

**AN ALGORITHM FOR FINDING A BASIS
OF THE FIXED POINT SUBGROUP OF AN AUTOMORPHISM
OF A FREE GROUP**

OLEG BOGOPOLSKI

*Mathematisches Institut, Universität Düsseldorf,
Universitätsstraße 1, 40225 Düsseldorf, Germany,
and Sobolev Institute of Mathematics, Novosibirsk,
4 Acad. Koptuyug avenue, 630090, Novosibirsk, Russia
Oleg.Bogopolski@yahoo.com*

OLGA MASLAKOVA

*Sobolev Institute of Mathematics, Novosibirsk,
4 Acad. Koptuyug avenue, 630090, Novosibirsk, Russia
o.s.maslakova@gmail.ru*

We describe an algorithm which, given an automorphism φ of a free group F of finite rank, computes a basis of the fixed point subgroup $\text{Fix}(\varphi)$.

1. Introduction

Let F_n be the free group of finite rank n . For any automorphism φ of F_n the *fixed point subgroup* of φ is

$$\text{Fix}(\varphi) = \{x \in F_n \mid \varphi(x) = x\}.$$

In the seminal paper [2], Bestvina and Handel proved the Scott conjecture that

$$\text{rk Fix}(\varphi) \leq n.$$

However, the following problem has been open for almost 20 years.

Problem A. Find an algorithm for computing a basis of $\text{Fix}(\varphi)$, where φ is an automorphism of a free group F of finite rank.

A weaker form of this problem is formulated in [12, Problem (F1) (a)].

Problem A has been solved in three special cases: for positive automorphisms in the paper [5] of Cohen and Lustig, for special irreducible automorphisms in the paper of Turner [11, Proposition B], and for all automorphisms of F_2 in the paper of Bogopolski [3].

In 1999, Maslakova attempted to solve this problem in general case. However, her proof published in [8], see also [9], was not complete. So, we have decided to give a full and correct proof. The main result of this paper is the following.

Theorem 1.1. *Let F_n be the free group of a finite rank n . There exists an algorithm which, given an automorphism φ of F_n finds a basis of its fixed point subgroup $\text{Fix}(\varphi) = \{x \in F_n \mid \varphi(x) = x\}$.*

As in [2], we use the relative train track techniques. A relative train track is a homotopy equivalence $f : \Gamma \rightarrow \Gamma$ of a finite graph Γ with certain good properties, see Sec. 3. In [2, Theorem 5.12], Bestvina and Handel proved that every outer automorphism \mathcal{O} of F_n can be represented by a relative train track $f : \Gamma \rightarrow \Gamma$. However, to start our algorithm, we need to represent the automorphism φ (and not its outer class) by a relative train track $f : (\Gamma, v) \rightarrow (\Gamma, v)$. This is done in Theorem 4.4.

We use f to define an auxiliary graph D_f (first introduced in [7] in another setting, see also [11]). The fundamental group of one of the components of D_f , denoted $D_f(\mathbf{1}_v)$, can be identified with $\text{Fix}(\varphi)$ (see Sec. 5). Thus, to compute a basis of $\text{Fix}(\varphi)$, we need to construct the core $\text{Core}(D_f(\mathbf{1}_v))$ of this component.

At all but finite number of the vertices of D_f there is a preferable outgoing direction. This determines a flow on almost all of D_f . The inverse automorphism φ^{-1} determines its own flow on almost all of D_f . According to [11] (see also [5]), there is a procedure for constructing a part of $\text{Core}(D_f)$ which contains $\text{Core}(D_f(\mathbf{1}_v))$: one should start from a finite number of computable exceptional edges and follow the first flow for sufficiently long. Theoretically we could arrive at a dead vertex, or get a loop, or arrive at a vertex where two rays of this flow meet, or none of these may occur. To convert this procedure into an algorithm, we must detect at the beginning which possibility occurs. For that, we must solve the Finiteness and the Membership problems for vertices and certain subgraphs of D_f (see Sec. 5). We solve these problems in this paper. Our algorithm for finding a basis of $\text{Fix}(\varphi)$ is given in Sec. 15.

2. Preliminaries

Let Γ be a finite connected graph, Γ^0 be the set of its vertices, and Γ^1 be the set of its edges. The initial vertex of an edge E is denoted by $\alpha(E)$, the terminal by $\omega(E)$, the inverse edge to E is denoted by \overline{E} .

The geometric realization of Γ is obtained by identification of each edge of Γ with a real segment $[a, b]$ of length 1. This realization is denoted again by Γ . Using this realization, we can work with partial edges and compute distances between points

inside edges without passing to a subdivision. *Partial edges* in Γ are identified with subsegments $[a_1, b_1] \subset [a, b]$. Let l be the corresponding metric on Γ .

We work only with piecewise linear maps. For brevity, we skip the wording *piecewise linear*, for example, we say a path instead of a piecewise linear path.

A *nontrivial path* in Γ is a continuous map $\tau : [0, 1] \rightarrow \Gamma$ with the following property: there exist numbers $0 = s_1 < s_2 < \dots < s_k < s_{k+1} = 1$ and a sequence of (possibly partial) edges E_1, E_2, \dots, E_k , such that $\tau|_{[s_i, s_{i+1}]}$ is a linear map onto E_i for each $i = 1, \dots, k$. We will not usually distinguish between τ and the concatenation of (partial) edges $E_1 E_2 \dots E_k$. The *length* of τ is $l(\tau) := \sum_{i=1}^k l(E_i)$. A *trivial path* in Γ is a map $\tau : [0, 1] \rightarrow \Gamma$ whose image consists of a single point. The trivial path whose image is $\{u\}$ is denoted by $\mathbf{1}_u$; we set $l(\mathbf{1}_u) = 0$. An *edge path* in Γ is either a path of the form $E_1 E_2 \dots E_k$, where all E_i are full edges, or a trivial path $\mathbf{1}_u$, where u is a vertex. The initial and the terminal points of a path τ are denoted by $\alpha(\tau)$ and $\omega(\tau)$, respectively. The inverse path to τ is denoted by $\bar{\tau}$. By $[\tau]$ we denote the reduced path in Γ which is homotopic to τ relative to the endpoints of τ . Let $[[\tau]]$ be the class of paths homotopic to τ relative to the endpoints of τ . For two paths τ, μ , we write $\tau = \mu$ if these paths are homotopic relative to endpoints and we write $\tau \equiv \mu$ if they graphically coincide. The concatenation of τ and μ (if it exists) is denoted by $\tau\mu$ or $\tau \cdot \mu$.

Let \mathcal{PLHE} be the class of all homotopy equivalences $f : \Gamma \rightarrow \Gamma$ such that Γ is a finite connected graph, $f(\Gamma^0) \subseteq \Gamma^0$, and for each edge E the following is satisfied: $f(E) \equiv E_1 E_2 \dots E_k$, where each E_i is an edge and E has a subdivision into segments, $E \equiv e_1 e_2 \dots e_k$, such that $f|_{e_i} : e_i \rightarrow E_i$ is surjective and linear with respect to the metric l . The abbreviation \mathcal{PLHE} stands for *piecewise linear homotopy equivalence*. If $f : \Gamma \rightarrow \Gamma$ belongs to \mathcal{PLHE} , then, for every path τ in Γ , the map $f \circ \tau$ is also a path in Γ . We denote it by $f(\tau)$.

The homotopy equivalence $f : \Gamma \rightarrow \Gamma$ is called *tight* (respectively, *nondegenerate*) if for each edge E in Γ the path $f(E)$ is reduced (respectively, nontrivial). The *norm* of f is the number $\|f\| := \max\{l(f(E)) \mid E \text{ is an edge of } \Gamma\}$. We use the following bounded cancellation lemma from [6], where it is credited to Grayson and Thurston.

Lemma 2.1. *Let Γ be a finite connected graph and $f : \Gamma \rightarrow \Gamma$ be a homotopy equivalence sending edges to edge paths. Let τ_1, τ_2 be reduced paths in Γ such that $\omega(\tau_1) = \alpha(\tau_2)$ and the path $\tau_1 \tau_2$ is reduced. Then*

$$l([f(\tau_1 \tau_2)]) \geq l([f(\tau_1)]) + l([f(\tau_2)]) - 2C_*,$$

where $C_* > 0$ is an algorithmically computable constant which depends only on f .

3. Relative Train Tracks for Outer Automorphisms of Free Groups

First we recall the definition of a relative train track from [2]. All auxiliary notions given before Definition 3.1 first appeared in [2, Secs. 1 and 5]. Since we are interested in algorithmic problems, we will work only with homotopy equivalences from the class \mathcal{PLHE} .

Let Γ be a finite connected graph and let $f : \Gamma \rightarrow \Gamma$ be a tight and nondegenerate homotopy equivalence from the class \mathcal{PLHE} .

A *turn* in Γ is an unordered pair of edges of Γ originating at a common vertex. A turn is *nondegenerate* if these edges are distinct, and it is *degenerate* otherwise. The map $f : \Gamma \rightarrow \Gamma$ induces a map $Df : \Gamma^1 \rightarrow \Gamma^1$ which sends each edge $E \in \Gamma^1$ to the first edge of the path $f(E)$. This induces a map Tf on turns in Γ by the rule $Tf(E_1, E_2) = (Df(E_1), Df(E_2))$. A turn (E_1, E_2) is *legal* if the turns $(Tf)^n(E_1, E_2)$ are nondegenerate for all $n \geq 0$; a turn is *illegal* if it is not legal. An edge path $E_1 E_2 \cdots E_m$ in Γ is *legal* if all its turns $(\overline{E_i}, E_{i+1})$ are legal. Clearly, a legal edge path is reduced.

From each pair of mutually inverse edges of the graph Γ we choose one edge. Let $\{E_1, \dots, E_k\}$ be the ordered set of chosen edges. The *transition matrix* of the map f (with respect to this ordering) is the matrix $M(f)$ of the size $k \times k$ such that the ij th entry of $M(f)$ is equal to the total number of occurrences of the edges E_i and $\overline{E_i}$ in the path $f(E_j)$.

A *filtration* for $f : \Gamma \rightarrow \Gamma$ is a strictly increasing sequence of (not necessarily connected) f -invariant subgraphs $\emptyset = G_0 \subset \cdots \subset G_N = \Gamma$. The subgraph $H_i = \text{cl}(G_i \setminus G_{i-1})$ is called the i th *stratum*. Edges in H_i are called *i -edges*. A turn with both edges in H_i is called an *i -turn*. A turn with one edge in H_i and another in G_{i-1} is called *mixed* in (G_i, G_{i-1}) . We assume that the edges of Γ are ordered so that the edges from H_i precede the edges from H_{i+1} . The edges from H_i define a square submatrix $M_{[i]}$ of $M(f)$.

The filtration is called *maximal* if each matrix $M_{[i]}$ is irreducible or zero. We will assume that all filtrations are maximal. If $M_{[i]}$ is nonzero, then it has the associated Perron–Frobenius eigenvalue $\lambda_i \geq 1$. If $\lambda_i > 1$, then the stratum H_i is called *exponential*. If $\lambda_i = 1$, then H_i is called *polynomial*. In this case $M_{[i]}$ is a permutation matrix, hence for every edge $E \in H_i^1$ the path $f(E)$ contains exactly one edge of H_i . A stratum H_i is called a *zero stratum* if $M_{[i]}$ is a zero matrix. In this case $f(E)$ lies in G_{i-1} for every edge $E \in H_i^1$.

Definition 3.1 (see [2, p. 38]). Let Γ be a finite connected graph and let $f : \Gamma \rightarrow \Gamma$ be a tight and nondegenerate homotopy equivalence from the class \mathcal{PLHE} . The map f is called a *PL-relative train track* if there exists a maximal filtration $\emptyset = G_0 \subset \cdots \subset G_N = \Gamma$ for f such that each exponential stratum H_r of this filtration satisfies the following conditions:

- (RTT-i) Df maps the set of 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 reduced nontrivial edge path with endpoints in $H_r \cap G_{r-1}$, then $[f(\rho)]$ is a nontrivial edge path with endpoints in $H_r \cap G_{r-1}$.
- (RTT-iii) For each legal edge path $\rho \subset H_r$, the path $f(\rho)$ does not contain any illegal turns in H_r .

Definition 3.2. We use the above notations.

- (1) A path ρ in Γ is said to be of *height* r if it lies in G_r , but not in G_{r-1} .
- (2) Let H_r be an exponential stratum. A nontrivial reduced path ρ in G_r is called *r-legal* if the minimal edge path containing ρ does not contain any illegal turns in H_r (see [2, p. 39]).

The following proposition will be used often in the subsequent proofs.

Proposition 3.3 (see [2, Lemma 5.8]). *Suppose that $f : \Gamma \rightarrow \Gamma$ is a PL-relative train track and H_r is an exponential stratum of Γ . Let ρ be a reduced r -legal path:*

$$\rho \equiv b_0 \cdot a_1 \cdot b_1 \cdot \dots \cdot a_k \cdot b_k,$$

where $k \geq 1$, a_1, \dots, a_k are paths in H_r , and b_0, \dots, b_k are paths in G_{r-1} , and all these paths except maybe b_0 and b_k are nontrivial. Then

$$[f(\rho)] \equiv [f(b_0)] \cdot f(a_1) \cdot [f(b_1)] \cdot \dots \cdot f(a_k) \cdot [f(b_k)]$$

and this path is r -legal. Moreover, for all $i \geq 1$ we have

$$[f^i(\rho)] \equiv [f^i(b_0)] \cdot [f^i(a_1)] \cdot [f^i(b_1)] \cdot \dots \cdot [f^i(a_k)] \cdot [f^i(b_k)]$$

and these paths are r -legal.

The r -length functions ℓ_r and L_r . Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track and H_r be an exponential stratum. Choose a positive vector \vec{v} satisfying $\vec{v}M_{[r]} = \lambda_r \vec{v}$. Since $M_{[r]}$ is an integer matrix, we can choose \vec{v} so that the coordinates of \vec{v} are rational functions of λ_r over \mathbb{Q} . First we define the *r-length* ℓ_r on edges of G_r . If E is the i th edge of H_r , we set $\ell_r(E) = v_i$; if E is an edge of G_{r-1} , we set $\ell_r(E) = 0$. For every edge path $\tau = E_1 E_2 \dots E_k$ in G_r , we define $\ell_r(\tau) = \sum_{i=1}^k \ell_r(E_i)$. Then we have $\ell_r(f(\tau)) = \lambda_r \ell_r(\tau)$.

We extend this definition to all paths (not necessarily edge paths) in G_r , as was done in [2, Lemma 5.10]. For an arbitrary path μ in G_r , let $\ell_r^\bullet(\mu)$ be the sum of r -lengths of full r -edges which occur in μ if they exist and zero if not. For any path ρ in G_r , we set

$$L_r(\rho) := \lim_{k \rightarrow \infty} \lambda_r^{-k} \ell_r^\bullet(f^k(\rho)).$$

Observe that $L_r(\rho) = \ell_r(\rho)$ if ρ is an edge path in G_r .

Lemma 3.4 (see [2, Lemma 5.10]). *Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track and H_r be an exponential stratum. The function L_r has the following properties:*

- (1) $L_r(f(\rho)) = \lambda_r L_r(\rho)$ for any path ρ in G_r .
- (2) $L_r([f(\rho)]) = \lambda_r L_r(\rho)$ for any reduced r -legal path ρ in G_r .
- (3) If ρ is a nontrivial initial or terminal segment of an r -edge, then $L_r(\rho) > 0$.

Proof. Property (1) follows from the definition of L_r ; (2) follows from Proposition 3.3. Property (3) is stated in [2, Lemma 5.10]. \square

A representation of an outer automorphism of F_n by a PL-relative train track. The rose with n petals R_n is the graph with one vertex $*$ and n geometric edges. We assume that the free group on n letters F_n is identified with $\pi_1(R_n, *)$. Obviously, every automorphism φ of F_n can be represented by a homotopy equivalence $R_n \rightarrow R_n$.

In [2, Theorem 5.12]), Bestvina and Handel proved that every outer automorphism \mathcal{O} of F_n can be represented by a relative train track $f : \Gamma \rightarrow \Gamma$. One can show that this proof can be organized in a constructive way. Also, we may assume that f is a PL-relative train track. Thus, we have the following starting point for our algorithm.

Theorem 3.5 (see [2, Theorem 5.12]). *Let F_n be the free group of finite rank n . There is an algorithm which, given an outer automorphism \mathcal{O} of F_n , constructs a PL-relative train track $f : \Gamma \rightarrow \Gamma$ and a homotopy equivalence (a marking) $\tau : R_n \rightarrow \Gamma$ such that f represents \mathcal{O} with respect to τ .*

The latter means that for any homotopy equivalence $\sigma : \Gamma \rightarrow R_n$ which is a homotopy inverse to τ , the map $(\sigma \circ f \circ \tau)_* : \pi_1(R_n, *) \rightarrow \pi_1(R_n, *)$ represents \mathcal{O} .

4. Relative Train Tracks for Automorphisms of Free Groups

Let F be a free group of finite rank, φ be an automorphism of F , and \mathcal{O} be the outer automorphism class of φ . Theorem 3.5 gives a representation of \mathcal{O} by a PL-relative train track. However this is not sufficient for our aims. The purpose of this section is to show that φ itself can be represented by a PL-relative train track, see Theorem 4.4.

Notation 4.1. Let Γ be a finite connected graph and $f : \Gamma \rightarrow \Gamma$ be a homotopy equivalence. For each vertex $v \in \Gamma^0$ we define the isomorphism

$$f_v : \pi_1(\Gamma, v) \rightarrow \pi_1(\Gamma, f(v)), \\ [[\mu]] \mapsto [[f(\mu)]], \quad \text{where } [[\mu]] \in \pi_1(\Gamma, v).$$

For each path p in Γ from v to $f(v)$ we define the automorphism

$$f_{v,p} : \pi_1(\Gamma, v) \rightarrow \pi_1(\Gamma, v), \\ [[\mu]] \mapsto [[pf(\mu)\bar{p}]], \quad \text{where } [[\mu]] \in \pi_1(\Gamma, v).$$

Remark 4.2. By Theorem 3.5, given an automorphism φ of F , one can construct a finite connected graph Γ , a PL-relative train track $f : \Gamma \rightarrow \Gamma$, and an isomorphism $i : F \rightarrow \pi_1(\Gamma, v)$, where v is a vertex of Γ , such that the automorphism $i \circ \varphi \circ i^{-1} : \pi_1(\Gamma, v) \rightarrow \pi_1(\Gamma, v)$ coincides with $f_{v,p}$ for some path $p \subset \Gamma$ from v to $f(v)$.

We claim that p can be computed. Indeed, if q is an arbitrary path in Γ from v to $f(v)$, then $f_{v,q}$ differs from $f_{v,p}$ by an inner automorphism of $\pi_1(\Gamma, v)$. Comparing $f_{v,q}$ with $i \circ \varphi \circ i^{-1}$, we can compute this inner automorphism and hence p . This gives us the following form of Theorem 3.5.

Theorem 4.3. *Let F be a free group of finite rank. There is an algorithm which, given an automorphism φ of F , constructs a PL-relative train track $f : \Gamma \rightarrow \Gamma$ and indicates a vertex $v \in \Gamma^0$, a path $p \subset \Gamma$ from v to $f(v)$, and an isomorphism $i : F \rightarrow \pi_1(\Gamma, v)$ such that the automorphism $i \circ \varphi \circ i^{-1} : \pi_1(\Gamma, v) \rightarrow \pi_1(\Gamma, v)$ coincides with $f_{v,p}$.*

The following theorem says that in Theorem 4.3 we can provide $f(v) = v$ and choose p equal to the trivial path at v .

Theorem 4.4. *Let F be a free group of finite rank. There is an algorithm which, given an automorphism φ of F , constructs a PL-relative train track $f_1 : \Gamma_1 \rightarrow \Gamma_1$ with a vertex $v_1 \in \Gamma_1^0$ fixed by f_1 , and indicates an isomorphism $j : F \rightarrow \pi_1(\Gamma_1, v_1)$ such that $j \circ \varphi \circ j^{-1} = (f_1)_{v_1}$.*

Proof. Let $f : \Gamma \rightarrow \Gamma$, v , p , and $i : F \rightarrow \pi_1(\Gamma, v)$ be the PL-relative train track, the vertex, the path, and the isomorphism from Theorem 4.3, respectively. Then we have $i \circ \varphi = f_{v,p} \circ i$. Hence, for every $w \in F$, we have

$$i(\varphi(w)) = [[p]] [[f(i(w))]] [[\bar{p}]]. \quad (4.1)$$

Let Γ_1 be the graph obtained from Γ by adding a new vertex v_1 and a new edge E connecting v_1 and $f(v)$. We extend the homotopy equivalence $f : \Gamma \rightarrow \Gamma$ to a map $f_1 : \Gamma_1 \rightarrow \Gamma_1$ by the rule $f_1(v_1) = v_1$ and $f_1(E) := E[f(p)]$. Clearly, f_1 is a homotopy equivalence. We define a maximal filtration for f_1 by extending the maximal filtration for f with the help of the new top polynomial stratum consisting of the edges E and \bar{E} . Finally, we define the isomorphism $j : F \rightarrow \pi_1(\Gamma_1, v_1)$ by the rule

$$j(w) := [[E]] [[f(i(w))]] [[\bar{E}]], \quad w \in F. \quad (4.2)$$

To complete the proof, we verify that the automorphism $j \circ \varphi \circ j^{-1}$ of the group $\pi_1(\Gamma_1, v_1)$ coincides with the induced automorphism $(f_1)_{v_1} : \pi_1(\Gamma_1, v_1) \rightarrow \pi_1(\Gamma_1, v_1)$. It suffices to check that $(f_1)_{v_1}(j(w)) = j(\varphi(w))$ for every $w \in F$:

$$\begin{aligned} (f_1)_{v_1}(j(w)) &\stackrel{(4.2)}{=} (f_1)_{v_1}([[E]] [[f(i(w))]] [[\bar{E}]]) \\ &= [[f_1(E)]] [[f^2(i(w))]] [[f_1(\bar{E})]] \\ &= [[Ef(p)]] [[f^2(i(w))]] [[f(\bar{p})\bar{E}]] \\ &= [[E]] [[f(p f(i(w)) \bar{p})]] [[\bar{E}]] \\ &\stackrel{(4.1)}{=} [[E]] [[f(i(\varphi(w)))] [[\bar{E}]] \stackrel{(4.2)}{=} j(\varphi(w)). \quad \square \end{aligned}$$

Thus, for computing a basis of $\text{Fix}(\varphi)$, it suffices to compute a basis of the group

$$\overline{\text{Fix}}(f_1) = \{[[\mu]] \in \pi_1(\Gamma_1, v_1) \mid f_1(\mu) = \mu\},$$

where Γ_1 is the graph, v_1 is the vertex, and f_1 is the PL-relative train track from Theorem 4.4.

5. Graphs D_f and $\text{CoRe}(C_f)$ for a Homotopy Equivalence $f : \Gamma \rightarrow \Gamma$

Let Γ be a finite connected graph with a distinguished vertex v_* . Let $f : \Gamma \rightarrow \Gamma$ be a homotopy equivalence which maps vertices of Γ to vertices and edges to reduced edge paths, and suppose that f fixes v_* . We consider the group

$$\overline{\text{Fix}}(f) := \{[p] \in \pi_1(\Gamma, v_*) \mid f(p) = p\}.$$

In papers [7, 11], the authors suggest a procedure for computation of a basis of $\overline{\text{Fix}}(f)$ with the help of a graph D_f . This procedure is not an algorithm in the general case, since one cannot determine from the beginning whether it terminates or not. We give a description of this procedure. We also show that the procedure can be converted into an algorithm if the Membership and the Finiteness problems can be algorithmically solved.

First, we recall some constructions and facts from [7, 11, 5].

A. Definition of f -paths. An edge path μ in Γ is called an f -path if the last point of μ coincides with the first point of $f(\mu)$. Observe that

- the trivial path at a vertex u of Γ , denoted $\mathbf{1}_u$, is an f -path if and only if u is fixed by f ;
- if μ is an f -path, then $[\mu]$ is also an f -path; and
- 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.

B. Definition of the graph D_f . The vertices of D_f are reduced f -paths in Γ . Let μ be a reduced f -path in Γ and let E_1, \dots, E_n be all edges in Γ outgoing from $\alpha(\mu)$. Then we connect the vertex μ of D_f to the vertices $[\overline{E_1}\mu f(E_1)], \dots, [\overline{E_n}\mu f(E_n)]$ by edges with labels E_1, \dots, E_n , respectively, see Fig. 1. The label of a nontrivial edge path in the graph D_f is the product of labels of consecutive edges of this path. The label of the trivial edge path at a vertex μ of D_f is $\mathbf{1}_{\alpha(\mu)}$.

For a vertex μ of D_f , let $D_f(\mu)$ be the component of D_f containing μ .

Lemma 5.1 (see [7]). *The fundamental group of each component of D_f is finitely generated. Moreover, $\pi_1(D_f(\mathbf{1}_{v_*}), \mathbf{1}_{v_*}) \cong \overline{\text{Fix}}(f)$.*

The proof in [7] uses preferable directions at vertices of D_f .

C. Preferable directions at vertices of D_f , dead and alive vertices of D_f . For a reduced nontrivial f -path μ in Γ , we set $\hat{f}(\mu) := [\overline{E}\mu f(E)]$, where E is the first edge of μ . Then μ and $\hat{f}(\mu)$ are vertices of the graph D_f connected by the edge with the label E . The direction of this edge is called *preferable* at the vertex μ . We will put the symbol \triangleright on this edge near the vertex μ .

Note that only the vertices $\mathbf{1}_w$, where $w \in \Gamma^0$ and $f(w) = w$, do not admit preferable directions. We call such vertices *dead* and all other vertices of D_f *alive*. Observe that at each vertex of D_f , there is at most one outwardly \triangleright -directed edge.




Fig. 1. On the left we consider μ as a path in Γ , and on the right as a vertex in D_f . The symbol \triangleright on the right shows the preferable direction at the vertex μ .

D. Ordinary, repelling and attracting edges of D_f .

Definition 5.2 (see Fig. 2). Let e be an edge of D_f , let p, q be the initial and the terminal vertices of e , and let $E \in \Gamma^1$ be the label of e .

- (1) The edge e is called *ordinary* in D_f if one of the following holds:
 - (a) E is the first edge of the path p in Γ and \bar{E} is not the first edge of the path q in Γ .
 - (b) E is not the first edge of the path p in Γ and \bar{E} is the first edge of the path q in Γ .
- (2) The edge e is called *repelling* in D_f if E is not the first edge of the path p in Γ and \bar{E} is not the first edge of the path q in Γ . In other words, e is repelling if and only if $f(E)$ starts with $\bar{p}E$.

A vertex of D_f is called *repelling* if it is the initial or the terminal vertex of a repelling edge.

- (3) The edge e is called *attracting* in D_f if E is the first edge of the path p in Γ and \bar{E} is the first edge of the path q in Γ .

An edge of D_f is called *exceptional* if it is attracting or repelling.




Fig. 2. Two ordinary edges, an attracting edge, and a repelling edge in D_f .

Proposition 5.3 (see [7, 11, 5]). *The repelling edges of D_f are in one-to-one correspondence with the occurrences of edges E in $f(E)$, where $E \in \Gamma^1$. In particular, there exist at most $|\Gamma^1| \cdot \|f\|$ repelling edges in D_f and they can be algorithmically found.*

E. Definition of a μ -subgraph of D_f . Let μ be a vertex in D_f . If μ is not a dead vertex, that is, if $\mu \equiv E_1 E_2 \cdots E_m$ for some edges $E_i \in \Gamma^1$, $m \geq 1$, then we can pass from μ to the vertex $\hat{f}(\mu) \equiv [E_2 \cdots E_m f(E_1)]$ by using the direction which is preferable at μ .

The vertices of the μ -subgraph are the vertices μ_1, μ_2, \dots of D_f such that $\mu_1 = \mu$ and $\mu_{i+1} = \hat{f}(\mu_i)$ if the vertex μ_i is not dead, $i \geq 1$. The edges of the μ -subgraph are those which connect μ_i with μ_{i+1} and carry the preferable direction at μ_i .

Note that the μ -subgraph is finite if and only if, starting from μ and moving along the preferable directions, we will come to a dead vertex or to a vertex which we have seen earlier. If the μ -subgraph is infinite, we call it the μ -ray. Thus, any μ -subgraph is one of the four types given in Fig. 3.




Fig. 3. Types of μ -subgraphs in D_f .

Let μ and τ be two vertices of D_f . Clearly, if the μ -subgraph and the τ -subgraph intersect, then they differ only by their finite “initial subsegments”. Observe that

- if μ_0 is a vertex of the μ -subgraph, then the μ_0 -subgraph is contained in the μ -subgraph;
- if μ_0 is a vertex of the μ -subgraph and τ_0 is a vertex of the τ -subgraph, then the μ -subgraph and the τ -subgraph intersect if and only if the μ_0 -subgraph and the τ_0 -subgraph intersect; and
- if the μ -subgraph is infinite and the τ -subgraph is finite, then they do not intersect.

From this point, we start to develop the above approach.

F. Definitions of the graphs C_f and $\text{CoRe}(C_f)$. A component of D_f is called *repelling* if it contains at least one repelling edge (see Fig. 4). Let C_1, \dots, C_n be all repelling components of D_f . For each C_i , let $\text{CoRe}(C_i)$ be the minimal connected subgraph of C_i which contains all repelling edges of C_i and carries $\pi_1(C_i)$. We set $C_f := \bigcup_{i=1}^n C_i$ and $\text{CoRe}(C_f) := \bigcup_{i=1}^n \text{CoRe}(C_i)$.




Fig. 4. An example of a graph D_f with three repelling components.

G. To construct the graph $\text{CoRe}(C_f)$, it suffices to do the following:

- (1) Find all repelling edges of D_f .
- (2) For each alive 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.

To convert this procedure to an algorithm, we shall construct algorithms for steps (2) and (4). In papers [5, 11] these algorithms are given only in some special cases (for positive automorphisms and for irreducible automorphisms represented by train tracks for which each fixed point is a vertex). The main idea in these papers is to use an *inverse preferred direction* at vertices in the graph D_f . At each vertex, this direction can be constructed algorithmically (in the general case) with the help of a homotopic inverse to f . The resulting inverse flow has its own repelling edges and repelling and dead vertices; they can be algorithmically found.

H. Inverse preferred directions in D_f . We will realize the following plan. First we define a map $g : \Gamma \rightarrow \Gamma$ which is a homotopy inverse to $f : \Gamma \rightarrow \Gamma$. Then we show that there is a label preserving graph map $\Phi : D_f \rightarrow D_g$. Finally we define the inverse preferred directions at vertices in D_f by pulling back the preferred directions in D_g by Φ . This idea is due to Turner [11], and has sources in the paper of Cohen and Lustig [5]. Note that in [11], the map Φ is claimed to be locally injective (see Proposition in Sec. 3 there), and we claim that Φ is an isomorphism.

Definition 5.4. For the given homotopy equivalence $f : \Gamma \rightarrow \Gamma$, we can algorithmically construct a homotopy equivalence $g : \Gamma \rightarrow \Gamma$ such that g maps vertices of Γ to vertices, edges to edge paths, and the maps $h := g \circ f$ and $f \circ g$ are homotopic to the identity on Γ . This can be done as follows.

First we choose a maximal tree T in Γ and orient the edges e_1, \dots, e_k of $\Gamma^1 \setminus T^1$. Then $\pi_1(\Gamma, v_*)$ can be identified with the free group F with basis e_1, \dots, e_k . Let f_* be the automorphism of F induced by the map f . We compute f_*^{-1} by using Nielsen transformations. Finally, we realize f_*^{-1} by a homotopy equivalence $g : \Gamma \rightarrow \Gamma$ as

follows: For every vertex u of Γ , let t_u be the unique reduced path in T from v_* to u . For $e \in \{e_1, \dots, e_k\}$, if f_*^{-1} sends e to a word $W(e_1, \dots, e_k)$, then we define $g(e)$ to be the tightening of the path $\overline{t_{\alpha(e)}} \cdot W(t_{\alpha(e_1)}e_1\overline{t_{\omega(e_1)}}, \dots, t_{\alpha(e_k)}e_k\overline{t_{\omega(e_k)}}) \cdot t_{\omega(e)}$. For $e \in T^1$, we define $g(e) = e$. This g is a homotopy inverse to f .

From now on, we fix g . Let $H : \Gamma \times [0, 1] \rightarrow \Gamma$ be a homotopy from the identity id to $h := g \circ f$. For each point u in Γ , let p_u be the path from u to $h(u)$ determined by the homotopy H : namely $p_u(t) = H(u, t)$, $t \in [0, 1]$. We set

$$K_*(f) := \max\{l(p_u) : u \in \Gamma^0\}.$$

First we define a map Φ from the set of vertices of D_f to the set of vertices of D_g . Let μ be a vertex in D_f . We regard μ as a reduced f -path in Γ and let u be the initial vertex of μ . Then we set $\Phi(\mu) = [p_u g(\overline{\mu})]$. Clearly, $\Phi(\mu)$ is a reduced g -path in Γ . Hence $\Phi(\mu)$ can be regarded as a vertex in D_g .

Lemma 5.5. *The map Φ can be continued to a graph homomorphism $\Phi : D_f \rightarrow D_g$ preserving the labels of edges.*

Proof. Let μ and μ_1 be two vertices in D_f connected by an edge with label E , that is, $\mu_1 = [\overline{E}\mu f(E)]$. We must show that $\Phi(\mu)$ and $\Phi(\mu_1)$ are connected by an edge with the label E , that is, $\Phi(\mu_1) = [\overline{E}\Phi(\mu)g(E)]$. Let u and w be the initial and the terminal vertices of E . Then u and w are the initial vertices of μ and μ_1 , respectively. We have

$$\begin{aligned} \Phi(\mu_1) &= [p_w g(f(\overline{E})\overline{\mu}E)] \\ &= [p_w h(\overline{E})g(\overline{\mu})g(E)] = [\overline{E}p_u g(\overline{\mu})g(E)] = [\overline{E}\Phi(\mu)g(E)]. \end{aligned}$$

Here we use the fact that H is a homotopy and hence

$$[h(\ell)] = [\overline{p_{\alpha(\ell)}}\ell p_{\omega(\ell)}] \tag{5.1}$$

for any path ℓ in Γ . □

Remark 5.6. Let μ be a vertex in D_f . Then the following holds:

- (1) The f -path μ and the g -path $\Phi(\mu)$ have the same initial vertices in Γ .
- (2) Let E_1, \dots, E_n be the edges outgoing from $\alpha(\mu)$ in Γ . Then the vertices μ and $\Phi(\mu)$ of the graphs D_f and D_g have degree n and the labels of edges outgoing from each of these vertices are E_1, \dots, E_n .

Proposition 5.7. *The map $\Phi : D_f \rightarrow D_g$ is an isomorphism of graphs.*

Proof. By Lemma 5.5 and Remark 5.6, it suffices to show that Φ is bijective on vertices. First we show that Φ is injective on vertices. Let μ_1, μ_2 be two different vertices of D_f . If the f -paths μ_1 and μ_2 have different initial vertices in Γ , then, by Remark 5.6(1), the g -paths $\Phi(\mu_1)$ and $\Phi(\mu_2)$ have different initial vertices in Γ too, hence $\Phi(\mu_1) \neq \Phi(\mu_2)$.

Suppose that the initial vertices of the f -paths μ_1 and μ_2 coincide and equal u . Then their terminal vertices also coincide and equal $f(u)$. Since the f -paths μ_1, μ_2 are reduced, $\mu_1 \neq \mu_2$, and g is a homotopy equivalence, we have $[g(\mu_1)] \neq [g(\mu_2)]$, hence $\Phi(\mu_1) = [p_u g(\bar{\mu}_1)] \neq [p_u g(\bar{\mu}_2)] = \Phi(\mu_2)$.

Now we show that Φ is surjective on vertices. Let τ be a vertex in D_g , that is, τ is a reduced g -path in Γ . Let u be the initial vertex of the path τ . We will find a reduced f -path μ in Γ such that $\Phi(\mu) = \tau$. Let μ_1 be an arbitrary path in Γ from u to $f(u)$. Then the paths τ and $p_u g(\bar{\mu}_1)$ have the same endpoints, so $\bar{\tau} p_u g(\bar{\mu}_1)$ is a loop based at $g(u)$. Hence, there exists a loop σ in Γ based at u such that $g(\sigma) = \bar{\tau} p_u g(\bar{\mu}_1)$. We set $\mu := [\sigma \mu_1]$. Then μ is an f -path and $\Phi(\mu) = [p_u g(\bar{\mu})] = [p_u g(\bar{\mu}_1) g(\bar{\sigma})] = \tau$. \square

Definition 5.8. The *inverse preferred direction* at a vertex μ in D_f is the preimage of the preferred direction at the vertex $\Phi(\mu)$ in D_g under Φ .

We formulate this in more detail. Recall that $\Phi(\mu) = [p_u g(\bar{\mu})]$, where u is the initial vertex of the f -path μ . First suppose that the g -path $\Phi(\mu)$ is nontrivial and let E be the first edge of this path. Then the inverse preferred direction at the vertex μ of D_f is the direction of the edge of D_f which starts at μ and has the label E . If the g -path $\Phi(\mu)$ is trivial in Γ , the inverse preferred direction at μ in D_f is not defined.

Proposition 5.9. *The inverse preferred direction is defined at all but finitely many vertices of D_f .*

Proof. If the inverse preferred direction at a vertex μ in D_f is not defined, then $\Phi(\mu)$ lies in the finite set $\{\mathbf{1}_u \mid u \in \Gamma^0\}$. Since Φ is injective, the number of such μ is finite. \square

Definition 5.10. Preimages, with respect to Φ , of repelling edges, repelling vertices, and dead vertices of D_g are called *inv-repelling* edges, *inv-repelling* vertices, and *inv-dead* vertices of D_f , respectively.

By Proposition 5.3 applied to g , there are only finitely many inv-repelling edges, inv-repelling vertices, and inv-dead vertices in D_f , and they can be algorithmically found.

I. Normal vertices

Definition 5.11. A vertex of D_f is called *normal* if the preferred and the inverse preferred directions at this vertex exist and do not coincide.

The main purpose of this subsection is to prove Propositions 5.15 and 5.16; they will help us to decide whether two rays in D_f (given by their initial vertices) meet.

The following lemma easily follows from Lemma 2.1.

Lemma 5.12. *Let Γ be a finite connected graph and $f : \Gamma \rightarrow \Gamma$ be a homotopy equivalence sending edges to edge paths. Let p be an initial subpath of a reduced*

path q . Write $[f(p)] \equiv ab$, where a is the maximal common initial subpath of $[f(p)]$ and $[f(q)]$. Then $l(b) \leq C_*(f)$.

The source of the following lemma is [5, Proposition (4.3)].

Lemma 5.13. *Let R be the μ -subgraph with consecutive vertices $\mu = \mu_0, \mu_1, \dots$, and with labels of consecutive edges E_1, E_2, \dots . For each $j \geq 0$ with alive vertex μ_j , let $k(j)$ be the maximal natural number such that $\mu_j \equiv E_{j+1} \cdots E_{j+k(j)} \cdot Z_j$ for some Z_j . If $j > l(\mu_0)$ and R has at least $j + k(j) + 2$ vertices, then $l(Z_j) \leq C_*(f)$.*

Proof. With notation $X_j := E_1 E_2 \cdots E_j$, we have $\mu_j \equiv [\overline{X_j} \mu_0 f(X_j)]$. Hence, $f(X_j) = \overline{\mu_0} X_j \mu_j$. Therefore $[f(X_j)] \equiv [\overline{\mu_0} X_j] \cdot E_{j+1} \cdots E_{j+k(j)} \cdot Z_j$. Indeed, the condition $j > l(\mu_0)$ guarantees that the last edge of $[\overline{\mu_0} X_j]$ is E_j , which is not inverse to E_{j+1} . Applying the same arguments to $\mu_{j+k(j)}$, we have

$$\begin{aligned} [f(X_{j+k(j)})] &\equiv [\overline{\mu_0} X_{j+k(j)}] \cdot E_{j+k(j)+1} \cdots E_{j+k(j)+k(j+k(j))} \cdot Z_{j+k(j)} \\ &\equiv [\overline{\mu_0} X_j] \cdot E_{j+1} \cdots E_{j+k(j)} \cdot E_{j+k(j)+1} \cdots E_{j+k(j)+k(j+k(j))} \cdot Z_{j+k(j)}. \end{aligned}$$

From Lemma 5.12 applied to X_j and $X_{j+k(j)}$, we deduce that $l(Z_j) \leq C_*(f)$. \square

The source of the following lemma is [11, Proposition (4.10)]. The map g and the constant $K_*(f)$ were defined in Definition 5.4.

Lemma 5.14. *Let R be the μ -subgraph with consecutive vertices $\mu = \mu_0, \mu_1, \dots$, and with labels of consecutive edges E_1, E_2, \dots . Let j be a natural number such that $j > l(\mu_0)$ and $l(\mu_j) > C_*(f) \cdot (\|g\| + 1) + K_*(f)$. If R has at least $j + k(j) + 2$ vertices, then $\mu_{j+k(j)}$ is normal. (Here $k(j)$ is as in Lemma 5.13.)*

Proof. It suffices to show that the first edge of the g -path $\Phi(\mu_{j+k(j)})$ is $\overline{E_{j+k(j)}}$. Then, by Definition 5.8, the inv-preferred direction at $\mu_{j+k(j)}$ in D_f will coincide with the direction of the edge outgoing from $\mu_{j+k(j)}$ and having the label $\overline{E_{j+k(j)}}$. On the other hand, the preferred direction at $\mu_{j+k(j)}$ in D_f coincides with the direction of the edge outgoing from $\mu_{j+k(j)}$ and having the label $E_{j+k(j)+1}$. Since these labels do not coincide, the vertex $\mu_{j+k(j)}$ is normal. By Lemma 5.13,

$$\mu_j \equiv E_{j+1} \cdots E_{j+k(j)} \cdot Z_j \quad \text{with } l(Z_j) \leq C_*(f). \quad (5.2)$$

This implies

$$\mu_{j+k(j)} = \overline{E_{j+1} \cdots E_{j+k(j)}} \mu_j f(E_{j+1} \cdots E_{j+k(j)}) = Z_j f(E_{j+1} \cdots E_{j+k(j)}).$$

Recall that $\Phi(\mu) = [p_{\alpha(\mu)} g(\overline{\mu})]$. Then, using (5.1), where $h = g \circ f$, we have

$$\begin{aligned} \Phi(\mu_{j+k(j)}) &\equiv [p_{\omega(E_{j+k(j)})} (g \circ f) (\overline{E_{j+1} \cdots E_{j+k(j)}}) g(\overline{Z_j})] \\ &\equiv [\overline{E_{j+1} \cdots E_{j+k(j)}} p_{\alpha(E_{j+1})} g(\overline{Z_j})]. \end{aligned}$$

From (5.2) and the assumption in this lemma, we have

$$\begin{aligned} l(\overline{E_{j+1} \cdots E_{j+k(j)}}) &= k(j) \geq l(\mu_j) - C_*(f) \\ &> K_*(f) + C_*(f) \cdot \|g\| \\ &\geq l(p_{\alpha(E_{j+1})}) + l(g(\overline{Z_j})). \end{aligned}$$

Therefore the first edge of $\Phi(\mu_{j+k(j)})$ is $\overline{E_{j+k(j)}}$. □

Proposition 5.15. *There exists an algorithm which, given an f -path μ , either proves that the μ -subgraph R is finite or finds a normal vertex in R .*

Proof. Computing consecutive vertices of R , $\mu = \mu_0, \mu_1, \dots$, we either prove that R is finite, or find the first j with $j > l(\mu_0)$ and $l(\mu_j) > C_*(f) \cdot (\|g\| + 1) + K_*(f)$. If we find such j , then we compute $k(j)$ (note that $k(j) \leq l(\mu_j)$) and check whether $\mu_0, \mu_1, \dots, \mu_{j+k(j)+1}$ exist and are different. If the result is negative, then R is finite; if positive, then the vertex $\mu_{j+k(j)}$ is normal by Lemma 5.14. □

The following proposition is contained in Claim (b) in the proof of [11, Theorem A]. This claim was inspired by [5, Lemma (4.8) and Proposition (4.10)]. The proof of this proposition is valid in the general situation, that is, for any homotopy equivalence $f : \Gamma \rightarrow \Gamma$ sending edges to edge paths. We give it for completeness.

Proposition 5.16. *Let R_1 and R_2 be the μ_1 -ray and the μ_2 -ray in D_f , respectively. Suppose that they do not contain inv-repelling vertices and that their initial vertices μ_1 and μ_2 are normal. Then R_1 and R_2 are either disjoint or one of them is contained in the other.*

Proof. Suppose that the rays R_1 and R_2 intersect and neither of them is contained in the other. We indicate the preferred directions by white triangles and the inv-preferred directions by black triangles. Since μ_1 and μ_2 are normal, the white and the black triangles at μ_1 and at μ_2 look in different directions, see Fig. 5(a).

Since R_1 and R_2 do not contain inv-repelling vertices, we can inductively reconstruct the inv-preferred directions at the vertices of R_1 and R_2 until the first intersection point of these rays, see Fig. 5(b). We obtain two inv-preferred directions at this point, a contradiction. □




Fig. 5. Two intersecting rays.

J. How to convert the procedure in G into an algorithm

As was observed, it suffices to find algorithms for steps (2) and (4). Using Propositions 5.15 and 5.16, Step (4) can be replaced by the following three steps.

- (4.1) For each repelling vertex μ whose μ -subgraph is a ray, find in this μ -ray a vertex μ' such that the μ' -ray does not contain inv-repelling vertices.
- (4.2) Find a normal vertex μ'' in the μ' -ray.
- (4.3) For every two repelling vertices μ and τ whose μ - and τ -subgraphs are rays, verify whether τ'' is contained in the μ'' -ray or μ'' is contained in the τ'' -ray.

Step (4.2) can be done algorithmically by Proposition 5.15. Steps (4.1) and (4.3) can be done if we find an algorithm for the following problem.

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

Indeed, for Step (4.1) we first find all inv-repelling vertices in D_f . Then we detect those of them which lie in the μ -ray. Let I be the minimal initial segment of the μ -ray which contains all these vertices. We can take μ' equal to the first vertex in the μ -ray which lies outside I . Step (4.3) is a partial case of the above problem.

Step (2) can be done if we find an algorithm for the following problem:

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

Thus, to construct $\text{CoRe}(C_f)$ algorithmically, it suffices to find algorithms for these problems.

Remark 5.17. (1) In Sec. 14, we will present algorithms solving the Membership problem and the Finiteness problem for D_f in the case where f is a PL-relative train track satisfying (RTT-iv) (see Definition 8.12). Thus, $\text{CoRe}(C_f)$ is constructible in this case.

(2) We give a plan for computing a basis of $\text{Fix}(\varphi)$:

- By Theorem 4.4, we represent φ by a PL-relative train track $f_1 : (\Gamma_1, v_1) \rightarrow (\Gamma_1, v_1)$. As can be seen from the proof of this theorem, we may assume that $\mathbf{1}_{v_1}$ is a repelling vertex in D_{f_1} . Subdividing as in Corollary 8.13, we may additionally assume that f_1 satisfies (RTT-iv).
- Since $\text{CoRe}(C_{f_1})$ is constructible, we can find its component Δ containing $\mathbf{1}_{v_1}$ and compute $\pi_1(\Delta, \mathbf{1}_{v_1})$. Since $\pi_1(\Delta, \mathbf{1}_{v_1}) \cong \pi_1(D_{f_1}(\mathbf{1}_{v_1}), \mathbf{1}_{v_1}) \cong \overline{\text{Fix}}(f_1)$ by Lemma 5.1, we can compute a basis of $\overline{\text{Fix}}(f_1)$ and hence a basis of $\text{Fix}(\varphi)$.

6. n -Transversal Paths

Definition 6.1. (1) If p is a reduced path in D_f from τ to σ with the label $E_1 E_2 \cdots E_k$, then there is a reduced path in D_f from $[f(\tau)]$ to $[f(\sigma)]$ with the label $[f(E_1 E_2 \cdots E_k)]$. We denote this path by $f_\bullet(p)$.

(2) Let R be the μ -subgraph. For each $i \geq 0$ with alive vertex $\hat{f}^i(\mu)$, let e_{i+1} be the edge of D_f from $\hat{f}^i(\mu)$ to $\hat{f}^{i+1}(\mu)$. Let $n \in \mathbb{N}$. A reduced path p in D_f is called n -transversal to the μ -subgraph if $\omega(p) = \hat{f}^n(\mu)$ and the last edge of p is different from e_n and \bar{e}_{n+1} . Let $\mathcal{S}_n(\mu)$ be the set of vertices $\sigma \in D_f$ such that there exists a reduced path p starting at σ and n -transversal to the μ -subgraph. The set $\mathcal{S}_n(\mu)$ is called the n -sphere determined by μ .

Theorem 6.2. *Let R be the μ -subgraph with consecutive vertices $\mu = \mu_0, \mu_1, \dots$. Let s be a natural number such that $s > (l(\mu_0) + 1) \cdot (\|f\| + 3)$ and $l(\mu_s) > 2C_*$, and R contains at least $s + k(s) + 2$ vertices, where $k(s)$ is defined in Lemma 5.13. If $\sigma \in \mathcal{S}_s(\mu)$, then $[f(\sigma)] \in \mathcal{S}_t(\mu)$ for some computable t satisfying $s < t$.*

Moreover, given an s -transversal path connecting σ to the μ -subgraph, one can construct a t -transversal path connecting $[f(\sigma)]$ to the μ -subgraph.

Proof (illustrated by Fig. 6). Let E_1, E_2, \dots be the labels of consecutive edges of R . We set $j := l(\mu_0) + 1$. Using notations of Lemma 5.13 we have

$$j + k(j) \leq j + l(\mu_j) = j + l([\bar{X}_j \mu_0 f(X_j)]) \leq 2j + l(\mu_0) + \|f\| \cdot j < s.$$

Let $[\mu_j, \mu_s]$ be the segment of R from μ_j to μ_s with the label $X_{j,s} := E_{j+1} \cdots E_s$.




Fig. 6. An illustration to the proof of Theorem 6.2.

Then $\mu_s \equiv [\bar{X}_{j,s} \mu_j f(X_{j,s})]$, hence $[f(X_{j,s})] \equiv [\bar{\mu}_j X_{j,s} \mu_s]$. By Lemma 5.13, we have $\mu_j \equiv E_{j+1} \cdots E_{j+k(j)} \cdot Z_j$ and $\mu_s \equiv E_{s+1} \cdots E_{s+k(s)} \cdot Z_s$, where $l(Z_j) \leq C_*$ and $l(Z_s) \leq C_*$. Using this and the estimate $j + k(j) < s$, one can easily check that the label of the path $f_\bullet([\mu_j, \mu_s])$ is $[f(X_{j,s})] \equiv \bar{Z}_j \cdot E_{j+k(j)+1} \cdots E_{s+k(s)} \cdot Z_s$.

Let $[\sigma, \mu_s]$ be a reduced path from σ to μ_s with the last edge different from the last edge of the path $[\mu_j, \mu_s]$. By Lemma 2.1, the maximal common terminal segment of $f_\bullet([\sigma, \mu_s])$ and $f_\bullet([\mu_j, \mu_s])$ has length at most C_* . Then $[f(\sigma)] \in \mathcal{S}_t(\mu)$ for some $t \geq s + k(s) - C_*$. By Lemma 5.13, we have $l(\mu_s) \leq k(s) + C_*$, and by assumption we have $l(\mu_s) > 2C_*$. Then $t > s$ as desired. \square

7. Dynamics of Paths in Γ Under Iteration of f

This section is an introduction to the subsequent Secs. 8 and 9. To explain the ideas behind them, we recall some important notions from [1]. Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track with the maximal filtration $\emptyset = G_0 \subset \cdots \subset G_N = \Gamma$.

A. Indivisible periodic Nielsen paths and r -cancellation areas

A nontrivial reduced path $\sigma \subset \Gamma$ is called a *periodic Nielsen path* for f if $[f^k(\sigma)] = \sigma$ for some $k \geq 1$; if $k = 1$, then σ is called a *Nielsen path*. The periodic Nielsen path σ is called *indivisible* if it cannot be written as a concatenation of nontrivial periodic Nielsen paths (see [1, Definition 5.1.1]). Indivisible periodic Nielsen paths will be abbreviated as IPN paths.

Suppose that H_r is an exponential stratum. Our nearest aim is to describe the dynamics of reduced paths of height r under iteration of f . Obviously, the sets of r -legal and IPN paths of height r must play an important role in the description. However, the set of IPN paths of height r is in some sense small, therefore the authors of [1] introduce a larger set P_r (see the paragraph before Lemma 4.2.5 in [1]).

Definition 7.1. The set P_r consists of all reduced paths ρ of height r in Γ that satisfy the following properties:

- (i) Each $[f^k(\rho)]$ contains exactly one illegal turn in H_r .
- (ii) The initial and terminal (possibly partial) edges of each $[f^k(\rho)]$ are contained in H_r .
- (iii) The number of r -edges in $[f^k(\rho)]$ is bounded independently of k .

In particular, each path from P_r can be decomposed into the product of two r -legal subpaths such that the turn between them is an illegal r -turn. This turn produces cancellations under iteration of f (see Fig. 7). By [1, Lemma 4.2.5], the set P_r is f -invariant and finite.

In Sec. 8, we define in a constructive way *r -cancellation areas* and show that these areas are exactly the elements of P_r . The main result of Sec. 8 is Theorem 8.10 which says that the set of r -cancellation areas is computable.

B. Splittings

Suppose that σ is a nontrivial reduced path in Γ . A decomposition $\sigma \equiv \sigma_1\sigma_2 \cdots \sigma_n$ is called a *splitting* of σ if $[f^k(\sigma)] \equiv [f^k(\sigma_1)][f^k(\sigma_2)] \cdots [f^k(\sigma_n)]$ for all $k > 0$ (see [1, Sec. 4.1]).

To shorten formulations, we introduce the following terminology. Let H_r be an exponential stratum in Γ . A reduced path $\tau \subset G_r$ is called *r -stable* if each $[f^k(\tau)]$ has the same number of illegal r -turns.




Fig. 7. Dynamics of elements of P_r under iteration of f .




Fig. 8. Example: The path $\sigma := [y_0, y_4]$ of height r contains four maximal r -legal subpaths $[y_j, y_{j+1}]$, $j = 0, \dots, 3$. The subpaths $[x_i, z_i]$, $i = 1, 2, 3$, are r -cancellation areas. Then the r -decomposition of σ can be obtained by splitting at x_i and z_i .

Observe that for any reduced path $\sigma \subset G_r$, the number of illegal r -turns in $[f(\sigma)]$ does not exceed the number of illegal r -turns in σ . Therefore there exists a minimal k_0 such that $[f^{k_0}(\sigma)]$ is r -stable. We call this path *the r -stable image* of σ and denote it by $(\sigma)_{\text{stab}}$.

Lemma 7.2 (see [1, Lemma 4.2.6]). *Suppose that H_r is an exponential stratum and $\sigma \subset G_r$ is an r -stable path. Then σ has a splitting into subpaths that are either r -legal or elements of P_r .*

We may assume that no two consecutive factors of this splitting are r -legal. We call such splitting the *r -decomposition* of σ (see Fig. 8). The r -legal paths and the r -cancellation areas in the r -decomposition have simple dynamics under iteration of f : the r -legal paths stretch with the factor λ_r (measured by L_r), and the r -cancellation areas map to r -cancellation areas.

In Sec. 9, we show that, given a reduced edge path σ in G_r , one can compute its r -stable image $(\sigma)_{\text{stab}}$ and the r -decomposition of $(\sigma)_{\text{stab}}$.

8. r -Cancellation Areas

Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track with the maximal filtration $\emptyset = G_0 \subset \dots \subset G_N = \Gamma$.

Definition 8.1. (1) Let p, q be reduced paths in Γ with the same initial point. By $I(p, q)$ we denote the largest common initial subpath of p and q . Then $p \equiv I(p, q) \cdot p'$ and $q \equiv I(p, q) \cdot q'$ for some paths p', q' . We denote $\Lambda(p, q) := (p', q')$.

(2) Let $\tau \equiv \bar{p} \cdot q$ be a reduced path in Γ . For $k \geq 1$, we set

$$(p_k, q_k) \equiv \Lambda([f^k(p)], [f^k(q)]) \quad \text{and} \quad I_k \equiv I([f^k(p)], [f^k(q)]).$$

Then $[f^k(\tau)] \equiv \bar{p}_k \cdot q_k$. The occurrence $y^k := \alpha(p_k) = \alpha(q_k)$ in $[f^k(\tau)]$ is called the *k -successor* of $y := \alpha(q)$.

Definition 8.2. Let H_r be an exponential stratum. Let $\tau \equiv \bar{p} \cdot q$ be a reduced path in G_r , where p and q are r -legal paths. For $k \in \mathbb{N}$, let c_k be the maximal initial subpath of p such that $[f^k(c_k)]$ is a subpath of I_k and the terminal (possibly partial) edge of c_k lies in H_r if it exists. Clearly, $c_1 \subseteq c_2 \subseteq \dots \subseteq p$. Let $(p)_{\text{min}}$ be

the minimal path containing all c_k . Note that the terminal point of $(p)_{\min}$ is not necessarily a vertex. We define $(q)_{\min}$ analogously.

Definition 8.3. Let H_r be an exponential stratum. Let τ be a reduced path in G_r . An occurrence of a vertex y in τ is called an *r-cancellation point* in τ if τ contains a subpath $\bar{a}b$, where a and b are nontrivial partial edges such that $\alpha(a) = \alpha(b) = y$ and the full edges containing a and b form an illegal r -turn.

Lemma 8.4. *Let H_r be an exponential stratum. Suppose that $\tau \equiv \bar{p}q$ is a reduced path in G_r such that the paths p and q are r -legal and the common initial point of p and q is an r -cancellation point in τ . Then the following statements hold:*

- (1) *The initial and the terminal (possibly partial) edges of the paths $(p)_{\min}$ and $(q)_{\min}$ lie in H_r .*
- (2) *The number of r -edges in $(p)_{\min}$ and in $(q)_{\min}$, including their terminal (possibly partial) r -edges, is bounded from above by a computable natural number n_{critical} depending only on f . In particular, $L_r((p)_{\min}) < L_{\text{critical}}$ and $L_r((q)_{\min}) < L_{\text{critical}}$, where $L_{\text{critical}} := \max\{L_r(E) \mid E \in H_r^1\} \cdot n_{\text{critical}}$.*

Proof. (1) The initial (partial) edges of the paths $(p)_{\min}$ and $(q)_{\min}$ lie in H_r , since the common initial vertex of these paths is an r -cancellation point in τ . The terminal (partial) edges of $(p)_{\min}$ and $(q)_{\min}$ lie in H_r by definition of these paths.

(2) We will use Lemma 4.2.2 from [1] and the notion a *k-protected subpath* introduced in the paragraph preceding this lemma.

We claim that our statement holds for $n_{\text{critical}} := 2\lceil K \rceil + 2$, where K is the constant from [1, Lemma 4.2.2]. Suppose the contrary, that is, the number of r -edges in $(p)_{\min}$ is larger than $2\lceil K \rceil + 2$. Then there exists an r -edge E in $(p)_{\min}$ such that there are $\lceil K \rceil$ r -edges to the left and to the right of E in $(p)_{\min}$. It follows that the subpath E is $\lceil K \rceil$ -protected in τ . By this lemma, τ can be split at the endpoints of E , that is $[f^k(E)]$ remains uncanceled in $[f^k(\tau)]$ for each $k \in \mathbb{N}$. This contradicts Definition 8.2. \square

Definition 8.5. Let H_r be an exponential stratum. Suppose that $\tau \equiv \bar{p}q$ is a reduced path in G_r such that the paths p and q are r -legal and the common initial point y of p and q is an r -cancellation point in τ . We say that y is *non-deletable* in τ if for every $k \geq 1$ the k -successor y^k is an r -cancellation point in $[f^k(\tau)]$. We say that y is *deletable* in τ if this does not hold.

If y is non-deletable in τ , we call the path $A := \overline{(p)_{\min}}(q)_{\min}$ the *r-cancellation area* (in τ). The number $a := L_r((p)_{\min}) = L_r((q)_{\min})$ is called the *r-cancellation radius* of A .

Lemma 8.6. *Let H_r be an exponential stratum. The set \mathcal{A} of r -cancellation areas and the set P_r coincide. In particular, \mathcal{A} is f -invariant.*

Proof. Let $A \in \mathcal{A}$. By the above definition, A contains a unique non-deletable r -cancellation point. Therefore A satisfies the property (i) in Definition 7.1. Properties (ii) and (iii) are satisfied by Lemma 8.4. Thus $A \in P_r$.

Conversely, let $B \in P_r$. By (i), we can write $B = \bar{p}q$, where p and q are r -legal and the common initial vertex of p and q is a non-deletable r -cancellation point in B . We set $A := (\bar{p})_{\min}(q)_{\min}$. Clearly A is a subpath of B . Since one path from P_r cannot properly contain another one, we have $A = B$; hence $B \in \mathcal{A}$. \square

Proposition 8.7. *Let H_r be an exponential stratum.*

- (1) *Given two r -legal paths β, γ in G_r with $L_r(\beta) > 0$ and $L_r(\gamma) > 0$, there exists at most one r -cancellation area A such that $\bar{\beta}$ and γ are some initial and terminal subpaths of A .*
- (2) *The number of r -cancellation areas is at most $M_r := m_r^2 n_r^2$, where m_r is the number of edges in H_r and n_r is the number of sequences (p_1, p_2, \dots, p_s) , where all p_i are r -legal edge paths in H_r with $\sum_{i=1}^s L_r(p_i) \leq L_{\text{critical}}$, $s \in \{0\} \cup \mathbb{N}$.*

Proof. (1) Without loss of generality, we may assume that $L_r(\beta) = L_r(\gamma)$. Suppose that $A \equiv \bar{p}q$ is an r -cancellation area, where p and q are r -legal, and β and γ are terminal subpaths of p and q , respectively. Let k be the minimal natural number such that

$$\lambda_r^k \cdot L_r(\beta) > L_{\text{critical}}.$$

Then $L_r([f^k(\beta)]) = L_r([f^k(\gamma)]) > L_{\text{critical}}$. This implies that $[f^k(\bar{p}q)] = [\bar{b}c]$, where b is obtained from $[f^k(\beta)]$ by deleting the maximal initial subpath lying in G_{r-1} and c is obtained analogously from $[f^k(\gamma)]$. Hence $[f^k(\bar{p}q)]$ and so $\bar{p}q$ are completely determined by β and γ .

(2) We begin by setting some additional notation. For any reduced path τ in G_r , we can write $\tau \equiv c_0 \tau_1 c_1 \cdots \tau_s c_s$, where the paths c_1, c_2, \dots, c_{s-1} lie in G_{r-1} and are nontrivial, the paths $\tau_1, \tau_2, \dots, \tau_s$ lie in H_r and are nontrivial, and c_0, c_s lie in G_{r-1} or are trivial. We denote $\tau \cap H_r := (\tau_1, \tau_2, \dots, \tau_s)$.

Let τ'_s be the path obtained from τ_s by deleting the terminal partial edge of τ_s if it exists. We set $[\tau \cap H_r] := (\tau_1, \dots, \tau_{s-1}, \tau'_s)$ if τ'_s is not empty and $[\tau \cap H_r] := (\tau_1, \dots, \tau_{s-1})$ if τ'_s is empty.

The following claim is proved in the proof of Lemma 4.2.5 in [1]:

For any two sequences $\mu := (\mu_1, \mu_2, \dots, \mu_s)$, $\sigma := (\sigma_1, \sigma_2, \dots, \sigma_t)$, where $\mu_1, \dots, \mu_s, \sigma_1, \dots, \sigma_t$ are r -legal edge paths in H_r , and for any two edges E_1, E_2 in H_r , there exists at most one r -cancellation area $A \equiv \bar{p}q$ such that the paths p and q are r -legal, $[p \cap H_r] = \mu$, $[q \cap H_r] = \sigma$, and the terminal (possibly partial) edge of p is a part of E_1 , and the terminal (possibly partial) edge of q is a part of E_2 .

Clearly, this claim and Lemma 8.4(2) imply the statement (2). \square

Definition 8.8. Let H_r be an exponential stratum. Let x be a point in an r -edge E . The $(l, L_r)_E$ -coordinates of x is the pair $(l(p), L_r(p))$, where p is the initial segment of E with $\omega(p) = x$.

The following lemma follows from the fact that $f : \Gamma \rightarrow \Gamma$ lies in the class $\mathcal{P}\mathcal{L}\mathcal{H}\mathcal{E}$.

Lemma 8.9. Let H_r be an exponential stratum. For each r -edge E and each $m \in \mathbb{N}$, the set $\{x \in E \mid f^m(x) = x\}$ is finite. Given such E and m , we can compute the set of $(l, L_r)_E$ -coordinates of all points of this set.

Theorem 8.10. There is an algorithm finding all r -cancellation areas of f .

Proof. Let \mathcal{A} be the set of all r -cancellation areas. Let U be the set of all endpoints of all r -cancellation areas. The set U is f -invariant and lies in H_r by Lemma 7.1, and $|U| \leq M_r$ by Proposition 8.7. We consider the subset $U' := \{f^{M_r}(u) \mid u \in U\}$ of U . Then each point of U' is fixed by f^m for some $0 < m \leq M_r$. Therefore U' is contained in the set

$$\overline{U'} := \bigcup_{E \in H_r^1} \bigcup_{m=1}^{M_r} \{x \in E \mid f^m(x) = x\}.$$

This set is finite and computable by Lemma 8.9. Then the set

$$\overline{U} := \{u \in H_r \mid f^{M_r}(u) \in \overline{U'}\}$$

is finite and computable, and contains U .

Let \mathcal{P} be the set of nontrivial initial segments ρ of r -edges with $\omega(\rho) \in \overline{U}$. Suppose that A is an r -cancellation area. We write $A \equiv \bar{p}q$, where p and q are r -legal paths. Then p and q have terminal segments β and γ , respectively, which lie in \mathcal{P} . By Proposition 8.7(1), A is completely determined by β and γ . The proof of this proposition gives us the following algorithm constructing all elements of \mathcal{A} :

- (a) Compute $\mathcal{L} = \min\{L_r(\rho) \mid \rho \in \mathcal{P}\}$ and the minimal $k \in \mathbb{N}$ such that $\lambda_r^k \cdot \mathcal{L} > L_{\text{critical}}$. Denote $\mathcal{A}_k := \{[f^k(A)] \mid A \in \mathcal{A}\}$. Clearly, $\mathcal{A}_k \subseteq \mathcal{A}$.
- (b) Compute the set Ψ_k of all paths of the form $\bar{b}c$, where $\alpha(b) = \alpha(c)$ is a vertex, b and c are nontrivial terminal subpaths of $[f^k(\beta)]$ and of $[f^k(\gamma)]$ for some $\beta, \gamma \in \mathcal{P}$, $L_r(b) = L_r(c) \leq L_{\text{critical}}$, and the first (possibly partial) edges of b and c form a nondegenerate illegal r -turn. Then $\mathcal{A}_k \subseteq \Psi_k$.
- (c) Compute the set Ψ of reduced paths $d \subset G_r$ such that $[f^k(d)] \in \Psi_k$. Then $\mathcal{A} \subseteq \Psi$.
- (d) For each $d \in \Psi$ verify whether d is an r -cancellation area. For that compute $[f^s(d)]$ for $s = 0, \dots, M_r$. Then d is an r -cancellation area if and only if a repetition occurs, that is, if $[f^i(d)] = [f^j(d)]$ for some $0 \leq i < j \leq M_r$, and the points (i) and (ii) of Definition 7.1 are satisfied for $[f^s(d)]$, $s = 0, \dots, j$. Then \mathcal{A} consists of the elements of Ψ which are r -cancellation areas. \square

Remark 8.11. Let H_r be an exponential stratum and let N_r be the number of r -cancellation areas. Then, for each r -cancellation area A , the path $[f^{N_r}(A)]$ is an IPN path of height r , and each IPN path of height r has this form. Thus, by Theorem 8.10, we can find all IPN paths of height r .

Definition 8.12. Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track. We say that f satisfies (RTT-iv) if there is a natural number P such that for each exponential stratum H_r and each r -cancellation area A of f , the r -cancellation area $[f^P(A)]$ is an edge path.

Corollary 8.13. *Given a PL-relative train track $f : \Gamma \rightarrow \Gamma$, one can find a natural number P and subdivide Γ so that the induced map will be a PL-relative train track satisfying (RTT-iv) for this P .*

Proof. Let P be the maximum of numbers of r -cancellation areas in Γ over all r . Then (RTT-iv) can be arranged via subdivisions at endpoints of $[f^P(A)]$, where A runs over all r -cancellation areas. These endpoints can be found by Theorem 8.10. \square

9. r -Stable Paths and Their r -Decompositions

Let $f : \Gamma \rightarrow \Gamma$ be a PL-relative train track with the maximal filtration $\emptyset = G_0 \subset \dots \subset G_N = \Gamma$. Let H_r be an exponential stratum in Γ . Recall that a reduced path $\tau \subset G_r$ is called r -stable if each $[f^k(\tau)]$ has the same number of illegal r -turns.

Theorem 9.1. *Let H_r be an exponential stratum in Γ . There exists an algorithm which, given a reduced edge path $\tau \subset G_r$, computes the minimal $i_0 \geq 0$ such that the path $[f^{i_0}(\tau)]$ is r -stable. In particular, one can check whether τ is r -stable or not.*

Proof. It suffices to find some j_0 such that the path $[f^{j_0}(\tau)]$ is r -stable. Then i_0 is the minimal integer such that $0 \leq i_0 \leq j_0$ and the paths $[f^{i_0}(\tau)]$ and $[f^{j_0}(\tau)]$ have the same number of illegal r -turns.

Let y_1, \dots, y_k be all r -cancellation points in τ . Let y_0 and y_{k+1} be the initial and the terminal points of τ , respectively. Let τ_i be the subpath of τ from y_i to y_{i+1} , $i = 0, \dots, k$.

For each $1 \leq i \leq k$, we check (using Theorem 8.10) whether $\tau_{i-1}\tau_i$ contains an r -cancellation area or not. If some $\tau_{i-1}\tau_i$ does not contain an r -cancellation area, then y_i is deletable in $\tau_{i-1}\tau_i$ and we can find T_i such that $[f^{T_i}(\tau_{i-1}\tau_i)]$ is r -legal or trivial. Then the number of r -cancellation points in $[f^{T_i}(\tau)]$ is smaller than k and we proceed by induction.

Now suppose that each $\tau_{i-1}\tau_i$ contains an r -cancellation area A_i . If the interiors of each neighboring A_i, A_{i+1} do not overlap, then τ is r -stable. If the interiors of some neighboring A_i, A_{i+1} overlap, then there exists T such that the number of r -cancellation points in $[f^T(\tau)]$ is smaller than k and we proceed by induction. \square

Definition 9.2. Let H_r be an exponential stratum in Γ .

- (1) For any reduced path $\mu \subset G_r$, there exists a minimal i_0 such that the path $[f^{i_0}(\mu)]$ is r -stable. The latter path is denoted by $(\mu)_{\text{stab}}$.
- (2) Let τ be a reduced r -stable path in G_r . Then τ can be uniquely written as $\tau \equiv b_0 \cdot A_1 \cdot b_1 \cdot \dots \cdot A_k \cdot b_k$, where b_0, \dots, b_k are r -legal or trivial paths in G_r and A_1, \dots, A_k are all r -cancellation areas in τ . We call such decomposition the r -decomposition of τ and denote $\mathbf{N}_r(\tau) := k$. The subpaths b_i and A_i are called *blocks* of τ . Any subpath of τ which is a concatenation of blocks of τ is called a *block subpath* of τ .

Remark 9.3. Suppose that τ has the r -decomposition

$$\tau \equiv b_0 \cdot A_1 \cdot b_1 \cdot \dots \cdot A_k \cdot b_k.$$

Then, for every $i \geq 1$, the path $[f^i(\tau)]$ has the r -decomposition

$$[f^i(\tau)] \equiv b_0^i \cdot A_1^i \cdot b_1^i \cdot \dots \cdot A_k^i \cdot b_k^i,$$

where $b_j^i \equiv [f^i(b_j)]$ and $A_j^i \equiv [f^i(A_j)]$ for all possible j .

Lemma 9.4. *Let μ be an r -stable path such that the r -decomposition of μ starts (ends) with an r -cancellation area. Suppose that μ_1 is obtained from μ by deleting the first (last) (possibly partial) edge. Then $\mathbf{N}_r((\mu_1)_{\text{stab}}) = \mathbf{N}_r(\mu) - 1$.*

Proof. Let $\mu \equiv A_1 \cdot b_1 \cdot \dots \cdot A_k \cdot b_k$ be the r -decomposition of μ . If $k = 1$, then the statement follows from the definition of an r -cancellation area. If $k > 1$, one should additionally use the previous remark. \square

In Sec. 11, we need a stronger property than r -stability, namely the r -superstability.

Definition 9.5. *Let H_r be an exponential stratum. A reduced f -path $\tau \subset G_r$ is called r -superstable if τ and $[\tau f(\tau)]$ are r -stable and all r -cancellation areas in these paths are edge paths.*

Note that if τ is r -superstable, then $[f^i(\tau)]$ is r -superstable for all $i \geq 0$.

Lemma 9.6. *Suppose that $f : \Gamma \rightarrow \Gamma$ is a PL-relative train track satisfying the property (RTT-iv). Let H_r be an exponential stratum. For any reduced f -path $\tau \subset G_r$, one can compute a natural number $S = S(\tau)$ such that the path $[f^S(\tau)]$ is r -superstable.*

Proof. By Theorem 9.1, we can compute numbers S_1 and S_2 such that $[f^{S_1}(\tau)]$ and $[f^{S_2}([\tau f(\tau)])]$ are r -stable. We set $S = \max\{S_1, S_2\} + P$, where P is the (RTT-iv) constant of Definition 8.12. \square

10. r -Perfect and \mathcal{A}_r -Perfect Paths

Let H_r be an exponential stratum and $\tau \equiv E_1 E_2 \cdots E_k$ be an f -path of height r . In general, there may be cancellations in passing from τ to $E_2 \cdots E_k \cdot f(E_1)$, and in passing from $\hat{f}^i(\tau)$ to $\hat{f}^{i+1}(\tau)$ for $i \geq 1$.

Below we define r -perfect and \mathcal{A}_r -perfect paths. If τ is such a path, then these cancellations can be well controlled. In particular, we will show that all f -paths $[f^i(\tau)]$, $i \in \mathbb{N}$, considered as vertices of D_f , lie in the τ -subgraph. This will be used in Sec. 14 to solve the Finiteness Problem and the Membership Problem with the help of Corollaries 12.2 and 12.3.

If τ is an arbitrary reduced f -path of height r , then Theorem 11.6 helps to improve it so that the resulting path will be r -perfect, or \mathcal{A}_r -perfect, or will lie in G_{r-1} , or will have small length.

Definition 10.1. Let H_r be an exponential stratum. An edge path $\tau \subset G_r$ is called r -perfect if the following conditions are satisfied:

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

A vertex in D_f is called r -perfect if the corresponding f -path in Γ is r -perfect.

Note that these conditions imply that $[\tau f(\tau)]$ is r -legal. In the following proposition we formulate some important properties of r -perfect paths.

Proposition 10.2. Let H_r be an exponential stratum and let τ be an r -perfect path in G_r . Then the following statements are satisfied:

- (1) For every $i \geq 0$, the path $\hat{f}^{i+1}(\tau) \subset G_r$ is r -legal, contains edges from H_r , and $L_r(\hat{f}^{i+1}(\tau)) \geq L_r(\hat{f}^i(\tau))$.
- (2) For every $i \geq 0$, the path $[f^i(\tau)]$ is r -perfect.
- (3) For every $i \geq 0$, the vertex $[f^i(\tau)]$ of D_f lies in the τ -subgraph. Moreover, $[f^i(\tau)] \equiv \hat{f}^{m_i}(\tau)$ for some computable m_i satisfying $m_0 = 0$, $m_i < m_{i+1}$. In particular, $L_r(\hat{f}^{m_i}(\tau)) = \lambda_r^i L_r(\tau)$, and the τ -subgraph is infinite.

Proof. First we check (1) for $i = 0$. Let $\tau \equiv E_1 E_2 \cdots E_k$ be the edge decomposition of τ . By (iii), we have $\hat{f}(\tau) \equiv E_2 \cdots E_k \cdot [f(E_1)]$. This path is r -legal as a subpath of the r -legal path $[\tau f(\tau)]$. Since E_1 is an r -edge, we have $L_r(\hat{f}(\tau)) = L_r(\tau) + (\lambda_r - 1)L_r(E_1) > L_r(\tau)$.

To prove (1) in the general case, we write $\tau \equiv b_1 c_1 \cdots b_k c_k$, where b_1, \dots, b_k are the maximal subpaths of τ lying in H_r . As above, we can compute $\hat{f}^i(\tau)$ for $i = 1, \dots, l(b_1 c_1)$ and verify (1) for these i . In particular, we have $\hat{f}^{l(b_1 c_1)}(\tau) \equiv b_2 c_2 \cdots b_k c_k [f(b_1)][f(c_1)]$. This path is r -perfect, and we can proceed by induction.

(2) This statement follows from condition (RTT-i) and Proposition 3.3.

(3) Computing as in the proof of (1), we have $\hat{f}^{l(\tau)}(\tau) \equiv [f(\mu)]$. By (2), this path is r -perfect. Therefore $m_1 = l(\tau)$ and the numbers m_i , $i \geq 2$, can be computed inductively. Finally, $L_r(\hat{f}^{m_i}(\tau)) = L_r([f^i(\tau)]) = \lambda_r^i L_r(\tau)$. Since $\lambda_r > 1$, the τ -subgraph is infinite. \square

Definition 10.3. Let H_r be an exponential stratum. A reduced f -path $\tau \subset G_r$ containing edges from H_r is called \mathcal{A}_r -perfect if

- (i) τ is r -stable and the r -cancellation areas in τ are edge paths,
- (ii) the r -decomposition of τ begins with an r -cancellation area, that is, it has the form $\tau \equiv A_1 b_1 \cdots A_k b_k$, $k \geq 1$, and
- (iii) $[\tau f(\tau)] \equiv \tau \cdot [f(\tau)]$ and the turn at the point between τ and $[f(\tau)]$ is legal.

A vertex in D_f is called \mathcal{A}_r -perfect if the corresponding f -path in Γ is \mathcal{A}_r -perfect.

Note that the first edge of such τ lies in H_r .

Proposition 10.4. Let H_r be an exponential stratum and let τ be an \mathcal{A}_r -perfect path in G_r with the r -decomposition $\tau \equiv A_1 b_1 \cdots A_k b_k$.

For $1 \leq j \leq k$, we set $\tau_{0,j} \equiv [A_j b_j \cdots A_k b_k f(A_1 b_1 \cdots A_{j-1} b_{j-1})]$ and for $i \geq 1$ we set $\tau_{i,j} \equiv [f^i(\tau_{0,j})]$. Then the following statements are satisfied:

- (1) For any $1 \leq j \leq k$ and $i \geq 0$ the path $\tau_{i,j}$ is \mathcal{A}_r -perfect.
- (2) For any $1 \leq j \leq k$ and $i \geq 0$ the vertex $\tau_{i,j}$ of D_f lies in the τ -subgraph. Moreover, $\tau_{i,j} \equiv \hat{f}^{m_{i,j}}(\tau)$ for some computable $m_{i,j}$ satisfying $m_{0,1} = 0$, $m_{i,j} < m_{i,j+1}$, and $m_{i,k} < m_{i+1,1}$.
- (3) All \mathcal{A}_r -perfect vertices of the τ -subgraph are $\tau_{i,j}$, $1 \leq j \leq k$, $i \geq 0$.
- (4) For every vertex σ in the τ -subgraph, at least one of the paths σ , $\hat{f}(\sigma), \dots, \hat{f}^{l(\sigma)}(\sigma)$ coincides with $\tau_{i,j}$ for some i, j .

Proof. (1) For $i = 0$, the statement follows from Definition 10.3; for $i \geq 1$ it follows from the fact that the reduced form of the f -image of an \mathcal{A}_r -perfect path is \mathcal{A}_r -perfect.

Statement (2) For $i = 0$ follows from the facts that $\tau_{0,1} \equiv \tau$ and $\tau_{0,j+1} \equiv \hat{f}^{l(A_j b_j)}(\tau_{0,j})$, $1 \leq j \leq k-1$. Furthermore, $\tau_{1,1} \equiv \hat{f}^{l(A_k b_k)}(\tau_{0,k})$. For $i \geq 1$, the statement follows by induction with the help of Remark 9.3.

(3) We denote $A_j^i := [f^i(A_j)]$ and $b_j^i := [f^i(b_j)]$ for $i \geq 0$ and $1 \leq j \leq k$. Any vertex σ of the τ -subgraph has the form $[C'' A_{j+1}^i b_{j+1}^i \cdots A_k^i b_k^i f(A_1^i b_1^i \cdots A_{j-1}^i b_{j-1}^i C'')]$, where $C' C'' \equiv A_j^i b_j^i$. If C' or C'' is trivial, then $\sigma \equiv \tau_{i,j}$ or $\sigma \equiv \tau_{i,j+1}$, and σ is \mathcal{A}_r -perfect by (1). If C' and C'' are nontrivial, then σ does not begin with an r -cancellation area, hence σ is not \mathcal{A}_r -perfect.

(4) Using notation in (3), we have $\hat{f}^{l(C'')}(\sigma) \equiv \tau_{i,j+1}$ for $1 \leq j \leq k-1$ and $\hat{f}^{l(C'')}(\sigma) \equiv \tau_{i+1,1}$ for $j = k$. \square

11. r -Perfect and \mathcal{A}_r -Perfect Vertices in μ -Subgraphs

From here and to the end of the paper, we work with the PL-relative train track $f : \Gamma \rightarrow \Gamma$ which satisfies (RTT-iv). Let $R_\star = R_\star(f)$ be the maximum of l -lengths of r -cancellation areas over all r . The main result of this section is Theorem 11.6.

Lemma 11.1. *Let H_r be an exponential stratum. Let $\mu \equiv \sigma\tau$ be a reduced f -path such that σ is a nontrivial path in G_{r-1} with endpoints in H_r and τ is an r -legal path in G_r with the first and the last edges from H_r . Then $\hat{f}^{l(\sigma)}(\mu)$ is r -perfect.*

Proof. Denote $\mu' := \hat{f}^{l(\sigma)}(\mu)$. Clearly, $\mu' \equiv \tau[f(\sigma)]$ and μ' is r -legal. It remains to prove that the turn between μ' and $[f(\mu')]$ in $\mu' \cdot [f(\mu')]$ is legal. By (RTT-ii), $[f(\sigma)]$ is a nontrivial path in G_{r-1} . So, the last edge of μ' lies in G_{r-1} . On the other hand, the first edge of $[f(\mu')] \equiv [f(\tau)][f^2(\sigma)]$ lies in H_r . Therefore, the turn under consideration is mixed and hence legal. \square

Proposition 11.2. *Let H_r be an exponential stratum and let $\mu \subset G_r$ be a reduced f -path such that μ is r -legal and $[\mu f(\mu)]$ is r -legal or trivial. After several applications of \hat{f} , one can obtain an f -path μ' satisfying one of the following conditions:*

- (1) μ' lies in G_{r-1} ,
- (2) μ' is r -perfect, or
- (3) $l(\mu') \leq \|f\|$.

Proof. We induct on the number of r -edges in μ . We assume that μ contains an r -edge, otherwise we have (1) for $\mu' := \mu$. Write $\mu \equiv b_1 \cdot b_2$, where b_1 lies in G_{r-1} or is trivial, and the first edge of b_2 lies in H_r . Then $\hat{f}^{l(b_1)}(\mu)$ has the first edge in H_r , it satisfies all conditions for μ , and it has the same number of r -edges as μ . Thus, we may assume from the beginning that the first edge E of μ lies in H_r . Set $I := I(\bar{\mu}, [f(\mu)])$. Both I and $f(E)$ are initial subpaths of $[f(\mu)]$.

Case 1. Suppose that $f(E)$ lies in I .

We may assume that $l(\mu) \geq \|f\| + 1$. Then $\mu_1 := [\bar{E}\mu f(E)]$ is a subpath of μ . We may assume that μ_1 is nontrivial, otherwise we have (3) for $\mu' := \mu_1$. Moreover, $[\mu_1 f(\mu_1)]$ is a subpath of μ , $[f(\mu)]$, or $[\mu f(\mu)]$. Hence, μ_1 satisfies the assumptions of this proposition. Since μ_1 contains fewer r -edges than μ , we can apply induction.

Case 2. Suppose that $f(E)$ is longer than I (see Fig. 9).

We may assume that $l(\mu) \geq l(I) + 2$, otherwise $l(\mu) \leq 1 + l(I) \leq l(f(E)) \leq \|f\|$ and we have (3). Let E_1 be the second edge of μ . We set $\mu_1 := [\bar{E}\mu f(E)]$. The path μ_1 is r -legal as a subpath of $[\mu f(\mu)]$, it begins with E_1 and ends with the last edge of $f(E)$.

Suppose that E_1 is an r -edge. Then $[f(\mu_1)]$ begins with the first edge of $f(E_1)$. Therefore the turn between μ_1 and $[f(\mu_1)]$ coincides with the turn between $f(E)$




Fig. 9. An illustration to Case 2 in the proof of Proposition 11.2.

and $f(E_1)$. This turn is legal, since the turn between E and E_1 is an r -turn in the r -legal path μ . Hence, μ_1 is r -perfect and we have (2) for $\mu' := \mu_1$.

If E_1 lies in G_{r-1} , then we apply Lemma 11.1 to μ_1 and construct μ' satisfying (2). \square

Proposition 11.3. *Let H_r be an exponential stratum. Let $\mu \subset G_r$ be a reduced r -legal f -path. Suppose that $[\mu f(\mu)]$ contains a non-deletable r -cancellation point and that the r -cancellation area in $[\mu f(\mu)]$ is an edge path. After several applications of \hat{f} and f , one can obtain an f -path μ' satisfying one of the following conditions:*

- (1) μ' lies in G_{r-1} ,
- (2) μ' is r -perfect,
- (3) μ' is \mathcal{A}_r -perfect, or
- (4) $l(\mu') \leq R_\star + \|f\|$.

Proof. For $i = 0, 1$, let z^i be the unique non-deletable r -cancellation point in the path $[f^i(\mu)f^{i+1}(\mu)]$ and let $A(z^i)$ be the r -cancellation area in this path. We have $A(z^1) \equiv [f(A(z^0))]$. Let a and b be the initial and the terminal vertices of $A(z^0)$. Then $f(a)$ and $f(b)$ are the initial and the terminal vertices of $A(z^1)$. We induct on the number of r -edges in μ .

Let E be the first edge of μ . Arguing as in the proof of Proposition 11.2, we may assume that the first edge E of μ lies in H_r . Set $I := I(\bar{\mu}, [f(\mu)])$. Both I and $f(E)$ are initial subpaths of $[f(\mu)]$.

Case 1. Suppose that $f(E)$ lies in I .

Then $\mu_1 := [\bar{E}\mu f(E)]$ is a subpath of μ , hence μ_1 is r -legal or trivial. We may assume that μ_1 is nontrivial, otherwise we have (4) for $\mu' := \mu_1$. If $[\mu_1 f(\mu_1)]$ is r -legal or trivial, then we apply Proposition 11.2.

Suppose that $[\mu_1 f(\mu_1)]$ is not r -legal and nontrivial, so it contains a unique r -cancellation point, say x . If x is non-deletable, then $[\mu_1 f(\mu_1)]$ is a subpath of $[\mu f(\mu)]$, and the r -cancellation area in $[\mu_1 f(\mu_1)]$ is $A(z^0)$. In this case we apply induction to μ_1 . If x is deletable, then we find S such that $[f^S([\mu_1 f(\mu_1)])]$ is r -legal or trivial and apply Proposition 11.2 to $[f^S(\mu_1)]$ (as μ_1 , this path is r -legal or trivial; we may assume that it is nontrivial, otherwise we have (4) for $\mu' := [f^S(\mu_1)]$).




Fig. 10. An illustration to Case 2 in the proof of Proposition 11.3.

Case 2. Suppose that $f(E)$ is longer than I (see Fig. 10).

We may assume that E lies to the left of a , otherwise $l(\mu) \leq l(A(z^0)) + l(I) \leq R_* + l(f(E)) \leq R_* + \|f\|$ and we have (4). Then $[f(E)]$ lies in $[f(\mu)]$ to the left of $f(a)$, in particular, z^0 lies to the left of $f(a)$.

Let p be the initial subpath of μ until the vertex a . We claim that $\mu' := \hat{f}^{l(p)}(\mu)$ satisfies (3) or (4). To prove this, we first observe that μ' is the subpath of $[\mu f(\mu)]$ from a to $f(a)$. If $f(a)$ lies between z^0 and b , then $l(\mu') \leq R_*$, and we have (4). Thus, assume that $f(a)$ lies to the right of b in $[f(\mu)]$.

We prove that μ' is \mathcal{A}_r -perfect. First note that μ' has the r -decomposition $\mu' \equiv A(z^0) \cdot \ell$, where ℓ is the subpath of $[f(\mu)]$ from b to $f(a)$. We have $[f(\mu')] \equiv A(z^1) \cdot [f(\ell)]$. The turn in $\mu' \cdot [f(\mu')]$ at the point between μ' and $[f(\mu')]$ coincides with the turn in $[f(\mu)]$ at $f(a)$. This turn is legal, since $[f(\mu)]$ is r -legal and the first edge of the r -cancellation area $A(z^1)$ lies in H_r . Thus, μ' is \mathcal{A}_r -perfect. \square

Lemma 11.4. *Let H_r be an exponential stratum and let μ be a reduced r -stable f -path in G_r . After several applications of \hat{f} and f , one can obtain an f -path μ' with one of the following properties:*

- (1) μ' is trivial,
- (2) μ' is r -legal, or
- (3) μ' is r -superstable, the r -decomposition of μ' starts with an r -cancellation area, and $\mathbf{N}_r(\mu') \leq \mathbf{N}_r(\mu) + 1$.

Proof. We may assume that the paths which we obtain in the process below are nontrivial. Using Lemma 9.6, we may assume that μ is r -superstable. Let $\mu \equiv b_0 A_1 \cdots A_k b_k$ be the r -decomposition of μ . We shall analyze the case where $k \geq 1$ and b_0 is nonempty. Consider $\mu_1 := \hat{f}^{l(b_0)}(\mu) \equiv [A_1 \cdots A_k b_k f(b_0)]$.

First suppose that $[f(b_0)]$ completely cancels in $A_1 \cdots A_k b_k [f(b_0)]$. If A_1 remains there, then we set $\mu' := [f^S(\mu_1)]$, where S is as in Lemma 9.6. Otherwise, μ_1 is a proper initial segment of A_1 , hence $(\mu_1)_{\text{stab}}$ is r -legal, and we set $\mu' := (\mu_1)_{\text{stab}}$.

Now suppose that $[f(b_0)]$ does not completely cancel in $A_1 \cdots A_k b_k [f(b_0)]$. Then μ_1 is a subpath of $[\mu f(\mu)]$.

If A_1 remains in μ_1 , then μ_1 is a block subpath (see Definition 9.2) of the r -stable path $[\mu f(\mu)]$. Hence μ_1 is r -stable, its r -decomposition starts with A_1 , and $\mathbf{N}_r(\mu_1) \leq \mathbf{N}_r(\mu) + 1$. Applying a power of f , we may assume that μ_1 is r -superstable.

If the cancellations meet A_1 , then $\mu_1 \equiv A'_1[f(b_0)]'$, where A'_1 is a nontrivial initial proper subpath of A_1 and $[f(b_0)]'$ is a nontrivial terminal subpath of $[f(b_0)]$. Write $[f(b_0)]' \equiv [f(b'_0)]$ for some terminal subpath b'_0 of b_0 . Then $\hat{f}^{l(A'_1)}(\mu_1) \equiv [f(b_0)]'[f(A'_1)] \equiv [f(b'_0 A'_1)]$. By Lemma 9.4, $(b'_0 A'_1)_{\text{stab}}$ is r -legal. Hence $\mu' := (\hat{f}^{l(A'_1)}(\mu_1))_{\text{stab}}$ is r -legal. \square

Proposition 11.5. *Let H_r be an exponential stratum and let μ be a reduced f -path in G_r . After several applications of \hat{f} and f , one can obtain an f -path μ_1 with one of the following properties:*

- (1) μ_1 is trivial,
- (2) μ_1 is r -legal,
- (3) μ_1 is \mathcal{A}_r -perfect, or
- (4) μ_1 is an r -cancellation area.

Proof. We may assume that μ is r -stable and that the paths which we obtain in the process below are nontrivial. By Lemma 11.4, we can pass to a new μ and assume that μ is r -superstable and has the r -decomposition $\mu \equiv A_1 b_1 A_2 \cdots A_k b_k$; the value $\mathbf{N}_r(\mu)$ increases by at most 1. If b_k is a nontrivial path in G_{r-1} , then μ is A -perfect. So, we assume that b_k is empty or $L_r(b_k) > 0$. Applying f several times, we may assume that b_k is empty or $L_r(b_k) > 2\lambda_r L_{\text{critical}}$.

Case 1. Suppose that $k = 1$, that is, $\mu \equiv A_1 b_1$. We may assume that b_1 is nonempty, otherwise we are done. Since $L_r(b_1) > 2\lambda_r L_{\text{critical}}$, we have $\hat{f}^{l(A_1)}(\mu) \equiv [b_1 f(A_1)]$.

If there is no cancellation between b_1 and $[f(A_1)]$, then $[\mu f(\mu)] \equiv A_1 b_1 [f(A_1)] [f(b_1)]$. Since μ is r -superstable, the turn between b_1 and $[f(A_1)]$ is legal, hence μ is \mathcal{A}_r -perfect.

If there is a cancellation between b_1 and $[f(A_1)]$, then $\hat{f}^{l(A_1)}(\mu) \equiv b'_1 [f(A_1)]'$, where b'_1 is an initial subpath of b_1 and $[f(A_1)]'$ is a proper terminal subpath of $[f(A_1)]$. Write $[f(A_1)]' \equiv [f(A'_1)]$ for some proper terminal subpath A'_1 of A_1 . Then $\hat{f}^{l(A_1 b'_1)}(\mu) \equiv [f(A_1)]'[f(b'_1)] \equiv [f(A'_1 b'_1)]$. By Lemma 9.4, $(A'_1 b'_1)_{\text{stab}}$ is r -legal. Hence $\mu' := (\hat{f}^{l(A_1 b'_1)}(\mu))_{\text{stab}}$ is r -legal.

Case 2. Suppose that $k \geq 2$. First we suppose that b_k is nonempty. Since $L_r(b_k) > 2\lambda_r L_{\text{critical}}$, we have $\hat{f}^{l(A_1)}(\mu) \equiv [b_1 A_2 \cdots A_k b_k f(A_1)]$ and the last edge of $[f(A_1)]$ remains uncanceled. Then $\hat{f}^{l(A_1 b_1)}(\mu) \equiv [A_2 \cdots A_k b_k f(A_1) f(b_1)]$ is a block subpath of the r -stable path $[\mu f(\mu)]$, hence it is r -stable. Clearly, it is \mathcal{A}_r -perfect.

Now suppose that b_k is empty. If $[A_k f(A_1)]$ is nontrivial, then $\hat{f}^{l(A_1)}(\mu) \equiv [b_1 A_2 \cdots A_k f(A_1)]$, where the last edge of $[f(A_1)]$ remains uncanceled, and we can proceed as above. If $[A_k f(A_1)]$ is trivial, then $\mathbf{N}_r(\hat{f}^{l(A_1)}(\mu)) = \mathbf{N}_r(\mu) - 2$, and we can proceed by induction. \square

Theorem 11.6. *Let H_r be an exponential stratum. Suppose that μ is a reduced f -path in G_r . After several applications of \hat{f} and f , one can obtain an f -path μ_1 with one of the following properties:*

- (1) μ_1 lies in G_{r-1} ,
- (2) μ_1 is r -perfect,
- (3) μ_1 is \mathcal{A}_r -perfect, or
- (4) $l(\mu_1) \leq R_\star + \|f\|$.

Proof. In view of Proposition 11.5, it suffices to consider the case where μ is r -legal. By Lemma 9.6, we may additionally assume that μ is r -superstable. Applying Proposition 11.2 or Proposition 11.3, we complete the proof. \square

12. Auxiliary Results

Let F be a free group of finite rank with a fixed basis X . For any element $w \in F$ let $|w|$ be the length of w with respect to X . The following theorem was proved by Brinkmann in [4, Theorem 0.1].

Theorem 12.1. *There exists an algorithm which, given an automorphism φ of a free group F of finite rank and given elements $u, v \in F$, verifies whether there exists a natural number N such that $\varphi^N(u) = v$. If such N exists, then the algorithm computes it.*

Corollary 12.2. *There exists an algorithm which, given a finite connected graph Γ and a homotopy equivalence $f : \Gamma \rightarrow \Gamma$ with $f(\Gamma^0) \subseteq \Gamma^0$, and given two edge paths ρ, τ in Γ , decides whether there exists a natural number k such that $f^k(\rho) = \tau$. If such k exists, then the algorithm computes it.*

Proof. First we reduce the problem to the case, where f fixes the endpoints of ρ . Let u_i and v_i be the initial and the terminal vertices of $\rho_i := f^i(\rho)$. Since f acts on the finite set $\Gamma^0 \times \Gamma^0$, there exist natural numbers r, n such that $(u_i, v_i) = (u_{i+n}, v_{i+n})$ for $i \geq r$.

First we check, whether $f^k(\rho) = \tau$ for $k < r$. If yes, we are done; if no, we investigate the case $k \geq r$. Given such k , we can write $k = i + \ell n$ for some $\ell \geq 0$ and $r \leq i < r + n$. So, we have $f^k(\rho) = g^\ell(\rho_i)$, where $g := f^n$. Thus, we have to investigate n problems: does there exist $\ell \geq 0$ such that $g^\ell(\rho_i) = \tau$, $r \leq i < r + n$? Note that g fixes the endpoints of ρ_i .

So, from the beginning, we may assume that f fixes the endpoints of ρ and $\alpha(\rho) = \alpha(\tau)$, and $\omega(\rho) = \omega(\tau)$.

Let Γ_1 be the graph obtained from Γ by adding a new vertex v and two new oriented edges: E_1 from v to $\alpha(\rho)$ and E_2 from v to $\omega(\rho)$. Let $f_1 : \Gamma_1 \rightarrow \Gamma_1$ be the extension of f that maps E_1 to E_1 and E_2 to E_2 . Clearly, f_1 is a homotopy

equivalence which fixes v . Let $(f_1)_v : \pi_1(\Gamma_1, v) \rightarrow \pi_1(\Gamma_1, v)$ be the induced automorphism. We have

$$f^k(\rho) = \tau \Leftrightarrow f_1^k(E_1 \rho \bar{E}_2) = E_1 \tau \bar{E}_2 \Leftrightarrow (f_1)_v^k([[E_1 \rho \bar{E}_2]]) = [[E_1 \tau \bar{E}_2]].$$

Thus, the problem is solvable by Theorem 12.1. \square

Corollary 12.3. *There exists an algorithm which, given a finite connected graph Γ and a homotopy equivalence $f : \Gamma \rightarrow \Gamma$ with $f(\Gamma^0) \subseteq \Gamma^0$, and given two edge paths ρ, τ in Γ , decides whether there exist natural numbers $k > s$ such that $f^k(\rho) = f^s(\tau)$. If such k and s exist, then the algorithm computes them.*

Proof. Let u_i and v_i be the initial and the terminal vertices of $f^i(\rho)$, $i \geq 0$, and let u'_j and v'_j be the initial and the terminal vertices of $f^j(\tau)$, $j \geq 0$. First we decide whether there exist i, j such that $(u_i, v_i) = (u'_j, v'_j)$. If such i, j do not exist, then the desired k, s do not exist. If such i, j exist, we can algorithmically find natural i, j, n with the following properties:

- (1) $(u_i, v_i) = (u'_j, v'_j)$,
- (2) $(u_i, v_i) = (u_{i+n}, v_{i+n})$ and n is minimal,
- (3) $i > j$, and
- (4) $i - j$ is the minimal possible for (1)–(3).

So, we reduce the problem to the following: do there exist $p \geq q \geq 0$ such that $f^{i+pn}(\rho) = f^{j+qn}(\tau)$? We set $\rho_1 := f^i(\rho)$, $\tau_1 := f^j(\tau)$, $g := f^n$. Then the endpoints of ρ_1 and τ_1 coincide and are fixed by g . In this setting we have to decide whether there exist $p \geq q \geq 0$ such that $g^p(\rho_1) = g^q(\tau_1)$.

We extend Γ to Γ' by adding an edge E from v_i to u_i , and we extend g to $g' : \Gamma' \rightarrow \Gamma'$ by setting $g'|_{\Gamma} = g$ and $g'(E) = E$. Then the problem is equivalent to the following: do there exist $p \geq q \geq 0$ such that $g'^p(\rho_1 E) = g'^q(\tau_1 E)$?

Since g' is a homotopy equivalence and $\rho_1 E$ and $\tau_1 E$ are loops based at the same point, and this point is fixed by g' , we have

$$g'^p(\rho_1 E) = g'^q(\tau_1 E) \Leftrightarrow g'^{p-q}(\rho_1 E) = \tau_1 E.$$

Thus, the problem can be reformulated as follows: does there exist $m \geq 0$ such that $g'^m(\rho_1 E) = \tau_1 E$? This can be decided by Theorem 12.1. \square

13. \mathcal{E}_r -Perfect Vertices in μ -Subgraphs

Let H_r be a polynomial stratum. There exists a cyclic permutation σ on the set of r -edges and, for each r -edge E , there exists an edge path c_E (which is trivial or is an edge path in G_{r-1}) such that $f(E) = c_E \cdot \sigma(E) \cdot \bar{c}_E$. Then, for each $i \geq 0$ and for each r -edge E , one can compute a path $c_{i,E}$ (which is trivial or is an edge path in G_{r-1}) such that $f^i(E) \equiv c_{i,E} \cdot \sigma^i(E) \cdot \bar{c}_{i,E}$. For any edge path $\mu \subset G_r$, let $\mathcal{N}(\mu)$

be the number of r -edges in μ . Clearly, if μ is a reduced nontrivial f -path in G_r , then $\mathcal{N}(\hat{f}(\mu)) \leq \mathcal{N}(\mu)$.

Definition 13.1. Let H_r be a polynomial stratum. A vertex $\mu \in D_f$ is called \mathcal{E}_r -perfect if $\mu \equiv E_1 b_1 E_2 \cdots E_k b_k$, where $k \geq 1$, E_1, \dots, E_k are r -edges, b_1, \dots, b_k are paths which lie in G_{r-1} or are trivial, and $\mathcal{N}(\mu') = \mathcal{N}(\mu)$ for every vertex μ' in the μ -subgraph.

Proposition 13.2. Let H_r be a polynomial stratum. Let $\mu \equiv E_1 b_1 \cdots E_k b_k$ be a reduced f -path in G_r , where $k \geq 1$, E_1, \dots, E_k are r -edges, and b_1, \dots, b_k are paths which lie in G_{r-1} or are trivial. For $1 \leq j \leq k$ and $i \geq 1$, we set

$$\begin{aligned}\mu_{0,j} &\equiv [E_j b_j \cdots E_k b_k f(E_1 b_1 \cdots E_{j-1} b_{j-1})], \\ \mu_{i,j} &\equiv [\bar{c}_{i,E_j} f^i(\mu_{0,j}) f(c_{i,E_j})].\end{aligned}$$

Then the following statements are satisfied.

- (1) μ is \mathcal{E}_r -perfect if and only if $\mathcal{N}(\hat{f}(\mu)) = \mathcal{N}(\mu)$.
- (2) One can find a vertex in the μ -subgraph which is \mathcal{E}_r -perfect or lies in G_{r-1} (regarded as an f -path), or is dead.
- (3) If μ is \mathcal{E}_r -perfect, then $\mu_{i,j} \equiv \hat{f}^{m_{i,j}}(\mu)$ for some computable $m_{i,j}$ satisfying $m_{0,1} = 0$, $m_{i,j} < m_{i,j+1}$, and $m_{i,k} < m_{i+1,1}$. Moreover, all \mathcal{E}_r -perfect vertices of the μ -subgraph are $\mu_{i,j}$, $1 \leq j \leq k$, $i \geq 0$.

Proof. (1) If $k = 1$, then μ is \mathcal{E}_r -perfect. Suppose that $k \geq 2$. Then (1) follows by induction from the next claim as \mathcal{N} is a decreasing function.

Claim. The condition (a) below implies the condition (b).

- (a) $\mathcal{N}(\mu) = \mathcal{N}(\hat{f}(\mu)) = k$.
- (b) $\mathcal{N}(\mu') = \mathcal{N}(\hat{f}(\mu')) = k$, where $\mu' := \hat{f}^{1+l(b_1)}(\mu)$.

Proof of Claim. We have $\hat{f}(\mu) \equiv [b_1 E_2 b_2 \cdots E_k \cdot b_k c_{1,E_1} \cdot \sigma(E_1) \cdot \bar{c}_{1,\bar{E}_1}]$. If (a) is valid, then b_1 is an initial subpath of $\hat{f}(\mu)$ and we have

$$\mu' \equiv [E_2 b_2 \cdots E_k \cdot b_k c_{1,E_1} \cdot \sigma(E_1) \cdot \bar{c}_{1,\bar{E}_1} f(b_1)],$$

hence $\mathcal{N}(\mu') = \mathcal{N}(\hat{f}(\mu)) = k$. Then

$$\hat{f}(\mu') = [b_2 E_3 \cdots E_k \cdot b_k c_{1,E_1} \cdot \sigma(E_1) \cdot \bar{c}_{1,\bar{E}_1} f(b_1) c_{1,E_2} \cdot \sigma(E_2) \cdot \bar{c}_{1,\bar{E}_2}].$$

Suppose that (b) is not valid, that is, $\mathcal{N}(\hat{f}(\mu')) < k$. Then $[\sigma(E_1) \cdot \bar{c}_{1,\bar{E}_1} f(b_1) c_{1,E_2} \cdot \sigma(E_2)]$ is trivial. This is possible only if $E_2 = \bar{E}_1$ (hence b_1 is a loop) and $[\bar{c}_{1,\bar{E}_1} f(b_1) c_{1,\bar{E}_1}]$ is trivial. The latter is equivalent that b_1 is trivial. But then $[E_1 b_1 E_2]$ is trivial and μ is not reduced, a contradiction. \square

Statement (2) follows from (1).

(3) First we show that the equation $\mu_{0,j} \equiv \hat{f}^{m_{0,j}}(\mu)$ is satisfied for $m_{0,j} := l(E_1 b_1 \cdots E_{j-1} b_{j-1})$, $2 \leq j \leq k$. Since μ is \mathcal{E}_r -perfect, we have

$$\hat{f}(\mu) \equiv [b_1 E_2 \cdots E_k b_k f(E_1)] \equiv b_1 E_2 \cdots E_k [b_k c_{E_1}] \cdot \sigma(E_1) \cdot \bar{c}_{\overline{E_1}}.$$

Then

$$\hat{f}^{m_{0,2}}(\mu) \equiv [E_2 \cdots E_k b_k f(E_1) f(b_1)] \equiv E_2 \cdots E_k [b_k c_{E_1}] \cdot \sigma(E_1) \cdot [\bar{c}_{\overline{E_1}} f(b_1)].$$

Hence, $\mu_{0,2} \equiv \hat{f}^{m_{0,2}}(\mu)$ as desired. Inducting on j for $2 \leq j \leq k$, we deduce

$$\begin{aligned} \hat{f}^{m_{0,j}}(\mu) &\equiv [E_j \cdots E_k b_k f(E_1 b_1 \cdots E_{j-1} b_{j-1})] \\ &\equiv E_j \cdots E_k [b_k c_{E_1}] \cdot \sigma(E_1) \cdot [\bar{c}_{\overline{E_1}} f(b_1) c_{E_2}] \cdots \sigma(E_{j-1}) [\bar{c}_{\overline{E_{j-1}}} f(b_{j-1})]. \end{aligned}$$

Hence, $\mu_{0,j} \equiv \hat{f}^{m_{0,j}}(\mu)$ as desired. Now we set $m_{1,1} := m_{0,k} + l(E_k [b_k c_{E_1}])$. Using the above equation for $\hat{f}^{m_{0,k}}(\mu)$, we obtain

$$\begin{aligned} \hat{f}^{m_{1,1}}(\mu) &\equiv \sigma(E_1) \cdot [\bar{c}_{\overline{E_1}} f(b_1) c_{E_2}] \cdots \sigma(E_{k-1}) \\ &\quad \cdot [\bar{c}_{\overline{E_{k-1}}} f(b_{k-1}) c_{E_k}] \cdot \sigma(E_k) \cdot [\bar{c}_{\overline{E_k}} f(b_k) f(c_{E_1})]. \end{aligned}$$

Thus, $\hat{f}^{m_{1,1}}(\mu) \equiv [\bar{c}_{\overline{E_1}} f(\mu) f(c_{E_1})] \equiv \mu_{1,1}$ as desired. The path $\mu_{1,1}$ begins with an r -edge and lies in the μ -subgraph, hence it is \mathcal{E}_r -perfect. Considering $\mu_{1,1}$ instead of $\mu_{0,1}$ and using the above schema, one can define $m_{1,j}$ for $2 \leq j \leq k$ and $m_{2,1}$. Recursively, one can define all $m_{i,j}$.

The paths $\mu_{i,j}$ are exactly those in the sequence $\mu, \hat{f}(\mu), \hat{f}^2(\mu), \dots$, which begin with r -edges. Therefore only these paths in this sequence are \mathcal{E}_r -perfect.

Proposition 13.3. *Let H_r be a polynomial stratum. For every two \mathcal{E}_r -perfect vertices μ, τ in D_f , one can decide whether τ lies in the μ -subgraph.*

Proof. By Proposition 13.2(3), τ lies in the μ -subgraph if and only if $\tau \equiv \mu_{i,j}$ for some $i \geq 0$ and $1 \leq j \leq k$, where $k = \mathcal{N}(\mu)$.

Let m be the number of edges in H_r including the inverses. Since the filtration for f is maximal, we have $\sigma^m = \text{id}$. Then, for each r -edge E we have

$$f^m(E) \equiv c_{m,E} \cdot E \cdot \bar{c}_{m,\overline{E}}. \quad (13.1)$$

Since f , restricted to any edge, is a piecewise-linear map, we can find a subdivision $E = E' E''$ such that $f^m(E') \equiv c_{m,E} E'$ and $f^m(E'') \equiv E'' \bar{c}_{m,\overline{E}}$. This implies

$$c_{m,E} = f^m(E') \bar{E}''. \quad (13.2)$$

Claim. For any integers $a, b, s, t \geq 0$ and for each r -edge E the following is satisfied:

- (1) $c_{a+b,E} = f^b(c_{a,E}) c_{b,\sigma^a(E)}$.
- (2) $c_{ms,E} = f^{ms}(E') \bar{E}'$.
- (3) $c_{ms+t,E} = f^{ms+t}(E') f^t(\bar{E}') c_{t,E}$.

Proof of Claim. Claim (1) follows from the definition of $c_{i,E}$. From (1) and using $\sigma^m = \text{id}$, we have

$$c_{ms,E} = f^{m(s-1)}(c_{m,E}) \cdots f^m(c_{m,E})c_{m,E}.$$

This and Eq. (13.2) imply (2). Claim (3) follows from (1) and (2). \square

Using Claim (3), we deduce

$$\begin{aligned} \mu_{ms+t,j} &\equiv [\bar{c}_{ms+t,E_j} f^{ms+t}(\mu_{0,j}) f(c_{ms+t,E_j})] \\ &\equiv [\bar{c}_{t,E_j} f^t(E'_j) \cdot f^{ms+t}(\bar{E}'_j \mu_{0,j} f(E'_j)) \cdot f^{t+1}(\bar{E}'_j) f(c_{t,E_j})]. \end{aligned} \quad (13.3)$$

Thus, $\tau \equiv \mu_{i,j}$ for some $i \geq 0$ and $1 \leq j \leq k$ if and only if there exist $s \geq 0$, $0 \leq t < m$, and $1 \leq j \leq k$ such that

$$[f^t(\bar{E}'_j) c_{t,E_j} \tau f(\bar{c}_{t,E_j}) f^{t+1}(E'_j)] = [f^{ms}(f^t(\bar{E}'_j \mu_{0,j} f(E'_j)))].$$

For fixed t, j , and using Corollary 12.2 for f^m , one can decide whether there exists $s \geq 0$ satisfying the above equation. Hence, one can decide whether there exist i, j with $\tau \equiv \mu_{i,j}$.

14. The Finiteness and Membership Problems

We continue to work with the PL-relative train track $f : \Gamma \rightarrow \Gamma$ satisfying (RTT-iv). We prove Propositions 14.4 and 14.7 which solve the Finiteness and the Membership problems for such f .


Definition 14.1. A path p in D_f is called \triangleright -directed if either p is a vertex in D_f or $p \equiv e_1 e_2 \cdots e_n$, where e_1, \dots, e_n are edges in D_f and the preferable direction at $\alpha(e_i)$ is the direction of e_i for each $i = 1, 2, \dots, n$.

Recall the definitions of k -transversal paths and k -spheres $\mathcal{S}_k(\mu)$ from Definition 6.1.

Lemma 14.2. *Let μ be a vertex in D_f and let p be a k -transversal path to the μ -subgraph. If p is not \triangleright -directed, then there exists a terminal subpath p_1 of p such that p_1 is \triangleright -directed and $\alpha(p_1)$ is a repelling vertex of D_f .*

Proof. Let $p \equiv e_1 e_2 \cdots e_n$. If e_n is a repelling edge, then we set $p_1 := \omega(p)$. If not, then the preferable direction at $\alpha(e_n)$ is carried by e_n (see Fig. 11). If e_{n-1} is a repelling edge, then we set $p_1 := e_n$. If not, then the preferable direction at $\alpha(e_{n-1})$ is carried by e_{n-1} . We continue in this way. Since p is not \triangleright -directed, we will find a repelling edge e_i in p such that the subpath $e_{i+1} \cdots e_n$ is \triangleright -directed. \square

Remark 14.3. Suppose that the μ -subgraph R is infinite. Let p_1 and p_2 be \triangleright -directed paths which are k_1 -transversal and k_2 -transversal to R . If $k_1 \neq k_2$, then $\alpha(p_1) \neq \alpha(p_2)$.

Fig. 11. An example with $p = e_1 e_2 e_3 e_4$.

Proposition 14.4. *Given a vertex μ in D_f , one can decide whether the μ -subgraph is finite or not.*

Proof. Let r be the minimal number such that the f -path μ lies in G_r .

First suppose that H_r is an exponential stratum. Observe that if $\sigma \in \mathcal{S}_s(\mu)$ is an alive vertex, then $\hat{f}(\sigma) \in \mathcal{S}_s(\mu)$ or $\hat{f}(\sigma) \in \mathcal{S}_{s+1}(\mu)$. Moreover, if s satisfies the conditions of Theorem 6.2, then $[f(\sigma)] \in \mathcal{S}_t(\mu)$ for some computable $t > s$, and one can construct a t -transversal path from $[f(\sigma)]$ to the μ -subgraph.

This and Theorem 11.6 imply that we can either

- (a) prove that the μ -subgraph is finite, or
- (b) find $k \geq 1$ and a vertex $\mu' \in \mathcal{S}_k(\mu)$ with one of the following properties:
 - (1) the f -path μ' lies in G_{r-1} ,
 - (2) μ' is r -perfect,
 - (3) μ' is \mathcal{A}_r -perfect, or
 - (4) $l(\mu') \leq \max\{2C_\star, R_\star + \|f\|\}$.

Moreover, we can construct a k -transversal path p from μ' to the μ -subgraph.

It suffices to consider (b). If we have Case (4) or if the path p is not \triangleright -directed, then we restart with $\mu := \hat{f}^{(k+1)}(\mu)$ if it exists.

Let K be the number of f -paths τ in Γ with $l(\tau) \leq \max\{2C_\star, R_\star + \|f\|\}$ plus the number of repelling vertices in D_f . If the number of restarts exceeds K or the new μ does not exist, then the original μ -subgraph is finite (see Lemma 14.2 and Remark 14.3). Thus, we may assume that μ' satisfies one of (1)–(3) and the path p is \triangleright -directed. Hence the μ -subgraph and the μ' -subgraph intersect, thus they are either both finite or both infinite. Redenoting, we may assume that $\mu = \mu'$ satisfies one of (1)–(3).

In Case (1) we apply induction. In Case (2) the μ' -subgraph is infinite by Proposition 10.2. Consider Case (3). By Proposition 10.4(2), there exist natural numbers $m_{1,1} < m_{2,1} < m_{3,1} < \dots$ such that $\hat{f}^{m_{i,1}}(\mu') \equiv [f^i(\mu')]$, $i > 0$. Hence, the μ' -subgraph is finite if and only if there exist $0 < i < j$ such that $[f^i(\mu')] = [f^j(\mu')]$. This problem is decidable by Corollary 12.3.

Now suppose that H_r is a polynomial stratum. By Proposition 13.2(2), we can find a vertex μ' in the μ -subgraph with one of the following properties:

- (1) μ' is dead,
- (2) the f -path μ' lies in G_{r-1} , or
- (3) μ' is \mathcal{E}_r -perfect.

In Case (1) the μ -subgraph is finite. In Case (2) we apply induction. Consider Case (3). Let m be the number of edges in H_r including the inverses. Let $\mu' \equiv E_1 b_1 \cdots E_k b_k$, where $k \geq 1$, E_1, \dots, E_k are r -edges, and b_1, \dots, b_k are paths which lie in G_{r-1} or are trivial. By Proposition 13.2(3), the μ' -subgraph is finite if and only if there exist $0 \leq s_1 < s_2$ such that $\mu'_{m s_1, 1} \equiv \mu'_{m s_2, 1}$. (Recall that $\mu'_{0, 1} \equiv \mu'$.) By the formula (13.3), this is equivalent to

$$[f^{m s_1}(\overline{E_1} \mu' f(E_1))] = [f^{m s_2}(\overline{E_1} \mu' f(E_1))].$$

The problem of the existence of such s_1, s_2 is decidable by Corollary 12.3.

If H_r is a zero stratum, then following along the μ -subgraph at most $l(\mu)$ steps, we can find a vertex μ' in the μ -subgraph such that the f -path μ' lies in G_{r-1} or is trivial. Then we apply induction. \square

Proposition 14.5. *For every two vertices μ, τ in D_f , where μ is r -perfect, one can decide whether τ lies in the μ -subgraph.*

Proof. We may assume that the f -path τ lies in G_r (otherwise τ does not lie in the μ -subgraph). Let $\mu = \mu_0, \mu_1, \dots$, be consecutive vertices of the μ -subgraph. We apply Proposition 10.2(3) to μ to compute the minimal i such that $L_r(\mu_i) > L_r(\tau)$. Then τ lies in the μ -subgraph if and only if τ coincides with one of the vertices $\mu_0, \mu_1, \dots, \mu_{i-1}$. \square

Proposition 14.6. *For every two vertices μ, τ in D_f , where μ is \mathcal{A}_r -perfect, one can decide whether τ lies in the μ -subgraph.*

Proof. Due to Proposition 14.4, we may assume that the μ -subgraph is infinite. Let $\mu \equiv A_1 b_1 \cdots A_k b_k$ be the r -decomposition of μ . We use the following notation from Proposition 10.4:

$$\begin{aligned} \mu_{0,j} &\equiv [A_j b_j \cdots A_k b_k f(A_1 b_1 \cdots A_{j-1} b_{j-1})], \\ \mu_{i,j} &\equiv [f^i(\mu_{0,j})], \end{aligned} \tag{14.1}$$

where $1 \leq j \leq k$ and $i \geq 1$. By Proposition 10.4(4), for every vertex σ in the μ -subgraph, at least one of the paths $\sigma, \hat{f}(\sigma), \dots, \hat{f}^{l(\sigma)}(\sigma)$ coincides with $\mu_{i,j}$ for some i, j .

Thus, we first decide, whether one of the paths $\tau, \hat{f}(\tau), \dots, \hat{f}^{l(\tau)}(\tau)$ coincides with $\mu_{i,j}$ for some i, j . In view of (14.1), this can be done with the help of Corollary 12.2. If the answer is negative, then τ does not lie in the μ -subgraph. If it

is positive, then we can find t, i, j such that $\hat{f}^t(\tau) \equiv \mu_{i,j}$. Recall that by Proposition 10.4(2), $\mu_{i,j} \equiv \hat{f}^{m_{i,j}}(\mu)$ for computable $m_{i,j}$. Then τ lies in the μ -subgraph if and only if $m_{i,j} \geq t$ and $\tau \equiv \hat{f}^{m_{i,j}-t}(\mu)$. \square

Proposition 14.7. *Given two vertices μ, τ in D_f , one can decide whether τ lies in the μ -subgraph.*

Proof. By Proposition 14.4, we can decide whether the μ -subgraph and the τ -subgraph are finite or not. If the μ -subgraph is finite, then we can compute all its vertices and verify whether τ is one of them. Suppose that the μ -subgraph is infinite. If the τ -subgraph is finite, then the vertex τ cannot lie in the μ -subgraph. So, we may assume that the τ -subgraph is also infinite. Let r be the minimal number such that the f -path μ lies in G_r . We will induct on r .

First suppose that H_r is an exponential stratum. Then, as in the proof of Proposition 14.4, we can find $k \geq 1$, a vertex $\mu' \in \mathcal{S}_k(\mu)$, and a \triangleright -directed path p from μ' to $\mu_k := \hat{f}^k(\mu)$ such that one of the following properties is satisfied:

- (1) the f -path μ' lies in G_{r-1} ,
- (2) μ' is r -perfect, or
- (3) μ' is \mathcal{A}_r -perfect.

First we check whether τ belongs to the segment of the μ -subgraph from μ to μ_k . If yes, we are done. If not, we check whether τ belongs to the segment of the μ' -subgraph from μ' to μ_k . If yes, then τ does not belong to the μ -subgraph. If not, we replace μ by μ' and consider the above cases. In Case (1) we proceed by induction, in Case (2) by Proposition 14.5, and in Case (3) by Proposition 14.6.

Now suppose that H_r is a polynomial stratum. Then, by Proposition 13.2(2), we can find a vertex μ' in the μ -subgraph with one of the following properties:

- (1) the f -path μ' lies in G_{r-1} ,
- (2) μ' is \mathcal{E}_r -perfect.

In Case (1) we proceed by induction. Consider Case (2). We may assume that the f -path τ lies in G_r , otherwise τ does not lie in the μ -subgraph. By Proposition 13.2, we can find $k \geq 0$ such that either $\hat{f}^k(\tau)$ lies in G_{r-1} , or $\hat{f}^k(\tau)$ is \mathcal{E}_r -perfect. If $\hat{f}^k(\tau)$ lies in G_{r-1} , then τ does not lie in the μ -subgraph. Suppose that $\hat{f}^k(\tau)$ is \mathcal{E}_r -perfect. By Proposition 13.3, we can decide whether $\hat{f}^k(\tau)$ lies in the μ' -subgraph, and hence in the μ -subgraph (these subgraphs differ by a finite segment). If $\hat{f}^k(\tau)$ does not lie in the μ -subgraph, then τ does not lie in the μ -subgraph. If $\hat{f}^k(\tau)$ lies in the μ -subgraph, say $\hat{f}^k(\tau) = \hat{f}^t(\mu)$, then τ lies in the μ -subgraph if and only if $t \geq k$ and $\tau = \hat{f}^{t-k}(\mu)$.

Finally, if H_r is a zero stratum, we follow along the μ -subgraph at most $l(\mu)$ steps until we arrive at a vertex $\mu' \in D_f$ which, regarded as an f -path, lies in G_{r-1} . Then we apply induction. \square

15. The Main Algorithm

Our algorithm for finding a basis of $\text{Fix}(\varphi)$ is the following:

- (1) Represent the automorphism $\varphi : F \rightarrow F$ by a PL-relative train track $f : (\Gamma, v) \rightarrow (\Gamma, v)$ (see Theorem 4.4). Acting as in the proof of this theorem we may additionally assume that $\mathbf{1}_v$ is a repelling vertex in D_f .
- (2) Subdivide f as in Corollary 8.13 to ensure that f satisfies (RTT-iv).
- (3) Construct $\text{CoRe}(C_f)$. (A construction modulo the Finiteness and the Membership problems is explained in Sec. 5. Solutions to these problems in the case where f is a PL-relative train track satisfying (RTT-iv) are given in Sec. 14.)
- (4) Find the component Δ of $\text{CoRe}(C_f)$ containing $\mathbf{1}_v$ and compute $\pi_1(\Delta, \mathbf{1}_v)$. By Lemma 5.1, we have $\overline{\text{Fix}}(f) \cong \pi_1(D_f(\mathbf{1}_v), \mathbf{1}_v) \cong \pi_1(\Delta, \mathbf{1}_v)$.
- (5) Compute a basis of $\text{Fix}(\varphi)$ using (1) and (4).

References

- [1] M. Bestvina, M. Feighn and M. Handel, The Tits alternative for $\text{Out}(F_n)$ I: Dynamics of exponentially-growing automorphisms, *Ann. of Math.* **151**(2) (2000) 517–623.
- [2] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups, *Ann. of Math.* **135**(1) (1992) 1–53.
- [3] O. Bogopolski, Classification of automorphisms of the free group of rank 2 by ranks of fixed-point subgroups, *J. Group Theory* **3**(3) (2000) 339–351.
- [4] P. Brinkmann, Detecting automorphic orbits in free groups, *J. Algebra* **324**(5) (2010) 1083–1097.
- [5] M. M. Cohen and M. Lustig, On the dynamics and the fixed subgroup of a free group automorphism, *Invent. Math.* **96**(3) (1989) 613–638.
- [6] D. Cooper, Automorphisms of free groups have finitely generated fixed point sets, *J. Algebra* **111**(2) (1987) 453–456.
- [7] R. Z. Goldstein and E. C. Turner, Fixed subgroups of homomorphisms of free groups, *Bull. London Math. Soc.* **18**(5) (1986) 468–470.
- [8] O. S. Maslakova, Fixed point subgroup of an automorphism of a free group, *Algebra Logic* **42**(4) (2003) 237–265.
- [9] O. S. Maslakova, Fixed point subgroup of an automorphism of a free group, Ph.D Thesis (2004) (in Russian).
- [10] Open problems in combinatorial and geometric group theory, <http://www.grouptheory.info/>.
- [11] E. C. Turner, Finding indivisible Nielsen paths for a train tracks map, in *Combinatorial and geometric group theory*, London Mathematical Society Lectures Note Series, Vol. 204 (Cambridge University Press, 1994), pp. 300–313.