$:

$$\begin{aligned} \sigma &\equiv E_1 E_2 \ldots E_n, \\ \widehat{f}(\sigma) &\equiv E_2 E_3 \ldots E_n \cdot f(E_1), \end{aligned}$$

$\widehat{f}(\sigma)$ may be not perfect, but ...

μ-subgraphs in details (there are cancelations)

We define 3 types of *perfect* f-paths:

- r-perfect
- A-perfect
- E-perfect

Property. If σ is an r-perfect or A-perfect f-path, then there is no cancelation in passing from σ to $\widehat{f}(\sigma)$:

$$\begin{aligned} \sigma &\equiv E_1 E_2 \ldots E_n, \\ \widehat{f}(\sigma) &\equiv E_2 E_3 \ldots E_n \cdot f(E_1), \end{aligned}$$

$\widehat{f}(\sigma)$ may be not perfect, but ...

Theorem.

1) If a μ-subgraph is infinite, it contains ∞ many perfect vertices:

$$\widehat{f}^{\,n_1}(\mu), \widehat{f}^{\,n_2}(\mu), \widehat{f}^{\,n_3}(\mu) \ldots .$$

2) Perfectness is verifiable.

μ-subgraphs in details (there are cancelations)

Weak alternative. Moving along the μ-subgraph, we can detect one of:

- the μ-subgraph is finite,
- the μ-subgraph contains a perfect vertex v_0.

In the second case we still have to decide, whether the μ-subgraph is finite or not.

μ-subgraphs in details (there are cancelations)

Weak alternative. Moving along the μ-subgraph, we can detect one of:

- the μ-subgraph is finite,
- the μ-subgraph contains a perfect vertex v_0.

In the second case we still have to decide, whether the μ-subgraph is finite or not.

Case 1. If v_0 is r-perfect, then

(1) $L_r(\widehat{f}^{\,i+1}(v_0)) \geqslant L_r(\widehat{f}^{\,i}(v_0)) > 0$ for all $i \geqslant 0$.

(2) There exist computable natural numbers $m_1 < m_2 < \dots$, such that
$L_r(\widehat{f}^{\,m_i}(v_0)) = \lambda_r^i L_r(v_0)$ for all $i \geqslant 1$.

$\Rightarrow$ In this case the μ-subgraph is ∞ and the membership problem in it is solvable.

μ-subgraphs in details (there are cancelations)

Case 2. If v_0 is A-perfect, then we can find a finite set $\{v_0, v_1, \ldots, v_k\}$ of A-perfect vertices in the v_0-subgraph such that all A-perfect vertices in the v_0-subgraph are:

$$\begin{array}{llll} v_0, & v_1, & \ldots, & v_k, \\ [f(v_0)], & [f(v_1)], & \ldots, & [f(v_k)], \\ [f^2(v_0)], & [f^2(v_1)], & \ldots, & [f^2(v_k)], \\ \ldots & & & \end{array}$$

Moreover, given a vertex u in the v_0-subgraph, we can find a number ℓ, such that $\widehat{f}^{\,\ell}(u)$ is an A-vertex.

μ-subgraphs in details (there are cancelations)

Case 2. If v_0 is A-perfect, then we can find a finite set $\{v_0, v_1, \ldots, v_k\}$ of A-perfect vertices in the v_0-subgraph such that all A-perfect vertices in the v_0-subgraph are:

$$\begin{array}{llll} v_0, & v_1, & \ldots, & v_k, \\ [f(v_0)], & [f(v_1)], & \ldots, & [f(v_k)], \\ [f^2(v_0)], & [f^2(v_1)], & \ldots, & [f^2(v_k)], \\ \ldots & & & \end{array}$$

Moreover, given a vertex u in the v_0-subgraph, we can find a number ℓ, such that $\widehat{f}^{\,\ell}(u)$ is an A-vertex.

So the finiteness and the membership problems for the v_0-subgraph can be reduced to:

Problem FIN. Does there exist $m > n \geqslant 0$ such that

$$[f^n(v_0)] = [f^m(v_0)]?$$

Problem MEM. Given an f-path τ, does there exist $n \geqslant 0$ s.t.

$$[f^n(v_0)] = \tau?$$

Both can be answered with the help of a theorem of Brinkmann.

THANK YOU!