THE SET OF TRAIN TRACK REPRESENTATIVES OF AN IRREDUCIBLE FREE GROUP AUTOMORPHISM IS CONTRACTIBLE

Jerome Los and Martin Lustig

preliminary version 20. 12. 2004

Introduction and statement of results

One of the crucial tools of Thurston's celebrated analysis of surface homeomorphisms are so called *train tracks*, which are branched 1-dimensional manifolds (see also [17]) embedded in the surface, with certain (mild) restrictions on the complementary components. There have been several attempts to carry over this important concept to automorphisms of free groups F_n (see [10], [12]), most prominently the one by Bestvina-Handel [2]. They introduced relative train track maps and used them, among other, to prove Scott's Conjecture that $rk\{w \in F_n \mid \alpha(w) = w\}) \leq n$ for all $\alpha \in Aut(F_n)$. Train track maps have become a fundamental tool in the theory of free group automorphisms ever since.

Special attention has been given to Bestvina-Handel's *irreducible* automorphisms of F_n, which can be represented by an *absolute train track map*, i.e. a homotopy equivalence $f : \Gamma \to \Gamma$ of a graph Γ with fundamental group F_n which maps vertices to vertices and for which every edge e of Γ is *legal*: All positive powers f^k of f map e in a locally injective way.

Absolute train track representatives $f : \Gamma \to \Gamma$ of non-periodic irreducible automorphisms have irreducible transition matrices $M(f) = (m_{e,e'})_{e,e' \in Edges(\Gamma)}$, where $m_{e,e'}$ denotes the number of times $f(e)$ passes over e' or its inverse. The (up to rescaling) uniquely determined positive eigenvector $v = (v_e)_{e \in Edges(\Gamma)}$ with Perron-Frobenius eigenvalue $\lambda > 1$ defines edge lengths $L(e) = v_e$ which have the property that every legal path γ in Γ is mapped by f to a path of length $L(f(\gamma)) = \lambda L(\gamma)$.

We use this property here to generalize the notion of train track maps to maps $f : \Gamma \to \Gamma$ of metric graphs which do not nececessarily map vertices to vertices. The resulting set of such *efficient representatives* of an irreducible non-periodic $\alpha \in Aut(F_n)$, provided with a canonical topoloty, is denoted by $\mathcal{E}(\alpha)$. The main result of this paper is:

Theorem I. *For any irreducible non-periodic automorphism α of F_n the space of efficient representatives $\mathcal{E}(\alpha)$ of α is contractible.*

There is a canonical map from $\mathcal{E}(\alpha)$ to Culler-Vogtmann's Outer space CV_n (see [3]), defined in section 1, which is injective for most non-periodic irreducible automorphisms α (see section 2).

The main tool of our investigations is the (projectively) α-invariant $\mathbb{R}$-tree T with isometric F_n-action, on which α acts by an expanding homothety, as well as F_n-equivariant maps $i : \tilde{\Gamma} \to T$ which map every edge of $\tilde{\Gamma}$ isometrically to a segment in T. By introducing *canonical folding* we define a semi-flow Φ on $\mathcal{E}(\alpha)$. The space $\mathcal{E}(\alpha)$ decomposes in a natural way into finitely many strata, called *blow-up classes*.

Theorem II. *For any irreducible non-periodic automorphism α of F_n the semi-flow Φ defines finitely many α-invariant flow lines in the space of efficient representatives $\mathcal{E}(\alpha)$ of α. More precisely, there is exactly one such line in each blow-up class of $\mathcal{E}(\alpha)$.*

In fact, we show in section 3 that among the finitely many α-invariant flow lines in $\mathcal{E}(\alpha)$ there is one prefered *principal axis* of α. Thus, the canonical folding semi-flow Φ has a certain resemblance to Teichmüller flow on Teichmüller space. We hope that further investigations of the semi-flow Φ prove to be useful for a better understanding of the geometry of Outer space and its boundary.

Our investigations have been inspired by a "predecessor" of this paper: In [11] the set of all train track representatives of an irreducible $\alpha \in Aut(F_n)$ has been studied and interpreted as the 0-skeleton of a certain 1-dimensional complex, also defined by means of folding operations, which is shown there to be connected. An alternative proof of a slightly modified version of this result has been given in [8].

We know of several papers, all of them unpublished, which have tried (much) related approaches: Skora [14] (compare also [7]) and White [16] have both studied Outer space using foldings, but their "canonical folds" are not the same as ours. Fehrenbach [4] has introduced "super-efficient" representatives of geometric automorphisms of F_n, which coincide partially with the ones exhibited here, see section 3.

Acknowledgments. *The second author would like to thank G. Levitt for the many conversations which have influenced this paper. He also wants to thank Univerité de Nice - Sophia Antipolis for being invited for a research stay in March 1999, during which an essential part of the work presented here was done. Furthermore, both authors would like to thank the MPI in Bonn and the CRM in Barcelona for their generous support during research stays in the fall of 2000 and of 2004.*

1. The general setting

In their seminal paper [3] M. Culler and K. Vogtmann defined, in striking analogy to Teichmüller space for surfaces and their homeomorphisms, a space called *Outer space* for free groups and their automorphisms. We denote Outer space by CV_n, where $n \geq 2$ is the rank of the finitely generated free group F_n. A point in CV_n is given by a *marked metric graph* Γ, by which we mean a finite connected graph Γ without edges of valence 1, where every edge e has been given a length $L(e) > 0$, and which is provided with a *marking*, i.e. an isomorphism $\theta : F_n \to \pi_1\Gamma$. Two labeled metric graphs Γ and Γ' describe the same point $[\Gamma] = [\Gamma']$ in CV_n if and only if there exists a homothety $\Gamma \to \Gamma'$ that commutes on π_1 with the marking isomorphisms.

Below we also consider the *unprojectivized Outer space* cv_n, which is defined precisely as CV_n, except that the word "homothety" in the last paragraph has to be replaced by "isometry". Projectivization leads to a natural map $cv_n \to CV_n$, $\Gamma \mapsto [\Gamma]$, which is equivariant with respect to the natural action of outer automorphisms α of F_n on marked graphs given by the marking change $\theta \mapsto \theta\alpha$.

By passing from Γ to its universal covering, a metric tree $\tilde{\Gamma}$ on which F_n acts via the marking by (isometric) deck transformations, the space cv_n is seen to embed canonically into the space of F_n-actions on $\mathbb{R}$-trees. In fact, cv_n is precisely the subspace of minimal free simplicial such actions, where *minimal* means that there is no non-empty F_n-invariant proper subtree. The closure $\overline{cv}_n$ of this subspace is given precisely by all minimal very small F_n-actions on $\mathbb{R}$-trees T. For more background and references see [9].

We define the volume $vol(\Gamma)$ of a metric graph Γ to be the sum of its edge lengths $L(e)$. As usual, the edges of Γ are unoriented, but for the sake of notation one puts an orientation on every edge, which however can be reversed at any given time (without explicit warning, but usually accompanied by switching the label e to $\bar{e}$), and this is not considered to be a change of the graph Γ.

For the rest of this section we fix a tree T in $\overline{cv}_n$, We consider the set $\mathcal{I}_0(T)$ of all pairs $(\tilde{\Gamma}, i)$, where $\tilde{\Gamma}$ denotes as before the universal covering of a marked metric graph $\Gamma \in cv_n$, and i is an F_n-equivariant map $i : \tilde{\Gamma} \to T$ that maps every edge $e \subset \tilde{\Gamma}$ isometrically onto a segment $i(e) \subset T$. We call such a map i *edge-isometric.* Note that the minimality of T and the F_n-equivarience of i imply that the map i is surjective.

Clearly the edge-isometric map i as well as the length of the edges of Γ is well defined once the images of the vertices of $\tilde{\Gamma}$ are specified. Hence we obtain a canonical topology on $\mathcal{I}_0(T)$, defined by continuously varying the i-images of the vertices of $\tilde{\Gamma}$. For any non-metric marked graph Γ^{top}, with a fixed cellular structure of vertices and edges, this gives a subspace $\mathcal{C}_0(\Gamma^{top})$ of points $(\tilde{\Gamma}, i)$ in $\mathcal{I}_0(T)$ which is canonically homeomorphic to $T^{Vertices(\Gamma^{top})}$.

Note that there is a canonical map from $\mathcal{I}_0(T)$ to the space cv_n and hence to CV_n. This map, obtained by "forgetting i", is clearly continuous with respect to the topology of $\mathcal{I}_0(T)$.

For any vertex P of Γ, a *turn* (e, e') at P is given by edges e and e' in Γ which (when properly reoriented) have a common initial vertex P. The turn (e, e') is *degenerated* if $e = e'$. This terminology extends to $\tilde{\Gamma}$.

Two edges e and e' raying out of a vertex P of $\tilde{\Gamma}$ define an *illegal turn* (e, e') at P if they have non-trivial inital segments with identical i-images. We say that these initial segments are *folded by* i. "Defining an illegal turn" is an equivalence relation among the edges with common initial vertex P, and an equivalence class is called a *gate* at P. For notational convenience we will sometimes extend the notion of gates to points in the interior of an edge e of $\tilde{\Gamma}$: For such points x there are precisely two gates, each of them represented by a segment of e adjacent to x in one of the two possible directions on e. Obviously this gate structure on $\tilde{\Gamma}$ is F_n-equivariant and hence induces a well defined gate structure on the marked metric quotient graph Γ.

We define $\mathcal{I}(T)$ to be the subspace of $\mathcal{I}_0(T)$ which consists of all points $(\tilde{\Gamma}, i)$ where every vertex of $\tilde{\Gamma}$ has at least two gates. One can show that $\mathcal{I}(T)$ is a strong deformation retract of $\mathcal{I}_0(T)$, but this will not be used here. It turns out that in the case we are most concerned with the map i is almost always determined by $\tilde{\Gamma}$. For notational convenience we will thus refer from now on to the points of $\mathcal{I}(T)$ as "trees" rather than "pairs" and denote them by $\tilde{\Gamma}$.

Remark 1.1. The above defined subspaces $\mathcal{C}_0(\Gamma^{top}) \subset \mathcal{I}_0(T)$ define a "cell structure" on $\mathcal{I}(T)$, with *cells* $\mathcal{C}(\Gamma^{top}) = \mathcal{C}_0(\Gamma^{top}) \cap \mathcal{I}(T)$. Notice that this cell structure is locally finite.

In fact, if one introduces a new vertex of valence 2 by subdividing an edge of Γ^{top}, thus creating a new cellular graph Γ_1^{top}, the subspace $\mathcal{C}_0(\Gamma_1^{top})$ will be larger than $\mathcal{C}_0(\Gamma^{top})$, while the corresponding cells $\mathcal{C}(\Gamma_1^{top})$ and $\mathcal{C}(\Gamma^{top})$ are equal. As we are only interested in the space $\mathcal{I}(T)$ we will admit below at any time refinements of the cell structure of Γ without specifically refering to it.

The following definition is crucial for the whole paper; the reader should note that the distance in question is measured in T and not in $\tilde{\Gamma}$.

Definition 1.2. Let $\tilde{\Gamma}$ be an element of $\mathcal{I}(T)$. For any non-degenerated turn (e, e') at a vertex P of $\tilde{\Gamma}$ we define the *identification length*

$$il(e, e') = sup\{d(i(P), i(Q)) \mid Q \in dir_P(e),\ Q' \in dir_P(e'),\ i(Q) = i(Q')\} ,$$

where the *direction* $dir_P(e)$ denotes the connected component of $T - P$ determined by e.

This definition is F_n-equivariant and hence we get a well defined identification length $il(e, e')$ for every turn (e, e') of Γ. The Definition 1.2 extends directly to any point x of Γ, as we can simply make x into a vertex by subdividing the edge which it contains. For any such non-vertex point $x \in \Gamma$ the identification length of the only non-degenerated turn at x varies continuously with x. We denote by $il(\Gamma)$ the maximal identification length of any non-degenerated turn in Γ and call it the *identification length* of Γ.

For F_n-equivariant maps $i : \tilde{\Gamma} \to T$ (not necessalily edge-isometric) the *backtracking bound* $BBT(i) \geq 0$ has been defined in [5] as the smallest real number $C \geq 0$ such that, for any points $P, Q \in \tilde{\Gamma}$ the i-image of the geodesic segment $[P, Q] \subset \tilde{\Gamma}$ is contained in the C-neighborhood of the geodesic segment $[i(P), i(Q)] \subset T$. It has been shown in [5] that for $T \in \overline{cv}_n$ every map $i : \tilde{\Gamma} \to T$ as above possesses such a (finite) bound $BBT(i) \geq 0$. Note that $BBT(i)$ does not depend on the metric of Γ. Note also that the image of any path $[P, Q] \subset \tilde{\Gamma}$ with $i(P) = i(Q)$ (called an *i-backtracking path*) has diameter at most $2BBT(i)$.

Lemma 1.3. *Let T be an* **R***-tree with a minimal very small action of F_n (i.e. $T \in \overline{cv}_n$). Then any $(\tilde{\Gamma}, i) \in \mathcal{I}(T)$ satisfies:*

1. $0 \leq il(\Gamma) \leq BBT(i) \leq vol(\Gamma)$.
2. *If $il(\Gamma) = 0$ then $\tilde{\Gamma} = T$.*

Proof. Assertion (2) and the first two inequalities of assertion (1) follow straight from the definitions. The third inequality of (1) has been proved at different places in the literature with varying degrres of generality (see [9]). □

It turns out that statement (1) of Lemma 1.3 can be improved to $il(i) = BBT(i)$, but this will not be used in this paper.

Every turn (e, e') which is illegal satisfies $il(e, e') > 0$, but the converse is very much wrong. The proof of the following, however, shows that the illegal turns maximize locally the quantity $il(x)$. (Notice that the global maximality of $il(e, e')$ is not used in the proof.)

Lemma 1.4. *If $il(e, e')$ is maximal, then the turn (e, e') is illegal.*

Proof. If no initial segment of (e, e') is folded by i, then for any $\epsilon \geq 0$ there exist points $Q \in dir(e)$ and $Q' \in dir(e')$ which realize $il(e, e')$ as in Definition 1.2 up to some small $\epsilon' < \epsilon$, as well as a point $S \in e \cup e'$ outside an ϵ'-ball around the common initial vertex P of (e, e') which lies on the geodesic segment $[Q, Q']$ but is not mapped to the same connected component of $T - \{i(P)\}$ as $i(Q) = i(Q')$. But then segment $[Q, Q'] \subset \tilde{\Gamma}$ defines a turn at S with identification length strictly larger than $il(e, e')$, in contradiction to our maximality assumption. □

Consider the family $IT(\Gamma)$ of all illegal turns (e, e') in Γ, as well as a *folding vector* $x = (x_{(e,e')})_{(e,e') \in IT(\Gamma)} \in \mathbb{R}^{IT(\Gamma)}$ of real numbers $0 \leq x_{(e,e')} < \frac{1}{2} length(i(e) \cap i(e'))$ $(\leq il(e, e'))$. Then one can identify isometrically initial segments of e and e' of length $x_{(e,e')}$ to get a quotient graph Γ_x (also metric and with an induced marking) as well as an induced F_n-equivariant edge-isometric map $i_x : \tilde{\Gamma}_x \to T$. Hence we obtain a (multiple) *isometric folding map* $f_x : \tilde{\Gamma} \to \tilde{\Gamma}_x$ which is *T-compatible*, by which we mean that it is F_n-equivariant and that i splits over f_x (via $i = i_x f_x$). Notice that it follows from our above choice of the lengths of the cooridnates of the vector x that the folding map f_x induces a well defined map on turns, which maps legal to legal and illegal to illegal turns.

Suppose now that for Γ_x one repeats the procedure: One chooses a new folding vector x' for Γ_x and obtains a new quotient graph $\Gamma_{x,x'}$ and a T-compatible folding map $f_{x,x'} : \tilde{\Gamma}_x \to \tilde{\Gamma}_{x,x'}$. If this new folding map only folds turns in Γ_x that are images of illegal turns in Γ, then there is a folding vector x'' for Γ such that $\Gamma_{x,x'}$ can alternatively be obtained by the single folding map $f_{x''} = f_{x'} \circ f_x : \tilde{\Gamma} \to \tilde{\Gamma}_{x,x'}$. A slightly incorrect but unambiguous notation which turns out to be useful here is to state $x'' = x + x'$.

Conversely, by the same reasoning one can write the original folding vector x as sum of two non-negative folding vectors $x = x^* + x^{**}$ to obtain f_x as product $f_x = f_{x^{**}} \circ f_{x^*} : \tilde{\Gamma} \to \tilde{\Gamma}_x$. This gives the possibility to interpret folding as a continuous procedure, once one specifies a parametrisation of the folding vectors $[t_0, t_1] \to \mathbb{R}^{IT(\Gamma)}$, $t \mapsto x(t)$ with $x(t_0) = 0$ and $x(t_1) = x$.

Warning. It is quite possible that a folding at some illegal turn (e, e') cannot be extended to the full identification length of (e, e'), as the illegal turn "vanishes" already beforehand ! Worse, if the turn (e, e') is legal but $il(e, e') > 0$, then one cannot even start folding. This however can not happen in the situation where il is maximal, as was proved in Lemma 1.4.

Consider any two elements $\tilde{\Gamma}$ and $\tilde{\Gamma}'$ of $\mathcal{I}(T)$, and assume that there is an F_n-equivariant map $h : \tilde{\Gamma} \to \tilde{\Gamma}'$ such that $i = i'h$ (i.e. h is T-compatible). It follows from [15] that h is a finite composition of (cellular) foldings maps, which must all be isometric and T-compatible. One deduces easily:

Lemma 1.5. *If $h : \tilde{\Gamma} \to \tilde{\Gamma}'$ is T-compatible, then $vol(\Gamma) \geq vol(\Gamma')$. If $vol(\Gamma) = vol(\Gamma')$ then it follows that h is a marking preserving isometry: $\Gamma = \Gamma'$.* □

More important is the following observation:

Lemma 1.6. *If $h : \tilde{\Gamma} \to \tilde{\Gamma}'$ is T-compatible, then $il(\Gamma) \geq il(\Gamma')$.*

Proof. Consider (as in Definition 1.2) a vertex $P' \in \tilde{\Gamma}'$ and a turn (e'_1, e'_2) at P' with $il(e'_1, e'_2) = il(\tilde{\Gamma}')$. Let $Q'_1 \in dir_{P'}(e'_1)$ and $Q'_2 \in dir_{P'}(e'_2)$ be points which realize $il(e'_1, e'_2)$ up to some small $\epsilon \geq 0$. Since the F_n-action on $\tilde{\Gamma}$ and $\tilde{\Gamma}'$ is minimal (by assumption there are no valence 1 vertices !) it follows that h is surjective. Let Q_1 and Q_2 be arbitrary h-preimages of Q'_1 and Q'_2, and consider the geodesic segment $[Q_1, Q_2]$. Its image $h([Q_1, Q_2])$ contains the geodesic $[Q'_1, Q'_2]$, and hence there exists a point $P \in [Q_1, Q_2]$ with $h(P) = P'$. If e_1 and e_2 denote the two edge segments adjacent to P which lie on $[Q_1, Q_2]$ (which are not necessarily mapped to (e'_1, e'_2)), then it follows directly from Definition 1.2 that $il(e_1, e_2) \geq il(e'_1, e'_2)$, which gives $il(\Gamma) \geq il(\Gamma')$. □

We now define, for any $\tilde{\Gamma} \in \mathcal{I}(T)$ with $il(\Gamma) > 0$, a *canonical* folding vector $x = x(t)$ for Γ by $x_{(e,e')} = t$ if $il(e, e') = il(\Gamma)$ and $x_{(e,e')} = 0$ otherwise, where the parameter $t > 0$ is smaller or equal to the smallest edge length of Γ, and also smaller than the difference between $il(\Gamma)$ and the second largest $il(e, e')$ at any vertex of Γ. This last condition implies that in the folding process (thought of as done in a continous fashion for increasing $t' \in [0, t]$) a new illegal turn with maximal il can not arise. At the end of the canonical fold (i.e. at the graph Γ_x reached for the parameter value $t' = t$), we can proceed "continuously" with a subsequent canonical fold, and so on. If in this iterative procedure for the limit parameter value one has reached a point $\tilde{\Gamma}'$ which lies in $\mathcal{I}(T)$ and satisfies $il(\Gamma') > 0$, then there is another canonical folding vector defined at Γ'. By Lemma 1.3 (2) this defines a *canonical folding path* at every point of $\mathcal{I}(T)$ (other than possibly T itself, in the case $T \in cv_n$), and the parameter t gives a well defined *length* of any compact segment of such a path. We say that Γ and $\Gamma_t = \Gamma_{x(t)}$ have *distance* t. We obtain directly from the definition:

Lemma 1.7. *Moving a point along its canonical folding path with "constant speed" is additive:* $(\tilde{\Gamma}_t)_{t'} = \tilde{\Gamma}_{t+t'}$ *for any* $\tilde{\Gamma}$ *in* $\mathcal{I}(T)$ *and* t, t' *sufficiently small.* □

As the folding maps $f_t : \tilde{\Gamma} \to \tilde{\Gamma}_t$ along a canonical folding path are always T-compatible, Lemma 1.6 gives directly:

Proposition 1.8. *Given* $\tilde{\Gamma}$ *in* $\mathcal{I}(T)$, *if* Γ_t *denotes the graph obtained from* Γ *by moving along the canonical folding path with parameter* $t = dist(\Gamma, \Gamma_t)$, *then the value of* $il(\Gamma_t)$ *decreases strictly monotonically for increasing* $t \geq 0$. *More precisely, one has* $il(\Gamma_t) = il(\Gamma) - t$. *In particular, the canonical folding path starting at* Γ *has length smaller or equal to* $il(\Gamma)$. □

In the above definition of the canonical folding path $(\Gamma_t)_{t \in [0,t')}$ that starts at some point $\tilde{\Gamma} = \tilde{\Gamma}_0 \in \mathcal{I}(T)$ we have always assumed that the folding path stays within $\mathcal{I}(T)$. However, for $t \to t'$ the limit tree of the trees $\tilde{\Gamma}_t$ may well not be any longer free and simplicial, i.e. it may well be a point of $\overline{cv}_n - cv_n$. But in most cases this can not happen if $t' < il(\Gamma)$, as is shown in the following proposition.

Proposition 1.9. *If the* F_n*-action on* T *has trivial arc stibilizers, then for any* $\tilde{\Gamma} \in \mathcal{I}(T)$ *and any parameter value* t *with* $0 < t < il(\Gamma)$ *the canonical folding line starting at* $\tilde{\Gamma}$ *extends to a point* $\tilde{\Gamma}_t \in \mathcal{I}(T)$.

Proof. If the canonical folding line starting at $\tilde{\Gamma}$ does not extend to a point $\tilde{\Gamma}_t \in \mathcal{I}(T)$, we may assume without loss of generality that for all $0 \leq s < t$ one has $\tilde{\Gamma}_s \in \mathcal{I}(T)$, and $T' = lim_{s \to t} \tilde{\Gamma}_s$ is a tree in $\overline{cv}_n - cv_n$. We observe that the maps $i_s : \tilde{\Gamma}_s \to T$ define an F_n-equivariant limit map $T' \to T$, and that the F_n-action on T' must have trivial arc stibilizers, since this was assumed for the action on T. The condition $0 < t < il(\Gamma)$ implies via Lemma 1.3 (1) that the volume of the graphs Γ_s does not tend to 0 for $s \to t$. Hence T' quotients onto a (non-free) simplicial tree $\tilde{\Gamma}'$ with trivial edge stabilizers and $vol(\tilde{\Gamma}'/F_n) = lim_{s \to t} vol(\Gamma_s)$. We can lift any vertex P of $\tilde{\Gamma}'$ with non-trivial stabilizer G_P to a subtree T'_P of T' on which G_P acts minimally with dense orbits. Notice that the quotient graph-of-groups decomposition of F_n given by $\tilde{\Gamma}'/F_n$ has trivial edge groups. Hence the vertex group G_P is a proper free factor of F_n and thus of strictly smaller rank than n. The edges of $\tilde{\Gamma}'$ lift homemorphically to arcs without branch points in T'. A slightly shortened subarc of any such arc lifts similarly to an arc without branch point in any sufficiently close "approximation tree" $\tilde{\Gamma}_s$. Among the subtrees of $\tilde{\Gamma}_s$, which are complementary to these arcs without branch points, there is one fixed by G_P. This G_P-invariant subtree contains a G_P-invariant minimal subtree

$\tilde{\Gamma}_s^P$ with quotient graph Γ_s^P and surjective edge-isometric G_P-equivariant map $i_P : \tilde{\Gamma}_s^P \to T'_P$. By induction over the rank of G_P we can assume that $il(\Gamma_P)$ is small, But this contradicts the definition of the canonical folding procedure at Γ_s, since Γ_s^P is a subgraph of Γ_s, and hence, by definition of the identification length, the identification length of any turn in Γ_s^P measured in $\tilde{\Gamma}_s^P$ can not be smaller than measured in $\tilde{\Gamma}_s$. □

Remark 1.10. (a) In the last proof it is not obvious whether or not the limit tree T' of any canonical folding line is always equal to the given tree T.

(b) It seems that the hypothesis "trivial arc stabilizers" in the last Proposition is indeed necessary. In fact, it seems that independently of the choice of T the limit tree T' as above has always trivial arc stabilizers.

The following slightly technical result which will be of great use in the next section.

Proposition 1.11. *Let $h : \tilde{\Gamma} \to \tilde{\Gamma}'$ be an equivariant map with $i = i'h$, and let Γ_t and $\Gamma'_{t'}$ be obtained from Γ and Γ' by folding of length $t \geq 0$ and $t' \geq 0$ along the canonical folding paths. If $il(\Gamma) - t \geq il(\Gamma') - t'$ then h induces an F_n-equivariant T-compatible map $h_{t,t'} : \tilde{\Gamma}_t \to \tilde{\Gamma}'_{t'}$.*

Proof. Notice first that, through replacing Γ' by $\Gamma_{t'}$ and h by $f'_{t'}h$, it suffices to prove the case $t' = 0$. Let $\hat{t} \geq 0$ be the supremum of all $t \geq 0$ such that a map $h_t (= h_{t,0})$ as claimed does exist. From the continuity of the canonical folding process it follows that this supremum is actually a maximum, i.e. an equivariant T-compatible map $h_{\hat{t}} : \tilde{\Gamma}_{\hat{t}} \to \tilde{\Gamma}'$ exists.

Hence for this maximal $\hat{t}$ the induced map $h_{\hat{t}} : \tilde{\Gamma}_{\hat{t}} \to \tilde{\Gamma}'$ does not fold some illegal turn (e, e') in $\tilde{\Gamma}_{\hat{t}}$ with $il(e, e') = il(\Gamma_{\hat{t}})$. Hence we get $il(h_{\hat{t}}(e), h_{\hat{t}}(e')) = il(e, e')$, and thus $il(\Gamma') \geq il(\Gamma_{\hat{t}})$. By Lemma 1.6 applied to the T-compatible map $h_{\hat{t}}$, this implies the equality $il(\Gamma') = il(\Gamma_{\hat{t}})$. But this implies $il(\Gamma_{\hat{t}}) = il(\Gamma) - \hat{t}$ by Proposition 1.8. Thus the maximality in the definition of $\hat{t}$ implies that for any $t \geq 0$ with $il(\Gamma) - t \geq il(\Gamma')$ the map h_t does exist. □

Moving every point of $\mathcal{I}(T)$ along its canonical folding path by a small distance $t > 0$ does not give a flow in the classical sense: As $\mathcal{I}(T)$ is not a manifold it will in general not define a local homeomorphisms. However, we want to prove a property which is the best possible analogue of what in the manifold case corresponds to "local homeomorphism", namely that for any t the map $(\tilde{\Gamma}, t) \to \tilde{\Gamma}_t$, which moves $\tilde{\Gamma}$ by a distance t along its canonical folding path, is continuous. This shows that the canonical folding paths define a *semi-flow* on $\mathcal{I}(T)$ which we call Φ.

For this purpose we first need to study more closely the topology of $\mathcal{I}(T)$. For any tree $\tilde{\Gamma}$ in $\mathcal{I}(T)$ and for some $0 < \epsilon < \frac{1}{2}\, min\{L(e) \mid e \in Edges(\Gamma)\}$ we consider the set $\mathcal{N}_\epsilon(\tilde{\Gamma}) \subset \mathcal{I}(T)$ which consists of all trees $\tilde{\Gamma}'$ obtained from $\tilde{\Gamma}$ in the following way: For any vertex P of $\tilde{\Gamma}$ we consider the ϵ-neighborhood $B_\epsilon(P) \subset \tilde{\Gamma}$ and we replace this finite tree F_n-equivariantly by another finite tree $B'_\epsilon(P)$ with the same endpoints. More precisely:

We first chose the topological type of $B'_\epsilon(P)$, which may well be different from that of $B_\epsilon(P)$, but we require that there is a (fixed) bijection between the set $\partial B'_\epsilon(P)$ of endpoints of $B'_\epsilon(P)$, and the endpoint set $\partial B_\epsilon(P)$. We then choose F_n-equivariantly for any vertex Q of $B'_\epsilon(P) - \partial B'_\epsilon(P)$ an image $i'(Q)$ in $i(B_\epsilon(P))$ in such a way that along any embedded path in $B'_\epsilon(P)$ the order of vertices is preserved by i'. We define the edge lengths in $B'_\epsilon(P)$ as well as the map $i' : \tilde{\Gamma}' \to T$ by extending i' in the canonical way: Outside the $B'_\epsilon(P)$ we let i' agree with i, while the edges of $B'_\epsilon(P)$ are mapped by i' isometrically to geodesic segments of T. The fact that we chose the $i'(Q)$ "vertex-order preserving" inside the "old" $i(B_\epsilon(P))$ ensures that the "two-gate-condition" is satisfied, i.e. $\tilde{\Gamma}'$ is an element of $\mathcal{I}(T)$ and not just of $\mathcal{I}_0(T)$. It follows directly from the topology of $\mathcal{I}(T)$ defined at the beginning of this section that the set $\mathcal{N}_\epsilon(\tilde{\Gamma})$ of such trees is indeed a neighborhood of the point $\tilde{\Gamma}$. Note that the "vertex-order preserving" condition in $B'_\epsilon(P)$ implies that the diameter of $B'_\epsilon(P)$ is precisely equal to 2ϵ.

We now chose $0 < t < \frac{1}{3}\, min\{L(e) \mid e \in Edges(\Gamma)\}$ and $\epsilon > 0$ small with respect to t and with respect to $il(\tilde{\Gamma})$. For any point $\tilde{\Gamma}'$ in $\mathcal{N}_\epsilon(\tilde{\Gamma})$, and for $|t - t'|$ small with respect to ϵ we want to compare the trees $\tilde{\Gamma}_t$ and $\tilde{\Gamma}'_{t'}$ obtained from $\tilde{\Gamma}$ and $\tilde{\Gamma}'$ by canonical folding of length t and t' respectively. For any vertex $P \in \tilde{\Gamma}$ as above and any pair of points Q and Q' in $\partial B_\epsilon(P)$ which satisfy $i(Q) = i(Q')$ (and thus $(i'(Q) = i'(Q'))$ we define reals $t_{Q,Q'} \geq 0$ and $t'_{Q,Q'} \geq 0$ to be the parameter value when Q and Q' are first identified in the canonical folding process at $\tilde{\Gamma}$ and $\tilde{\Gamma}'$. For example, if Q and Q' define an illegal turn at P with maximal identification length, then one has $t_{Q,Q'} = \epsilon$.

Recall now that, by our definition of $\mathcal{N}_\epsilon(\tilde{\Gamma})$, the image in T of a complementary component of $B_\epsilon(P)$ does not change when passing from $\tilde{\Gamma}$ to $\tilde{\Gamma}'$. Since a vertex in the tree $B'_\epsilon(P)$ can not backtrack during the canonical folding process, this implies that for all pairs Q, Q' as above the values $t_{Q,Q'}$ and $t'_{Q,Q'}$ differ by at most ϵ. But this implies:

Proposition 1.12. *For any $\tilde{\Gamma} \in \mathcal{I}(T)$ and any small $t > 0$ there are constants $0 < \epsilon' < \epsilon$ such that for any $|t-t'| < \epsilon'$ and any $\tilde{\Gamma}' \in \mathcal{N}_\epsilon(\tilde{\Gamma})$ the tree $\tilde{\Gamma}'_{t'}$ is contained in $\mathcal{N}_{2\epsilon}(\tilde{\Gamma}_t)$. As a consequence, the canonical folding paths define a continuous semi-flow Φ on $\mathcal{I}(T)$.* □

Aside. It seems interesting (in particular in the train track situation considered in the next section) to ask whether perhaps for any $\tilde{\Gamma}$ in $\mathcal{I}(T)$ the number of illegal turns stabilizes eventually, if one moves $\tilde{\Gamma}$ along its canonical folding path.

2. Train tracks

In this section we will concentrate on a subset of $\mathcal{I}(T)$, for certain special T, which occurs naturally in the context of automorphisms of free groups: For any automorphism $\alpha \in Aut(F_n)$ we consider the set $\hat{\mathcal{E}}(\alpha)$ of *efficient representatives*[1], i.e. marked metric graphs Γ of fixed volume 1 and without vertices of valence 1, together with a map $f : \Gamma \to \Gamma$ that represents α and is *efficient*: every edge of Γ is stretched constantly by a dilatation factor $\lambda > 1$, and away from the vertices all positive iterates of f are immersions. Such representatives of α have been introduced and analyzed in [2], where it has been shown that every *non-periodic irreducible* outer automorphism of F_n has such an efficient representative (which in addition maps vertices to vertices; such cellular efficient representatives are called *train track representatives* of α).

Recall that the class of irreducible non-periodic automorphisms of F_n is a slight extension of the more important class of *irreducible automorphisms with irreducible powers ("iwip")*, which are characterized by the property that no proper non-trivial free factor of F_n is mapped by any non-zero power of the automorphism to a conjugate of itself.

In [5] it has been shown how one derives from an efficient representative $f : \Gamma \to \Gamma$ of $\alpha \in Aut(F_n)$ an $\mathbb{R}$-tree T and an F_n-equivariant and edge-isometric map $i : \tilde{\Gamma} \to T$ as in section 1, as well as a homothety $H : T \to T$ with stretching factor λ which satisfies $i\tilde{f} = Hi$, where $\tilde{f}$ is a lift of f with $\alpha(w)\tilde{f} = \tilde{f}w : \tilde{\Gamma} \to \tilde{\Gamma}$ for all $w \in F_n$. Furthermore, if α is irreducible and non-periodic, it has been shown in Theorem 4.3 (b) of [13] that T and H depend (up to uniform rescaling of T) only on α and not on Γ. Notice that for the approach persued in [13] it is important that vertices are mapped by f to vertices, but the resulting uniqueness of T and H can be exploited of course for any efficient representative of α.

Lemma 2.1. *Let $f : \Gamma \to \Gamma$ and $f' : \Gamma' \to \Gamma'$ be two efficient representatives of α with associated maps $i : \tilde{\Gamma} \to T$ and $i : \tilde{\Gamma}' \to T$, and let $\tilde{j} : \tilde{\Gamma} \to \tilde{\Gamma}'$ be an F_n-equivariant isometry which satifies $i'\tilde{j} = i : \tilde{\Gamma} \to T$. Then $f : \Gamma \to \Gamma$ and $f' : \Gamma' \to \Gamma'$ determine the same point of $\hat{\mathcal{E}}(\alpha)$: The map $\tilde{j}$ is induced by a marking preserving isometry $j : \Gamma \to \Gamma'$ which satisfies $f'j = jf$.*

Proof. The tree T has at least one branch point $Q \in T$, and hence it possesses for some power H^r and some $w \in F_n$ an eigenray ρ at Q (see [5]), i.e. an infinite embedded ray ρ with $wH^r(\rho) = \rho$. After possibly increasing the exponent r the wH^r-eigenray ρ lifts via i to a $w\tilde{f}^r$-eigenray $\tilde{R}$ in $\tilde{\Gamma}$ which starts at a fixed point $\tilde{q} \in \tilde{\Gamma}$, see [5]. Similarly, ρ lifts via i' to give a $w\tilde{f'}^r$-eigenray $\tilde{R}'$ in $\tilde{\Gamma}'$ which starts at a fixed point $\tilde{q}' \in \tilde{\Gamma}'$. From the property BBT of i' it follows that $\tilde{j}(\tilde{R})$ and $\tilde{R}'$ must coincide up to a finite initial segment. Hence $i'\tilde{j}\tilde{f} = i\tilde{f} = Hi = Hi'\tilde{j} = i'\tilde{f}'\tilde{j}$ implies $\tilde{j} \circ w\tilde{f}^r(\tilde{x}) = w\tilde{f'}^r \circ \tilde{j}(\tilde{x})$ for any point $\tilde{x}$ on $\tilde{R}$ minus some finite initial segment. Similarly, if one considers the total collection of all eigenrays acted upon by $w \in F_n$ and by $\tilde{f}$ or $\tilde{f}'$, one obtains $\tilde{j}\tilde{f}(\tilde{x}) = \tilde{f}'\tilde{j}(\tilde{x})$. But since α is irreducible,

[1] this terminology has been introduced by T. White in [16]

the image R in Γ of the eigenray $\tilde{R}$ passes over every edge of Γ infinitely often (see e.g. [13], section 3), and hence the F_n-equivariant isometry j induced by $\tilde{j}$ satisfies $f'j(x) = jf(x)$ for all $x \in \Gamma$. $\square$

We use the above uniqueness of T and H to fix T within its projective class, and to rescale each Γ in $\hat{\mathcal{E}}(\alpha)$ accordingly so that the map $i : \tilde{\Gamma} \to T$ is edge-isometric. This defines a subspace $\mathcal{E}(\alpha)$ of $\mathcal{I}(T)$ and a surjective map from $\hat{\mathcal{E}}(\alpha)$ to $\mathcal{E}(\alpha) \subset \mathcal{I}(T)$, and by Lemma 2.1 this map $\hat{\mathcal{E}}(\alpha) \to \mathcal{E}(\alpha)$ is injective. The automorphism α acts on an element Γ of $\mathcal{E}(\alpha)$ by twisting the marking isomorphism with α, while simultaneously rescaling Γ by the factor λ^{-1}. The resulting element $\alpha^*(\Gamma) \in \mathcal{E}(\alpha)$ has as universal covering a tree $\alpha^*(\tilde{\Gamma})$ that comes with a canonical F_n-equivariant edge-isometric map $i_\alpha : \alpha^*(\tilde{\Gamma}) \to T$.

Below we need the following specifiations in the construction of T from $f : \Gamma \to \Gamma$ (compare [5]).

Remark 2.2. (a) A turn (e, e') in Γ lifts to turn in $\tilde{\Gamma}$ that is illegal (with respect to $i : \tilde{\Gamma} \to T$) if and only if (e, e') is f-illegal (i.e. some positive iterate of f identifies non-trivial initial segments of e and e'). Furthermore, a reduced edge path γ in Γ is *f-legal* (i.e. all $f^m(\gamma)$ with $m \geq 1$ are reduced) if and only if any lift $\tilde{\gamma}$ of γ to $\tilde{\Gamma}$ is mapped by i injectively into T. We call such paths γ or $\tilde{\gamma}$ *legal.*

(b) The map $i : \tilde{\Gamma} \to T$ factors F_n-equivariantly over a map $\tilde{f}_\alpha : \tilde{\Gamma} \to \alpha^*(\tilde{\Gamma})$ which is, when properly interpreted via taking the twist and the rescaling into account, a lift of the map f. It satisfies $i = i_\alpha \tilde{f}_\alpha$. The tree T is obtained in [5] as F_n-equivariant limit of the trees ${\alpha^*}^n(\tilde{\Gamma})$, and i is the limit of the maps $\tilde{f}_{\alpha^n}$.

(c) The above stated fact that $\tilde{f}$ commutes via i with H implies that the map i maps every vertex of $\tilde{\Gamma}$ with at least 3 adjacent gates to a point $Q \in T$ which is a branch point. As a consequence of (b) we also observe that the lengths of the edges of Γ together with the map i determine the map f. This shows that no two distinct efficient representatives of α, i.e. two distinct points of $\hat{\mathcal{E}}(\alpha)$ (or $\mathcal{E}(\alpha)$), both with at least one vertex that has 3 or more adjacent gates, can determine the same point of cv_n or of CV_n.

Proposition 2.3. *For every $\Gamma \in \mathcal{E}(\alpha)$ the canonical folding path starting at Γ is completely contained in $\mathcal{E}(\alpha)$.*

Proof. We first observe that the efficient representative $f : \Gamma \to \Gamma$ (and hence by Remark 2.2 (b) also the map i) folds every turn (e, e') that has maximal identification length. This follows directly from the fact that the edges of Γ are stretched under f by the factor $\lambda > 1$, so that, if (e, e') was not folded, then its f-image in Γ would have identification length bigger than the maximal identification length $il(\Gamma)$.

Hence, for any Γ_t obtained from Γ by canonical folding of sufficiently small length $t > 0$, the map f splits over the canonical folding map $f_t : \Gamma \to \Gamma_t$, i.e. there exists a marking preserving map $g_t : \Gamma_t \to \Gamma$ with $f = g_t f_t$. But then $f_t g_t : \Gamma_t \to \Gamma_t$ stretches every edge homogeneously by the factor λ, and Remark 2.2

(a) implies that away from the vertices all positive iterates of $f_t g_t$ are immersions. Furthermore $f_t g_t$ is conjugate (via the homotopy equivalence f_t) to the map f, so that it too represents α. Thus Γ_t belongs to $\mathcal{E}(\alpha)$.

It is easy to verify that $\mathcal{E}(\alpha)$ is closed in $\mathcal{I}(T)$, so that the above argument suffices to imply the desired claim. □

An immersed path γ in Γ which has endpoints Q and Q' (not necessarily vertices !) that are fixed by f, such that $f(\gamma)$ is homotopic relativ endpoints to γ, and where γ decomposes into two maximal legal subpaths $\gamma = \gamma_1^{-1} \circ \gamma_2$, is called an *indivisible Nielsen path (INP)* for f. It is well known (and follows for example as direct consequence of the finiteness results in [6]) that the number of INP's for f and for all of its positive powers together is finite. For each of them we consider a lift $\tilde{\gamma}$ with endpoints $\tilde{Q}, \tilde{Q}'$ to $\tilde{\Gamma}$ and its image $i(\tilde{\gamma})$ in T. Let $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ be the two maximal legal subpaths of $\tilde{\gamma} = \tilde{\gamma}_1^{-1} \circ \tilde{\gamma}_2$, and let $\tilde{P}$ be their common initial vertex in the "middle" of $\tilde{\gamma}$. It follows from the construction of T (compare Remark 2.2(b)) that $i(\tilde{Q}) = i(\tilde{Q}')$, and that this point is a branch point of T. Recall that in [6] it has been shown that T has only finitely many F_n-orbits of branch points as well as finitely many F_n-orbits of directions at them. We associate to $\tilde{\gamma}$ the following three *Nielsen data*:

(1) The direction of $i(\tilde{\gamma})$ at $i(\tilde{Q}) = i(\tilde{Q}')$.
(2) The set of all directions at $i(\tilde{Q}) = i(\tilde{Q}')$ that contain i-image points of the component of $\tilde{\Gamma} - \tilde{P}$ given by $\tilde{\gamma}_1$.
(3) The set of all directions at $i(\tilde{Q}) = i(\tilde{Q}')$ that contain i-image points of the component of $\tilde{\Gamma} - \tilde{P}$ given by $\tilde{\gamma}_2$.

We say that two trees $\tilde{\Gamma}$ and $\tilde{\Gamma}'$ of $\mathcal{E}(\alpha)$, given by efficient representatives $f : \Gamma \to \Gamma$ and $f : \Gamma' \to \Gamma'$ of α respectively, are in the same *blow-up class*, if for each positive power f^m of f the set of the F_n-orbits of Nielsen data associated to the collection of INP's for f^m agrees with the analogous set for ${f'}^m$.

Remark 2.4. The Nielsen data encode, in particular, the structure of the periodic orbits of $f : \Gamma \to \Gamma$ that are Nielsen equivalent. For instance it captures the subtle notion of weak Nielsen equivalence (periodic points that are Nielsen equivalent for some iterate but might have different periods). This notion is directly related to the notion of "pattern" of periodic orbits studied in [1].

Lemma 2.5. *Let $f : \Gamma \to \Gamma$ any efficient representative of α, and let Γ_t, for any $0 < t < il(\Gamma)$ be obtained from Γ through canonical folding of length t as defined in section 1.*

(a) The trees $\tilde{\Gamma}$ and $\alpha^*(\tilde{\Gamma})$ lie in the same blow-up class of $\mathcal{E}(\alpha)$.

(b) The trees $\tilde{\Gamma}$ and $\tilde{\Gamma}_t$ lie in the same blow-up class of $\mathcal{E}(\alpha)$.

Proof. (a) This follows directly from Remark 2.2 (b), since $\tilde{f}$ commutes via i with H, and since the homotopy class rel. endpoints of any INP for f is fixed by f, while the homotopy classes of INP's for positive powers f^m are permuted by f.

(b) The same argument as in case (a) applies, where the F_n-equivariant map $\tilde{f}_\alpha$ from Remark 2.2 (b) is replaced by a lift $\tilde{f}_t$ of the folding map $f_t : \Gamma \to \Gamma_t$ as in the proof of Proposition 2.3. □

The following is a vital ingredient in the second author's solution of the conjugacy problem for automorphisms of free groups. For the convenience of the reader we explain in a subsequent remark the precise transition from the terminology used in [13] to the terminology used here.

Proposition 2.6 [13]. *If $\alpha \in Aut(F_n)$ is irreducible and non-periodic, then for any $\Gamma \in \mathcal{E}(\alpha)$ there exists an $\epsilon = \epsilon(\Gamma) > 0$ such that for any $\Gamma' \in \mathcal{E}(\alpha)$ in the same blow-up class as Γ, with $il(\Gamma') \leq \epsilon$, there exists a T-compatible map $h : \tilde{\Gamma} \to \tilde{\Gamma}'$.*

Remark 2.7. We first note that Section 5 of [13] is written for "relative pseudo-Anosov maps" $f : \mathcal{G} \to \mathcal{G}$, but these include as a special case (namely the case where the relative part of $\mathcal{G}$ consists only of the vertices of the graph Γ and hence every edge of Γ belongs to the train track part) the particular case of (absolute) train track maps. As the condition that vertices are mapped by f to vertices is not used in the arguments from [13] needed here, they apply directly to efficient representatives $f : \Gamma \to \Gamma$ as considered here.

In Proposition 5.3 of [13] the existence of a map $h : \Gamma \to \Gamma'$ is shown that commutes with the two marking isomorphisms, up to precomposition with a suitably large positive power of α. As this twisting with iterates of α^* is only used to guaratee that Γ' has small $BBT(i')$, which is ensured already by the hypothesis $il(\Gamma') \leq \epsilon$, no such marking twist is needed here.

The map $h : \Gamma \to \Gamma'$ obtained from Proposition 5.3 of [13] furthermore satiesfies $hf = f'h$ "up to Nielsen faces". A Nielsen face is a 2-cell Δ glued to Γ along two of its three faces, namely along the two maximal legal subpaths γ_1 and γ_2 of an INP $\gamma = \gamma_1^{-1} \circ \gamma_2$ for $f : \Gamma \to \Gamma$. The third face of Δ is added as new "auxiliary edge" e to Γ, and f extends to a map which sends Δ to itself union the f-image of γ. For more details see section 3 of [13]. The statement that two maps agree "up to Nielsen faces" means that they are equal after possibly a homotopy which moves the image of some edge mapped to γ_1 over Δ to γ_2, or conversely.

In Lemma 5.5 of [13] this situation is improved, to get $hf = f'h$ honestly (i.e. not just up to Nielsen faces), by modifying Γ (and f) via introducing new Nielsen paths and new Nielsen faces. This "blow-up" procedure is essentially the same as presented in detail in section 3. However, the hypothesis in Proposition 2.6 above, that f and f' belong to the same blow-up class, means precisely that no blow-up of INP's in Γ' is required, so that already the map h given by Proposition 5.3 of [13] satisfies $hf = f'h$, for the original efficient representatives Γ and Γ' considered in Proposition 2.6.

The equation $hf = f'h$, together with the commutativity of h with the marking isomorphismss of Γ and Γ', is equivalent to the statement that suitable lifts $\tilde{f}$, $\tilde{f}'$ and $\tilde{h}$, of the maps f, f' and h respectively, satisfy $\tilde{h}\tilde{f} = \tilde{f}'\tilde{h}$. From the uniqueness of T and from the construction of T detailed in Remark 2.2 (b) it follows that $\tilde{h}$ is T-compatible.

Proposition 2.8. *For any $\Gamma, \Gamma' \in \mathcal{E}(\alpha)$ in the same blow-up class the canonical folding paths starting at Γ and at Γ' will eventually meet (and hence coincide thereafter).*

Proof. One considers the canonical folding path, as defined in the last section, which starts at Γ'. After canonically folding Γ' for a sufficiently long time t (in other words: after moving along an initial subpath of distance t on the canonical folding path that starts at Γ') we reach by Proposition 1.9 a graph Γ'_s with $0 < il(\Gamma'_s) \leq \epsilon(\Gamma)$. By Lemma 2.5 (b) the blow-up classes of Γ' and of Γ'_s coincide. Thus we obtain from Proposition 2.6 a T-compatible map $h : \tilde{\Gamma} \to \tilde{\Gamma}'_s$. Similarly, sufficient large canonical folding at Γ gives a graph Γ_t and a T-compatible map $g : \tilde{\Gamma}'_s \to \tilde{\Gamma}_t$.

We then apply Proposition 1.11 to g to obtain a graph $\Gamma'_{s'}$ and a map $g' : \tilde{\Gamma}'_{s'} \to \tilde{\Gamma}_t$ with $il(\Gamma'_{s'}) = il(\Gamma_t)$. A second application of Proposition 1.11 at Γ to the map h gives a graph $\Gamma_{t'}$ and a map $h' : \tilde{\Gamma}_{t'} \to \tilde{\Gamma}_{s'}$ with $il(\Gamma_{t'}) = il(\Gamma'_{s'})$.

But this implies $t = t'$, and hence Lemma 1.5 applies to show that h' and g' must be equivariant isometries: $\Gamma_{t'} = \Gamma'_{s'} = \Gamma_t$. □

Corollary 2.9. *Every irreducible non-periodic automorphisms α of F_n has a uniquely determined α^*-invariant canonical folding line in each of its blow-up classes.*

Proof. For any efficient representative Γ in $\mathcal{E}(\alpha)$ the canonical folding paths starting at $\tilde{\Gamma}$ and at $\alpha^*(\tilde{\Gamma})$ must meet, by Lemma 2.5 (a) and Proposition 2.8. Since canonical folding and application of α^* both (strictly) decrease the volume of the quotient graph, the folding path starting at their meeting point is mapped by α^* into itself, so that iterating with ${\alpha^*}^{-1}$ gives the desired α^*-invariant folding line. Its uniqueness follows directly from Proposition 2.8 and the α^*-invariance just derived. From Lemma 2.5 (b) it follows that canonical folding does not change the blow-up class, which shows that the just derived α^*-invariant folding line is indeed contained in the blow-up class of the original arbitrarily chosen representative Γ. □

As any irreducible non-periodic $\alpha \in Aut(F_n$ acts on the unique α^*-invariant canonical folding line by "translation" (one has to logarithmize the folding parameter t), it seems appropriate to call this line the *axis* of α in the given blow up class. Since there are only finitely many blow-up classes, it follows that we obtain finitely many such axes in Outer space CV_n, thus proving Theorem II from the introduction. In the next section we will exhibit among them a "best choice" which we will call the "principal axis" of α^*.

Corollary 2.10. *Every blow-up class in $\mathcal{E}(\alpha)$, for $\alpha \in Aut(F_n)$ non-periodic and irreducible, is contractible.*

Proof. From Proposition 1.12 and Proposition 2.8 it follows that each blow-up class can be contracted along the flow lines of the canonical semi-flow Φ (see section 1) to the unique α^*-invariant folding line, which is itself homeomorphic to $\mathbb{R}$ and hence contractible. Notice here that the invariant folding line is properly embedded in $\mathcal{E}(\alpha)$ since the identification length is a continuous function on $\mathcal{E}(\alpha) \subset \mathcal{I}(T)$ which decreases strictly when moving forward along any folding path, see Lemma 1.5. $\square$

3. Unfolding of INP's

The goal of this section is to introduce a techique which permits us to change the blow-up class of elements of $\mathcal{E}(\alpha)$ in a continuous fashion, similar to the canonical semiflow Φ as defined in section 1 through the canonical folding lines. This is done in two steps: We first show how to modify a given efficient representative by introducing new INP's . In a second step we elaborate this precedure in order to ensure that the modifications can be done in a continuous fashion throughout $\mathcal{E}(\alpha)$.

In order to change the blow-up class of a given $\tilde{\Gamma} \in \mathcal{E}(\alpha)$ we need to modify Γ by introducing new INP's: We consider a branch point Q in Γ which is fixed by f. Let e be an *eigen edge* of f, i.e. an edge of Γ raying out of Q such that the edge path $f(e)$ starts with e. We first need to modufy Γ slightly by folding initial segments of all edges in the same gate as e to guarantee that e is the only edge in its gate at Q. This *preliminary folding* is defined as follows:

As Q is fixed and $f(e)$ starts with e, any edge e' with initial vertex Q belongs the same gate as e if and only if some positive iterate $f^k(e')$ starts with e. Let $k(e')$ be the smallest such integer $k > 0$. These $k(e')$ are bounded above by a number M which depends only on the largest possible number of edges in that gate, and hence only on the rank n of F_n. We subdivide e by new vertices $P_0, P_1 = f(P_0), P_2 = f(P_1), \ldots, P_{M-1} = f(P_{M-2})$, such that $P_M = f(P_{M-1})$ is the midpoint of e. It follows now that any edge e' in the same gate as e has an initial segment that can be folded isometrically onto the segment $[Q, P_{m-k(e')}]$, such that f induces a well defined efficient map on the quotient graph. In this quotient graph, from now on again called Γ, the edge e is the only edge in its gate at Q. Since we only folded up to the middle of e, no INP of f or any positive iterate can possibly have been folded, so that the blow-up class is preserved by the preliminary folding

In order to introduce a new INP *in the direction of* e we need the following definition: Two edges e', e'', also with initial vertex Q, but different from e, are *elementary N-equivalent*, if there exists some edge e''' in Γ (not necessarily adjacent to Q) and some $k \geq 1$ such that $f^k(e''')$ crosses over $\bar{e}'_* e''_*$ or over $\bar{e}''_* e'_*$, where e'_* is any edge in the same gate as e', and e''_* is in the same gate as e''. We consider the N-equivalence classes $N_1, N_2, \ldots$ of edges at Q generated by this elementary N-equivalence. We now replace the initial segment $e_0 = [Q, P_0]$ of e by finitely many copies $e_1, e_2, \ldots$ of the same length, one for each N-equivalence class N_i, all with terminal vertex P_0, but with a new initial vertex Q_i. Any other edge e_* in Γ with initial vertex Q belongs to some N-equivalence class N_j: We disconnect e_* from Q and connect it to Q_j. This modification, together with the above defined preliminary folding, is called *the total blow-up* of the edge e of length $\epsilon = length(e_0)$. Notice that there is a canonical ("folding") projection $\pi : \Gamma' \to \Gamma$ from the blown-up graph Γ' back to the graph Γ, which maps every copy e_i of e_0 homeomorphically back to e_0 (and hence every Q_i to Q) and is induced by the identity map elsewhere.

If Q' is any vertex of Γ with $f^k(Q') = Q$, and if e' is an edge with initial vertex

Q' such that $f^k(e')$ starts with e (and no smaller positive k satisfies this property), we pull back the total blow-up of e (including the preliminary folding !) by f^k and do the anlogous modification there, to get a *preimage blow-up of length* $\frac{\epsilon}{\lambda^k}$ at Q'.

We can then define an efficient map f' which represents the same automorphism as f on the graph Γ' obtained from total blow-up of the edge e as above, and of all preimage blow-ups as considered in the previous paragraph: Let $x \in \Gamma'$ be an arbitrary point, and let $\pi : \Gamma' \to \Gamma$ denote again the canonical folding projection. Whenever $f(\pi(x))$ does not belong to e_0 or to one of its "preimage blow-up initial segments" of e', then we define $f'(x) = \pi^{-1} f \pi(x)$. If $f(\pi(x))$ belongs to e_0 then we consider the $\frac{\epsilon}{\lambda}$-neighborhood η of $\pi(x)$: By the above definition of the total blow-up of e it follows that in this case there is precisely one of the N-equivalence classes, say N_i, such that all of η is mapped by f to the union of e with all edges that belong to the N-equivalence class N_i. Hence there is a canonical choice of which of the copies of $f(\pi(x))$ in the blown-up graph Γ' is the right one to be defined as $f'(x)$. The analogous argument applies if $f(\pi(x))$ belongs to one of the preimage blow-up initial segments of an edge e' as considered in the last paragraph.

This definition yields $\pi f' = f \pi$, and it is easy to verify that the modified map f' is also an efficient representative and that it represents the same automorphism α as the unmodified one. We say that (Γ', f') is obtained from (Γ, f) by *maximal blow-up of length* ϵ of the edge e.

Similarly, we can define the analogous maximal blow-up (including the preliminary folding) for a vertex Q and a direction defined by an edge e which are fixed by some f^k, if we simply blow up all of its (finite) f-orbit of directions (each by the same length $\epsilon > 0$). Since there are only finitely many vertices with more than two gates in Γ, an iterative sequence of maximal blow-ups of length ϵ must eventually terminate, and the obtained graph Γ' is said to be obtained from Γ by *maximal blow-up* of length ϵ. It is not hard to verify that this maximally blown-up Γ' does not depend on the order in which the single blow-ups have been performed, but only on the length ϵ. Hence the blow-up class of a maximally blown-up graph Γ' is uniquely determined the originally given graph Γ. But an even stronger statment is true:

Lemma 3.1. *All maximally blown-up graphs of $\mathcal{E}(\alpha)$ belong to the same blow-up class, called "principal blow-up class" of α.*

Proof. We consider any branch point Q in $\tilde{\Gamma}$ and any eigen edge e with initial vertex Q. The above defined N-equivalence for the directions at Q defines an equivalence (also called N-equivalence) for their images at $i(Q)$. We observe that the preliminary folding doesn't change this N-equivalence at $i(Q)$, and that the N-equivalence classes at $i(Q)$ are in 1:1 correspondence with set of vertices $Q_1, Q_2, \ldots$ newly introduced by the above blow-up procedure. But N-equivalence in T can be defined independently of the particular efficient representative Γ considered in the blow-up: It suffices to observe that every edge in $\tilde{\Gamma}$ is (by irreducibility of α) contained in some F_n-translate of any in $\tilde{\Gamma}$: this is an infinite (legal) ray $\tilde{R}$ with the property $\tilde{f}(\tilde{R}) = \tilde{R}$, compare [5]. But $\tilde{f}$-eigenrays in $\tilde{\Gamma}$ are mapped by i to

H-eigenrays in T, so that elementary N-equivalence in T is alternatively given by the analogously defined elementary "eigenray equivalence". This notion is clearly independent of the particular efficient representative $f : \Gamma \to \Gamma$ in question, and hence the uniqueness of T and H described at the beginning of section 2 shows the desired claim. $\square$

For the proof of the next proposition we need to be able to increase (a) the length $\epsilon > 0$ of the above defined total blow-up of an edge e, or (b) the length of an INP $\gamma = \gamma'' \circ {\gamma'}^{-1}$ which exists already in Γ, to a length that varies continuously with the point in $\mathcal{E}(\alpha)$ determined by the graph Γ and hence should not depend on the cell structure of Γ. For this purpose we define a new parameter $K_0 = K_0(\Gamma)$ which is the minimal length of any legal path which has as endpoints vertices with 3 or more gates. We consider (in case (b)) the initial points Q' and Q'' of the two legal branches γ' and γ'' of γ. For any real K with $0 < K < K_0$ we define $\mathcal{L}_K$ to be in case (a) the set of legal paths of length K which start at Q at the gate of the edge e. Similarly, in case (b) $\mathcal{L}_K$ is the set of legal paths of length K which start at Q' or Q'' at the gate determined by γ' or γ'' respectively. Clearly the set $\mathcal{L}_K$ is finite.

We now observe that, since any vertex P in the interior of any of the paths η from $\mathcal{L}_K$ possesses precisely two gates, namely the "entrance" and the "exit" gate of η at P, it follows in case (a) that there is precisely one *eigen path* γ_* in $\mathcal{L}_K$, i.e. a path with the property that $f(\gamma_*)$ has γ_* as initial subpath. Similarly, in case (b) there is precisely one such path γ'_* for γ', and one such path γ''_* for γ''. Either K is smaller than the length of γ' (equal to that of γ''), a case which we do not need to consider any further, as the original INP γ was already long enough. Otherwise the union of γ'_* with γ''_* (which we denote below sometimes by γ_*, in case (b)) is precisely the original INP γ together with a finite path γ_0 attached at its cusp, which is the canonical prolongation of γ defined by the dynamics of the map f.

We now mimick the preliminary folding procedure introduced at the beginning of our above defined maximal blow-up, but with γ_* playing the role of e, and $\mathcal{L}_K$ playing the role of the set of edges in the same gate as e. We observe that the cardinality of $\mathcal{L}_K$ is bounded above by a constant which depends only on the rank n of F_n. Furthermore, for any $\eta \in \mathcal{L}_K$ there is a positive integer k such that $f^k(\eta)$ has γ_* in case (a), or one of γ'_* or γ''_* in case (b), as initial subpath. As in the above preliminary folding procedure, it follows that there is a common such exponenent k for all $\eta \in \mathcal{L}_K$, and the smallest such common exponent is bounded above by some integer $M > 0$ which only depends on the rank n. One divides γ_* into M subintervalls in precisely the same fashion as done above for e by the P_i introduced in the preliminary folding procedure. As in that case we can then fold isometrically a non-trivial initial segment of any of the $\eta \in \mathcal{L}_K$ onto some initial segment of γ_* in case (a), or of γ'_k or γ''_k in case (b), and f induces a well defined efficient map on the folded quotient graph. The crucial difference to the above preliminary folding procedure is that the length of the resulting unique edge raying out of Q is now $\frac{K}{2\lambda^M}$, which does not any more depend on the length of the original edge e on which the blow-up started.

One then continues in precisely the same fashion as explained above to obtain a *maximal blow-up of length* $\frac{K}{2\lambda^M}$ of the edge e in case (a), or a *canonical extension to length* $\frac{K}{2\lambda^M}$ of the INP γ in case (b). If we perform, at a given efficient representative $f : \Gamma \to \Gamma$, and for some fixed $0 < K < \frac{K_0(\Gamma)}{2\lambda^M}$ and all $k \geq 1$, the maximal blow-up of length K of any f^k-eigen edge e, and a canonical extension to length K of any INP γ for f^k, we obtain the *K-maximally blown-up graph* Γ^K, and an induced efficient map $f_K : \Gamma^K \to \Gamma^K$ that represents α. It follows from Lemma 3.1 that Γ^K lies in the principal blow-up class of α.

Proposition 3.2. *There is a continuous function $K_1 : \mathcal{E}(\alpha) \to \mathbb{R}$ with $0 < K_1(\Gamma) < \frac{K_0(\Gamma)}{2\lambda^M}$ in $\mathcal{E}(\alpha)$, such that, for $0 \leq s \leq 1$ the map $(\Gamma, s) \mapsto \Gamma^{sK_1(\Gamma)}$ describes a continuous deformation of $\mathcal{E}(\alpha)$ into its principal blow-up class.*

Proof. We consider the embedding $\mathcal{E}(\alpha) \subset \mathcal{I}(T)$ and recall the cellular structure of $\mathcal{I}(T)$ introduced at the beginning of section 1. The function $K_0 : \Gamma \to K_0(\Gamma)$ is not continuous on $\mathcal{E}(\alpha)$: One can produce easily a situation where a small edge is continuously been folded so that (only) at the end of the folding process a vertex of valence 3 or larger is folded onto a legal path in some $\mathcal{L}_{K_0}$, thus discontinously decreasing the value of the function K_0. On the other hand, the function K_0 is continuous in the interior of any of the cells $\mathcal{C}(\Gamma^{top}) \subset \mathcal{I}(T)$ defined in Remark 1.1: In fact, the value of the function K_0 depends within each cell $\mathcal{C}(\Gamma^{top})$ only on T and not on the particular metric graph $\Gamma \in \mathcal{C}(\Gamma^{top})$, so that K_0 is constant on $\mathcal{C}(\Gamma^{top})$. Hence, as the cellular structure of $\mathcal{I}(T)$ is locally finite (compare Remark 1.1), we can easily find a continuous function K_1 defined on all of $\mathcal{E}(\alpha)$, with values strictly between 0 and $\frac{K_0(\Gamma)}{2\lambda^M}$. Thus there exists a well defined function Ψ from $[0, 1] \times \mathcal{E}(\alpha)$ into $\mathcal{E}(\alpha)$ as claimed in the statement of the Proposition, and the value of this function at time $s = 1$ is an efficient representative of α in the principal blow-up class. The fact that Ψ is continuous can be verified with the technique introduced in section 1 to prove Proposition 2.10; since it does not need any further ingredients it is left here to the reader. $\square$

The combination of Proposition 3.2 together with Lemma 3.1 and Corollary 2.10 now gives an immediate proof of our Theorem I from the introduction.

References

1. J. Los et al., preprint.
2. M. Bestvina, M. Handel, *Train tracks for automorphisms of the free group*, Ann. Math. **135** (1992), 1–51.
3. M. Culler, K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), 91–119.
4. D. Gaboriau, A. Jaeger, G. Levitt, M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. Jour. **93** (1998), 425–452.
5. D. Gaboriau, G. Levitt, *The rank of actions on* **R***-trees*, Ann. Sc. ENS **28** (1995), 549–570.
6. J. Fehrenbach, Ph.D.-thesis, Nice 1998.
7. V. Guirardel, G. Levitt, preliminary preprint.
8. A. Jäger, Ph.D.-thesis, Bochum 1998.
9. G. Levitt, M. Lustig, *Irreducible automorphisms of F_n have North-South dynamics on compactified Outer space*, Journal of the Inst.of Math.Jussieu **2** (2003), 1–14.
10. J. Los, unpublished preprint.
11. J. Los, *On the conjugacy problem for automorphism of free groups*, Topology **35** (1996), 779–808. With an addendumby the author.
12. M. Lustig, *Automorphismen von freien Gruppen*, Habilitationsschrift, Bochum 1992.
13. M. Lustig, *Structure and conjugacy for automorphisms of free groups I, II*, MPI- Preprint Series **130** (2000) and **4** (2001).
14. R. Skora, *Deformations of length functions in groups*, preprint 1989.
15. J. Stallings, *Topology of finite graphs*, Invent.Math. **71** (1983), 551–565.
16. T. White, unpublished preprint.
17. Williams, *One-dimensional non-wandering sets*, Topology **6** (1967), 473–487.

J.L.: LATP, Univ. d'Aix-Marseille I, 39 rue F. Joliot Curie, 13453 Marseille 13, France.

E-mail address: los@cmi.univ-mrs.fr

M.L.: LATP, Univ. d'Aix-Marseille III, Ave Esq. Normandie-Niemen, 13397 Marseille 20, France.

E-mail address: Martin.Lustig@univ.u-3mrs.fr