Max-Planck-Institut für Mathematik
Bonn

Seven steps to happiness

by

Ilya Kapovich
Martin Lustig

Max-Planck-Institut für Mathematik
Preprint Series 2011 (21)

Seven steps to happiness

Ilya Kapovich
Martin Lustig

Max-Planck-Institut für Mathematik
Vivatsgasse 7
53111 Bonn
Germany

Department of Mathematics
University of Illinois at Urbana-Champaign
1409 West Green Street
Urbana, IL 61801
USA

Mathématiques (LATP)
Université Paul Cézanne - Aix Marseille III
52 av. Escadrille Normandie-Niémen
13397 Marseille 20
France

MPIM 11-21

SEVEN STEPS TO HAPPINESS

ILYA KAPOVICH AND MARTIN LUSTIG

ABSTRACT. For every atoroidal iwip automorphism φ of F_N (i.e. the analogue of a pseudo-Anosov mapping class) it is shown that the algebraic lamination dual to the forward limit tree $T_+(\varphi)$ is obtained as "diagonal closure" of the support of the backward limit current $\mu_-(\varphi)$. This diagonal closure is obtained through a finite procedure in analogy to adding diagonal leaves from the complementary components to the stable lamination of a pseudo-Anosov homeomorphism. We also give several new characterizations as well as a structure theorem for the dual lamination of $T_+(\varphi)$, in terms of Bestvina-Feighn-Handel's "stable lamination" associated to φ.

1. INTRODUCTION

For a closed surface Σ with Euler characteristic $\chi(\Sigma) < 0$ geodesic laminations play an important theoretical role, for many purposes. In particular, if such a lamination $\mathfrak{L}$ is equipped with a transverse measure μ, it becomes a powerful tool, for example in the analysis of the mapping class $[h]$ of a homeomorphism $h : \Sigma \to \Sigma$. The set of projective classes $[\mathfrak{L}, \mu]$ of such *measured laminations* carries a natural topology, and the issuing space $\mathcal{PML}(\Sigma)$ of projective measured laminations is homeomorphic to a high dimensional sphere: One of Thurston's fundamental results shows that $\mathcal{PML}(\Sigma)$ serves naturally as boundary to the Teichmüller space $\mathcal{T}(\Sigma)$ of Σ, and the action of the mapping class group $\mathrm{Mod}(\Sigma)$ extends canonically to the compact union $\mathcal{T}(\Sigma) \cup \mathcal{PML}(\Sigma)$.

In this paper we are interested in the "cousin world", where Σ (or rather $\pi_1\Sigma$) is replaced by a free group F_N of finite rank $N \geq 2$, the mapping class group $\mathrm{Mod}(\Sigma)$ is replaced by the outer automorphism group $\mathrm{Out}(F_N)$, and the role of Teichmüller space $\mathcal{T}(\Sigma)$ is (usually) taken on by Outer space CV_N. These and the other terms used in this introduction will be explained in section 2.

Date: April 10, 2011.

2000 *Mathematics Subject Classification.* Primary 20F36, Secondary 20E36, 57M05.

Key words and phrases. laminations, free groups, $\mathbb{R}$-trees, iwip automorphisms.

The first author was supported by the NSF grant DMS-0904200.

In the $\mathrm{Out}(F_N)$-setting, both, topological laminations $\mathfrak{L}$ and measured laminations $(\mathfrak{L}, \mu)$ have natural analogues, but the situation is quite a bit more intricate: There are two competing analogues of measured lamination, namely $\mathbb{R}$-trees T with isometric F_N-action, and currents μ over F_N. Both, the space of such $\mathbb{R}$-trees as well as the space of such currents have projectivizations which can be used to compactify CV_N (in two essentially different ways, see [11]), and both are used as important tools to analyze single automorphisms of F_N. Also, both, an $\mathbb{R}$-tree T and a current μ, determine an *algebraic lamination*, denoted $L^2(T)$ (the *dual lamination* of T) and $L^2(\mu) = \mathrm{supp}(\mu)$ (the *support* of μ) respectively. Algebraic laminations L^2 for a free group F_N are the natural analogues of (non-measured) geodesic laminations $\mathfrak{L}$ on a surface Σ (or rather, of their lift $\widetilde{\mathfrak{L}} \subset \widetilde{\Sigma}$ to the universal covering $\widetilde{\Sigma}$ of Σ: the set L^2 consists of F_N-orbits of pairs $(X, Y) \in \partial^2 F_N := \partial F_N \times \partial F_N \smallsetminus \{(Z, Z) \mid Z \in \partial F_N\}$).

The two spaces which serve (after projectivization) as analogues of $\mathcal{T}(\Sigma) \cup \mathcal{PML}(\Sigma)$, namely the space $\overline{\mathrm{cv}}_N$ of very small $\mathbb{R}$-tree actions of F_N, and the space $\mathrm{Curr}(F_N)$ of currents over F_N, are related naturally by an $\mathrm{Out}(F_N)$-equivariant continuous *intersection form* $\langle \cdot, \cdot \rangle : \overline{\mathrm{cv}}_N \times \mathrm{Curr}(F_N) \to \mathbb{R}_{\geq 0}$ (see [12]). In [13] it has been shown that $\langle T, \mu \rangle = 0$ if and only if $\mathrm{supp}(\mu) \subset L^2(T)$. Sometimes this inclusion is actually an equality: for example, if $(\mathfrak{L}, \mu)$ is a measured lamination of a surface Σ with boundary, and T_μ is the $\mathbb{R}$-tree transverse to the lift of $(\mathfrak{L}, \mu)$ to the universal covering of Σ, and if $\mathfrak{L}$ is *maximal* (i.e. every complementary component has either 3 cusps, or 1 boundary component plus 1 cusp), then the current $\widehat{\mu}$ defined by $(\mathfrak{L}, \mu)$ satisfies $\mathrm{supp}(\widehat{\mu}) = L^2(T_\mu)$.

However, in general one can not deduce from $\langle T, \mu \rangle = 0$ the equality $\mathrm{supp}(\mu) = L^2(T)$, since the dual laminations of $\mathbb{R}$-trees are by nature always *diagonally closed*, i.e. the *diagonal closure* $\overline{\mathrm{diag}}(L^2(T))$ is equal to $L^2(T)$ for all $T \in \overline{\mathrm{cv}}_N$. On the other hand, the support of a current may well be minimal and thus typically a proper subset of its diagonal closure. Moreover, it is easy to find *perpendicular* pairs (T, μ) (i.e. T and μ satisfy $\langle T, \mu \rangle = 0$) where even the diagonal closure $\overline{\mathrm{diag}}(\mathrm{supp}(\mu))$ is only a proper subset of $L^2(T)$. A class of examples with even stronger properties is described in section 7; alternatively, any measured lamination $(\mathfrak{L}, \mu)$ on a surface which has a complementary component with non-abelian π_1 defines a perpendicular pair $(T_\mu, \widehat{\mu})$ as above, with $\overline{\mathrm{diag}}(\mathrm{supp}(\widehat{\mu})) \neq L^2(T_\mu)$.

For pseudo-Anosov mapping classes $[h]$ of Σ it is well known that the induced action on $\mathcal{T}(\Sigma) \cup \mathcal{PLM}(\Sigma)$ has uniform North-South dynamics. In particular, $[h]$ has precisely one attracting and one repelling fixed point on $\mathcal{PLM}(\Sigma)$. These very statements are also true for the induced actions of *atoroidal iwip automorphisms* of F_N on both "cousin spaces", $\overline{\mathrm{CV}}_N$ and $\mathbb{P}\mathrm{Curr}(F_N)$, so that the (much studied) class of atoroidal iwip automorphisms should be considered as strict analogue of the class of pseudo-Anosov mapping classes. In [13] it has been shown that for any atoroidal iwip $\varphi \in \mathrm{Out}(F_N)$ the *forward limit tree* $T_+ = T_+(\varphi) \in \overline{\mathrm{cv}}_N$, which defines the attracting fixed point $[T_+] \in \overline{\mathrm{CV}}_N$, is perpendicular to the *backward limit current* $\mu_- = \mu_-(\varphi) \in \mathrm{Curr}(F_N)$, which gives the repelling fixed point $[\mu_-] \in \mathbb{P}\mathrm{Curr}(F_N)$:

$$\langle T_+(\varphi), \mu_-(\varphi)\rangle = 0 \quad \text{for any atoroidal iwip} \quad \varphi \in \mathrm{Out}(F_N)$$

The main result of this paper can now be stated as follows:

Theorem 1.1. *Let φ be an atoroidal iwip automorphism of F_N. Let T_+ be its forward limit tree, and let μ_- be its backward limit current. Then the dual lamination $L^2(T_+)$ and the diagonal closure of the support supp(μ_-) satisfy:*

$$L^2(T_+) = \overline{diag}(supp(\mu_-))$$

This result raises the question, what the precise conditions are, which allow one to deduce for perpendicular pairs $(T, \mu) \in \overline{\mathrm{cv}}_N \times \mathrm{Curr}(F_N)$ the analogous conclusion, i.e. that $L^2(T) = \overline{\mathrm{diag}}(\mathrm{supp}(\mu))$. A necessary condition for such a conclusion is that μ "fills out" enough of the available room in F_N which is "potentially dual" to T. Making such an indication precise takes more room than available here in the introduction. We will discuss those matters below in section 7.

A second result of this paper, which is stronger than Theorem 1.1, concerns the structure of the lamination $L^2(T_+)$. The proof of this result is really the main purpose of this paper; we show below how to derive Theorem 1.1 from the following:

Theorem 1.2. *Let φ be an atoroidal iwip automorphism of F_N, and let T_+ be its forward limit tree. Then there exists a sublamination $L^2_{BFH} \subset L^2(T_+)$ which satisfies:*

(1) *L^2_{BFH} is the "stable lamination" exhibited by Bestvina-Feighn-Handel in [2] for φ^{-1}. In particular, L^2_{BFH} is minimal and non-empty.*
(2) *L^2_{BFH} is the only minimal and non-empty sublamination of $L^2(T_+)$.*
(3) *$L^2(T_+) = \overline{diag}(L^2_{BFH})$*
(4) *$L^2(T_+) \smallsetminus L^2_{BFH}$ is a finite union of F_N-orbits of pairs $(X, Y) \in \partial^2 F_N$.*

The arguments given below prove also that L^2_{BFH} is contained in $\mathrm{supp}(\mu_-)$. Indeed, it can be shown (by a direct argument) that the two laminations are equal.

In the course of our proof we also give several alternative characterizations of the dual lamination $L^2(T_+)$, based on (absolute) train track representatives of φ or of φ^{-1}. The precise terminology of the terms used is given in sections 4 and 5 below.

Proposition 1.3. *Let φ be an atoroidal iwip automorphism of F_N, and let T_+ be its forward limit tree. Let $f_+ : \tau_+ \to \tau_+$ and $f_- : \tau_- \to \tau_-$ be stable train track maps that represent φ and φ^{-1} respectively.*

(1) The dual lamination $L^2(T_+)$ consists precisely of all pairs $(X, Y) \in \partial^2 F_N$ such that for some $C > 0$ the whole φ-orbit $\varphi^t(X, Y)$ is totally C-illegal with respect to f_+.

(2) The dual lamination $L^2(T_+)$ consists precisely of all pairs $(X, Y) \in \partial^2 F_N$ such that the whole φ-orbit is uniformly φ-contracting (or, equivalently, uniformy φ^{-1}-expanding).

(3) The dual lamination $L^2(T_+)$ consists precisely of all pairs $(X, Y) \in \partial^2 F_N$ such that the whole φ-orbit is used legal with at most one singularity, with respect to f_-.

This proposition is a direct consequence of the fact that in the course of our proof of Theorem 1.2 we show first that $L^2(T_+)$ is included in the set of pairs from $\partial^2 F_N$ defined by the the property stated in part (1) of Theorem 1.2. One then shows that this set of pairs is contained in the set defined by part (2), and that the latter is contained in the set defined by part (3). In a final step one shows the set of (X, Y) defined in part (3) is contained in the diagonal closure of L^2_{BFH}. But by part (3) of Theorem 1.2 the lamination $L^2(T_+)$ is equal to this diagonal closure, so that all of the previous inclusions must be in fact euqualities.

Using the same type of arguments one deduces Theorem 1.1 from Theorem 1.2: From the perpendicularity of T_+ with μ_- we derived above that $\mathrm{supp}(\mu_-)$ is a subset of $L^2(T_+)$, which implies (compare Definition-Remark 2.5) $\overline{\mathrm{diag}}(\mathrm{supp}(\mu_-)) \subset \overline{\mathrm{diag}}(L^2(T_+))$ Since $L^2(T_+)$ is diagonally closed (see Proposition 2.10), we obtain:

$$\overline{\mathrm{diag}}(\mathrm{supp}(\mu_-)) \subset L^2(T_+)$$

On the other hand, from the uniqueness property given by part (2) of Theorem 1.2 one deduces that L^2_{FBH} is contained in $\mathrm{supp}(\mu_-)$. Hence (again by Definition-Remark 2.5) $\overline{\mathrm{diag}}(L^2_{FBH})$ must be contained in $\overline{\mathrm{diag}}(\mathrm{supp}(\mu_-))$, so that part (3) of Theorem 1.2 yields

$$L^2(T_+) \subset \overline{\mathrm{diag}}(\mathrm{supp}(\mu_-)) \, .$$

Acknowledgements: This paper has been influenced by many discussions with colleagues in our scientific neighborhood. In particular we would like to thank T. Coulbois and A. Hilion, and also point to their forthcoming joint work with P. Reynolds, where much related (but harder) questions are treated with different (and distinctly deeper) methods.

Personal comment: The title of our paper goes back to the very origin of this paper, where the more optimistic of the two authors proposed a proof-scheme in seven steps, and the other, the more skeptical between us, decided that it reminded him of one of these advertised programs which guarantee personal fulfillment, wealth, or religious enlightenment. Since then, the result has been internally referred to by the title used here, and in particular the headers of the later chapters are still a reminiscence to the beginning of this particular collaboration.

Organisation of the paper:

- In §2 we recall the basic facts and definitions about $\mathbb{R}$-trees, Outer space and iwip automorphisms. We also review some of the basics of algebraic laminations, define the diagonal closure of a lamination, and recall some of the known facts about the dual lamination of an $\mathbb{R}$-tree.
- In §3 we recall the basic facts and definitions of train track maps. We purposefully employ space and care to make this section accessible for, say, a graduate student who is only partially an expert of the subject.
- In §4 we show that the dual lamination $L^2(T_+)$ of the forward limit tree T_+ of an atoroidal iwip automorphism φ has a strong contraction property with respect to the action of φ. This implies that $L^2(T_+)$ has a strong expansion property with respect to the action of φ^{-1}
- In §5 we will show that any lamination with a strong expansion property with respect to φ^{-1}, when realized as geodesic paths on a train track representative $f_- : \tau_- \to \tau_-$ of φ^{-1}, consists only of paths which are entirely "used legal", or else have precisely one "singularity" (the terms will be defined there). It will also be shown that the sublamination of $L^2(T_+)$ which is realized by used legal geodesics on τ_- is contained in Bestvina-Feighn-Handel's "attracting" lamination $L^2_{BFH}(f_-)$ associated to the train track map f_-.
- In §6 we consider elements $(X, Y) \in L^2(T_+)$ which are realized on τ_- by used legal geodesics with precisely one singularity. We

show that such geodesics have a rather special form: they are (essentially) the concatenation of two eigenrays of the map f_-. As a consequence, we show that such pairs (X, Y) must lie in the diagonal closure of $L^2_{BFH}(f_-)$. Also, the finiteness of such pairs (up to the F_N-action) as well as the fact that they are non-rational is a direct consequence of this characterization.
- In section 7 we discuss some question issuing from our main result and the surrounding facts.

2. Definition of terms and background

The purpose of this section is to properly define the terms used in the introduction, and to give some background with references about them.

Throughout the paper we fix a "model" free group F_N of finite rank $N \geq 2$. A finite connected graph τ is called a *marked graph*, if it is equipped with a *marking isomorphism* $\theta : F_N \overset{\cong}{\longrightarrow} \pi_1(\tau)$. Here we purposefully suppress the issue of choosing a basepoint of τ, as in this paper we are interested in automorphims of F_N only up to inner automorphisms.

2.1. **Outer space.**

We give here only a brief overview of basic facts related to Outer space. We refer the reader to [9] for more detailed background information.

The *unprojectivized Outer space* cv_N consists of all minimal free and discrete isometric actions on F_N on $\mathbb{R}$-trees (where two such actions are considered equal if there exists an F_N-equivariant isometry between the corresponding trees). There are several different topologies on cv_N that are known to coincide, in particular the equivariant Gromov-Hausdorff convergence topology and the so-called *length function* topology. Every $T \in \mathrm{cv}_N$ is uniquely determined by its *translation length function* $||.||_T : F_N \to \mathbb{R}$, where $||g||_T := \min\{d(gx, x) \mid x \in T\}$ is the *translation length* of g on T. Two trees $T_1, T_2 \in \mathrm{cv}_N$ are close if the functions $||.||_{T_1}$ and $||.||_{T_1}$ are close pointwise on a large ball in F_N. The closure $\overline{\mathrm{cv}}_N$ of cv_N in either of these two topologies is well-understood and known to consists precisely of all the so-called *very small* minimal isometric actions of F_N on $\mathbb{R}$-trees, see [1] and [4].

The outer automorphism group $\mathrm{Out}(F_N)$ has a natural continuous right action on $\overline{\mathrm{cv}}_N$ (that leaves cv_N invariant) given at the level of length functions as follows: for $T \in \overline{\mathrm{cv}}_N$ and $\varphi \in Out(F_N)$ we have

$||g||_{T\varphi} = ||\Phi(g)||_T$, with $g \in F_N$ and $\Phi \in Aut(F_N)$ representing $\varphi \in \mathrm{Out}(F_n)$.

The *projectivized Outer space* $\mathrm{CV}_N = \mathbb{P}\mathrm{cv}_N$ is defined as the quotient $\mathrm{cv}_N/\sim$ where for $T_1 \sim T_2$ whenever $T_2 = cT_1$ for some $c > 0$. One similarly defines the projectivization $\overline{\mathrm{CV}}_N = \mathbb{P}\overline{\mathrm{cv}}_N$ of $\overline{\mathrm{cv}}_N$ as $\overline{\mathrm{cv}}_N/\sim$ where $\sim$ is the same as above. The space $\overline{\mathrm{CV}}_N$ is compact and contains CV_N as a dense $\mathrm{Out}(F_N)$-invariant subset. The compactification $\overline{\mathrm{CV}}_N$ of CV_N is a free group analogue of the Thurston compactification of the Teichmüller space. For $T \in \overline{\mathrm{cv}}_N$ its $\sim$-equivalence class is denoted by $[T]$, so that $[T]$ is the image of T in $\overline{\mathrm{CV}}_N$.

2.2. **Laminations and the diagonal closure.**

For the free group F_N we define the *double boundary*

$$\partial^2 F_N := \{(X, Y) \in \partial F_N \times \partial F_N \mid X \neq Y\}.$$

The set $\partial^2 F_N$ comes equipped with a natural topology, inherited from and $\partial F_N \times \partial F_N$, and with a natural translation action of F_N by homeomorphisms. There is also a natural *flip* map $\partial^2 F_N \to \partial^2 F_N$, $(X, Y) \mapsto (Y, X)$, interchanging the two coordinates on $\partial^2 F_N$.

The following definition has been introduced and systematically studied in [6], where also further background material concerning this subsection can be found.

Definition 2.1. An *algebraic lamination* on F_N is a non-empty subset $L^2 \subseteq \partial^2 F_N$ which is (i) closed, (ii) F_N-invariant, and (iii) flip-invariant.

In analogy to laminations on surfaces, the elements $(X, Y) \in L^2$ are sometimes also referred to as the *leaves* of the algebraic lamination L^2.

In some circumstances it is useful to admit the empty set as algebraic lamination. We will formally stick to the classical non-empty notion, but occasionally informally include the empty set in our discussions about algebraic lamination.

Example 2.2. Let Ω be a set of conjugacy classes $[w]$ with $w \in F_N$. Denote by w^∞ and $w^{-\infty}$ the attracting and the repelling fixed points respectively of the action of w on ∂F_N. Then

$$L^2(\Omega) := \overline{\{(vw^\infty, vw^{-\infty}) \mid v \in F_N, [w] \in \Omega \cup \Omega^{-1}\}}$$

is an algebraic lamination.

Remark 2.3. (a) Note that, in the above Example 2.2, for finite Ω taking the closure on the right hand side of the displayed equality can be omitted without changing the set $L^2(\Omega)$.

(b) Thus for finite Ω one obtains a lamination $L^2(\Omega)$ which consists only of finitely many F_N-orbits of leaves. Indeed, it is an easy exercise to show that any such *finite* lamination occurs precisely in this way. (For the less experienced reader we recommend for this exercise the transition to *symbolic laminations*, as described in detail in [6].)

Definition 2.4. An algebraic lamination $L^2 \neq \emptyset$ is called *minimal* if it doesn't contain any proper sublamination (other than $\emptyset$). This is equivalent to requiring that for any $(X, Y) \in L^2$ the set $F_N \cdot (X, Y) \cup F_N \cdot (Y, X)$ is dense in L^2.

The following terminology is inspired by geodesic laminations on surfaces, and isolated leaves which cross diagonally through one of the complementary components.

Definition-Remark 2.5. Let S be a subset of $\partial^2 F_N$.

(a) The *diagonal extension* $\mathrm{diag}(S)$ of S is the set of all $(X, Y) \in \partial^2 F_N$ such that for some integer $m \geq 1$ there exist elements $X = Z_0, Z_1, \ldots, Z_m = Y$ in ∂F_N such that $(Z_{i-1}, Z_i) \in S$ for $i = 1, \ldots, m$.

(b) It is easy to see that $S \subseteq \mathrm{diag}(S)$ and that $\mathrm{diag}(\mathrm{diag}(S)) = \mathrm{diag}(S)$.

(c) A subset $S \subseteq \partial^2 F_N$ is said to be *diagonally closed* if $S = \mathrm{diag}(S)$.

(d) If $S' \subseteq S \subseteq \partial^2 F_N$ then $\mathrm{diag}(S') \subseteq \mathrm{diag}(S)$.

(e) If S is F_N-invariant, then so is $\mathrm{diag}(L)$. Similarly, if S is flip-invariant then so is $\mathrm{diag}(S)$. However, if S is a closed subset of $\partial^2 S$, then we can deduce that $\mathrm{diag}(S)$ is closed only of $\mathrm{diag}(S) \smallsetminus S$ is finite.

(f) We denote by $\overline{\mathrm{diag}}(S)$ the closure of $\mathrm{diag}(S)$ in $\partial^2 F_N$. For any algebraic lamination L^2 the *diagonal closure* $\overline{\mathrm{diag}}(L^2)$ is again an algebraic lamination. If $\mathrm{diag}(L^2) \smallsetminus L^2$ is finite, then $\overline{\mathrm{diag}}(L^2) = \mathrm{diag}(L^2)$.

Remark 2.6. Examples of laminations where the diagonal extension is not closed are easy to come by. For example, if L^2 is given by a geodesic lamination on a surface, then $\mathrm{diag}(L^2)$ is closed if and only if all complementary components are simply connected.

If one picks a basis $\mathcal{A}$ for F_N, one can associate naturally to any pair $(X, Y) \in F_N$ a biinfinite reduced word $Z = X^{-1} \cdot Y$ in $\mathcal{A} \cup \mathcal{A}^{-1}$, and to Z the set $\mathcal{L}$ of all finite subwords. In this way one can associate to every algebraic lamination L^2 a *symbolic lamination* $L_{\mathcal{A}}$ and a laminary language $\mathcal{L}_{\mathcal{A}}$, and both of these translations are inversible. For details see [6].

Interpreting words in $\mathcal{A}^{\pm 1}$ as paths in the the associated rose $\rho_{\mathcal{A}}$, or in its universal covering, the Cayley graph $\Gamma(F_N, \mathcal{A})$, one can easily generalize the concepts of a "symbolic lamination" or a "laminary language" to more general graphs τ with an identification $F_N = \pi(\tau)$, or to their universal coverings. For the purposes of this paper, however, the following suffices:

Definition 2.7. Let τ be a graph with a marking $\theta : F_N \xrightarrow{\cong} \pi_1\tau$, and let $\partial\theta : \partial F_N \to \partial\tau$ be the induced homeomorphisms on the Gromov boundaries.

Let $(X, Y) \in \partial^2 F_N$, and let $\widetilde{\gamma}$ be the biinfinite reduced path (the "geodesic") in the universal covering $\widetilde{\tau}$ of τ which joins the boundary point $\partial\theta(X)$ to the boundary point $\partial\theta(Y)$.

The reduced biinfinite path γ in τ which is the image of $\widetilde{\gamma}$, under the universal covering map $\widetilde{\tau} \to \tau$, is called *the geodesic realization* (in τ) of the pair (X, Y), and is denoted by $\gamma_\tau(X, Y)$.

Assume now that the marked graph τ comes with a homotopy equivalence $f : \tau \to \tau$ which *represents* $\varphi \in \mathrm{Aut}(F_N)$, i.e. $\varphi \circ \theta = f_* \circ \varphi$ (up to inner automorphisms of $\pi_1(\tau)$). In this case it follows directly from the above definition that, for any integer $t \geq 0$ and any $(X, Y) \in \partial^2 F_N$, the geodesic realizations satisfy the equality

$$[f^t(\gamma_\tau(X, Y))] = \gamma_\tau(\partial\varphi^t(X), \partial\varphi^t(Y)) \,,$$

where for any (possibly non-reduced) path η we denote by $[\eta]$ the reduced path obtained from η by reduction.

We say that a lamination L^2 *is generated* by an infinite set of edge paths γ_i in a marked graph τ, if a pair $(X, Y) \in \partial^2 F_N$ belongs to L^2 if and only if any finite subpath of the geometric realization $\gamma_\tau(X, Y)$ is also a subpath of some γ_i or $\overline{\gamma}_i$.

It is easy to see (compare [6]) that any infinite set of edge paths in τ generates an algebraic lamination.

2.3. The algebraic lamination dual to an $\mathbb{R}$-tree.

To any tree $T \in \overline{\mathrm{cv}}_N$ in [7] there has been naturally associated a *dual algebraic lamination* (also called the *zero lamination*) $L^2(T)$ of T:

Definition 2.8. Let T be an $\mathbb{R}$-tree from $\overline{\mathrm{cv}}_N$.

(1) For any $\varepsilon > 0$ let $\Omega_\varepsilon(T)$ be the set of elements $w \in F_N$ with translation length $||w||_T \leq \varepsilon$, and let $L^2_\varepsilon(T) =: L^2(\Omega_\varepsilon(T))$.

(2) Define $L^2(T) := \bigcap_{\varepsilon>0} L^2_\varepsilon(T)$.

In [7] it has been shown that set $L^2(T)$ is an algebraic lamination on F_N, if $T \in \partial\mathrm{cv}_N$, and $L^2(T) = \emptyset$ otherwise (i.e. $T \in \mathrm{cv}_N$). It has also been shown in [7], for any basis $\mathcal{A}$ of F_N, that $L^2(T)$ is precisely the set of all $(X, Y) \in \partial^2 F_N$ such that the associated reduced biinfinite word Z in $\mathcal{A}^{\pm 1}$ has the property that for every finite subword w of Z and any $\varepsilon > 0$ there is a cyclically reduced word v which contains w as subword and satisfies $||v||_T \leq \varepsilon$.

As a slight extension of the latter, we can formulate the criterion which alternatively could serve as definition for this paper:

Remark 2.9. For any marked graph τ the following characterization of $L^2(T)$ holds:

A finite geodesic path γ' is a subpath of the geodesic realization $\gamma_\tau(X, Y)$ for some $(X, Y) \in L^2(T)$ if and only if for any $\varepsilon > 0$ there is an element $w \in F_N$ with translation length on T of seize $||w||_T \leq \varepsilon$, such that the conjugacy class of w in F_N is represented by a geodesic loop $\widehat{\gamma}$ which contains γ' as subpath.

It is well known that every $T \in \partial\mathrm{cv}_N$ either contains points $x \in T$ with non-trivial stabilizer $\mathrm{stab}(x) \subset F_N$, or else T has *dense orbits*: This means that every orbit $F_N x$ of a point $x \in T$ is a dense subset of T. Note that this means that in particular the set of *branch points* (i.e. points with 3 or more complementary components) is dense in T. Of course, there are also trees $T \in \partial\mathrm{cv}_N$ which have both, non-trivial point stabilizers, and dense orbits.

For $T \in \overline{\mathrm{cv}}_N$ with dense F_N-orbits one can give an equivalent characterization of $L^2(T)$ to that given in Definition 2.8 above. For such a tree T there is a canconical F_N-equivariant map $Q : \partial F_N \to \widehat{T} = \overline{T} \cup \partial T$, constructed in [7], where ∂T is the metric completion of T. The precise nature of how the map Q is defined is not relevant for the present paper. However, it is proved in [7] that in this case for $(X, Y) \in \partial^2 F_N$ we have $(X, Y) \in L^2(T)$ if and only if $Q(X) = Q(Y)$. This fact immediately implies:

Proposition 2.10. *Let $T \in \overline{cv}_N$ have dense F_N-orbits. Then $L^2(T)$ is diagonally closed, that is diag$(L^2(T)) = L^2(T)$.*

2.4. **Iwip automorphisms.**

Definition 2.11. (a) An outer automorphism $\varphi \in \mathrm{Out}(F_N)$ is called *irreducible with irreducible powers (iwip)* if no positive power of φ preserves the conjugacy class of a proper free factor of F_N.

(b) An outer automorphism $\varphi \in \mathrm{Out}(F_N)$ is called *atoroidal* if it has no non-trivial periodic conjugacy classes, i.e. if there do not exist $t \geq 1$ and $w \in F_N - \{1\}$ such that φ^t fixes the conjugacy class $[w]$ of w in F_N.

It was proved in [3] that for $N \geq 2$ an iwip automorphisms $\varphi \in \mathrm{Out}(F_N)$ is non-atoroidal if and only if φ is induced, via an identification of F_N with the fundamental group of a compact surface S with a single boundary component, by a pseudo-Anosov homeomorphisms $h : S \to S$.

Remark 2.12. The terminology "iwip" derives from the groundbreaking paper [3]: Bestvina-Handel call an element $\varphi \in \mathrm{Out}(F_N)$ is *reducible* if there exists a free product decomposition $F_N = C_1 * \dots C_k * F'$, where $k \geq 1$ and $C_i \neq \{1\}$, such that φ permutes the conjugacy classes of subgroups $C_1, \dots, C_k$ in F_N. An element $\varphi \in \mathrm{Out}(F_N)$ is called *irreducible* if it is not reducible.

It is not hard to see that an element $\varphi \in \mathrm{Out}(F_N)$ is an iwip if and only if for every $n \geq 1$ the power φ^n is irreducible (sometimes such automorphisms are also called *fully irreducible*).

It is known by a result of Levitt and Lustig [17] that iwips have a simple "North-South" dynamics on the compactified Outer space $\overline{\mathrm{CV}}_N$:

Proposition 2.13. [17] *Let $\varphi \in Out(F_N)$ be an iwip. Then there exist unique $[T_+] = [T_+(\varphi)], [T_-] = [T_-(\varphi)] \in \overline{CV}_N$ with the following properties:*

1. *The elements $[T_+], [T_-] \in \overline{CV}_N$ are the only fixed points of φ in $\overline{CV}_N$.*
2. *For any $[T] \in \overline{CV}_N$, $[T] \neq [T_-]$ we have $\lim_{n\to\infty}[T\varphi^n] = [T_+]$ and for any $[T] \in \overline{CV}_N$, $[T] \neq [T_+]$ we have $\lim_{n\to\infty}[T\varphi^{-n}] = [T_-]$.*
3. *We have $T_+\varphi = \lambda_+ T$ and $T_-\varphi^{-1} = \lambda_- T_-$ where $\lambda_+ > 1$ and $\lambda_- > 1$.*
4. *Both, T_+ and T_- are "intrinsically non-simplicial" $\mathbb{R}$-trees: Every F_N-orbit of points is dense in the tree.*

In [17] it is also proved that convergence in (2) is locally uniform and hence uniform on compact subsets. For more information see [15]. A description of T_+ in terms of train tracks is given below in section 3.7.

3. Train track technology

This section will be a reference section, organized by subheaders as a kind of glossary. Almost everything in this section is known, or within ε-neighborhood of known facts. It is only for the convenience of the reader that we assemble in this section what is needed later from basic train track theory.

The expert reader is encouraged to completely skip this section and only go back to it if he or she needs an additional explanation in the course of reading the later sections.

Notation 3.1. In order to avoid confusion, in this section the automorphisms in question will be denoted by α, as in the later sections we will have to consider both cases, $\alpha = \varphi$ and $\alpha = \varphi^{-1}$. Similarly, a train track map will be called $f : \tau \to \tau$ (to be specified later to $f_+ : \tau_+ \to \tau_+$ or $f_- : \tau_- \to \tau_-$, the forward limit tree will be simply called T (rather than T_+ or T_-), etc.

3.1. **Graphs, paths and graph maps.**

In this paper τ will always denote a graph, specified as follows:

Convention 3.2. By a *graph* τ we mean a non-empty topological space which is equipped with a cell structure, consisting of vertices and edges. Furthermore, τ satisfies the following conditions, unless explicitly otherwise stated:

- τ is connected.
- τ is a finite graph, i.e. it consists of finitely many vertices and edges.

  (Of course, the universal covering $\widetilde{\tau}$ of a finite graph τ is in general not finite, but this counts as "explicitly stated exception".)
- There are no vertices of valence 1 in τ (but we do admit vertices of valence 2).

We systematically denote vertices of τ by v, v' or v_i, and edges by e, e', e_j etc. A point which can be a vertex or an interior point of an edge is usually denoted by P or Q.

An edge e of τ is always oriented, and we denote by $\overline{e}$ the conversely oriented edge. In an *edge path* $\ldots e_k e_{k+1} e_{k+2} \ldots$ we always write the edges e_i in the direction in which the path runs. We occasionally use $\mathrm{Edges}(\tau)$ to denote the *edge set* of τ. We use the convention that its elements are, as before, *oriented* edges, so that the edge set $\mathrm{Edges}(\tau)$ contains for each pair $e, \overline{e}$ only the element e. If need be, we use the notation $\overline{\text{Edges}}(\tau) := \{\overline{e} \mid e \in \text{Edges}(\tau)\}$. Note that one is free, at any convenient time, to reorient edges: this operation does not change the graph, in our understanding, only the extra information about notation of edges.

However: to be specific: if we say "an edge of τ", this may well be the edge $\overline{e}$, for $e \in \text{Edges}(\tau)$.

We will denote the *simplicial length* of an edge path γ, i.e. the number of edges traversed by γ, with $L(\gamma)$. For the edge path $\overline{\gamma}$, by which we mean the path γ with reversed orientation, one has of course $L(\overline{\gamma}) = L(\gamma)$. In particular, for any single edge e one has $L(e) = L(\overline{e}) = 1$.

Convention-Definition 3.3. (a) If we use the terminology *path* without further specification, we always mean a finite path. The reader is free to formalize paths according to his or her own preferences; in most cases we recommend the viewpoint where a path γ means an *edge path* $\gamma = e_1 \circ e_2 \circ \ldots \circ e_q$, where the e_i are edges of τ such that the terminal vertex of any e_i coincides with the initial vertex of e_{i+1}. According to circumstances, we generalize this concept sometimes slightly by allowing that e_1 (or e_q) is a non-degenerate terminal (or initial) segment of an edge (or, if $q = 1$, any segment of a single edge).

A path γ is called *trivial* or *degenerate* if it consists of a single point. Otherwise, γ is said to be *non-trivial* or *non-degenerate*.

We will also sometimes suppress the distinction between an edge (or edge segment) e and a path which traverses precisely e.

(b) A path γ' is called a *backtracking path* if its endpoints coincide, and if the issuing loop is contractible in τ. If γ' occurs as subpath of a path γ, then γ' is called a *backtracking subpath* of γ.

(c) A (possibly infinite or biinfinite) path γ is *reduced* if it doesn't contain any non-degenerate backtracking subpath. We denote by $[\gamma]$ the reduced path obtained from a possibly non-reduced path γ by iteratively contracting all non-degenerate backtracking subpaths. (It is well known that $[\gamma]$ depends only on γ and not on the sequence of iteratively contracted backtracking subpaths).

(d) Two subpaths γ_1 and γ_2 of a path γ are *disjoint* if in the course of traversing γ one first traverses γ_1 completely without starting γ_2, or conversely. Similarly we define the *overlap* of the two subpaths as the maximal subpath of γ which is also a subpath of both γ_i.

Recall from the beginning of section 2 that a finite connected graph τ is called a *marked graph*, if it is equipped with a *marking isomorphism* $\theta : F_N \overset{\cong}{\longrightarrow} \pi_1(\tau)$. Recall also that we purposefully suppress the issue of choosing a basepoint of τ, as we are only interested in automorphims of F_N up to inner automorphisms.

Convention-Definition 3.4. (a) A *map* $f : \tau \to \tau'$ between graphs τ and τ' will always map vertices to vertices and edges to edge paths, which a priori may be non-reduced.

(b) A self-map $f : \tau \to \tau$ of a graph τ, provided with a marking isomorphism $\theta : F_N \to \pi_1(\tau)$, *represents* an automorphism α of F_N if the induced automorphisms $f_* : \pi_1(\tau) \to \pi_1(\tau)$ satisfies $\theta \circ \alpha = f_* \circ \theta$ up to inner automorphisms.

Remark 3.5. It is well known [5] that for any map $f : \tau \to \tau'$ between graphs τ and τ', which induces an isomorphism on $\pi_1(\tau)$ (or, for the matter, a monomorphism), there is an upper bound to the length of any backtracking path γ' which is contained as subpath in the (non-reduced) image $f(\gamma)$ of a reduced path γ in τ.

Definition 3.6. (a) A self-map $f : \tau \to \tau$ is called *expanding* if for every edge e of τ there is an exponent $t \geq 1$ such that the edge path $f^t(e)$ has length $L(f^t(e)) \geq 2$.

(b) If $f : \tau \to \tau$ is expanding, then there is a well defined self-map Df on the set $\text{Edges}(\tau) \cup \overline{\text{Edges}}(\tau)$ which associates to every edge e the initial edge of the edge path $f(e)$.

Remark 3.7. If a self-map $f : \tau \to \tau$ represents an iwip automorphism of F_N, then the hypothesis that f be expanding is always easy to satisfy: It suffices to contract all edges which are not expanded by any iterate f^t to an edge path of length ≥ 2: The issuing contracted subgraph must be a forest, as otherwise some f^t will fix (up to conjugacy) a non-trivial proper free factor of $\pi_1(\tau)$.

Alternatively, one can reparamatrize the map f along each edge so that it becomes a true homothety with respect to the Perron-Frobenius metric introcuded in subsection 3.7 below.

Since in this paper we will concentrate on self-maps of graphs that represent iwip automorphisms, we will almost exclusively consider self-maps $f : \tau \to \tau$ that are expanding.

Lemma 3.8. *Let the self-map $f : \tau \to \tau$ of the graph τ be expanding, and assume that f represents an iwip automorphism of F_N. Then we have:*

(1) *Every f-invariant subgraph $\tau_0 \subset \tau$ is either a single vertex (or a collection of vertices, if one relaxes the convention that a graph is connected), or equal to all of τ.*

(2) *For every edge e there exists an exponent $t \geq 0$ such that the (possibly non-reduced) edge path $f^t(e)$ crosses three or more times over e or $\overline{e}$.*

Proof. (1) Assume that $\tau_0 \subset \tau$ contains at least one edge. From the assumptions that f is expanding, and that $f(\tau_0) \subset \tau_0$, it follows that $\pi_1(\tau_0)$ is non-trivial. From the iwip hypothesis it follows that the inclusion $\tau_0 \subset \tau$ induces an isomorphisms on π_1. Hence our convention on graphs (see subsection 3.1), that τ is finite and does not contain vertices of valence 1, implies $\tau_0 = \tau$.

(2) By the expansiveness of f there is at least one edge e with the property described in statement (2). We consider a subgraph τ_0 which (i) contains at least one such edge, (ii) is f-invariant, and (iii) is minimal with respect to the properties (i) and (ii). If τ_0 contains edges which do not satisfy statement (2), then the union of those must be a non-empty subcomplex which is f-invariant (since any edge e, which lies outside of this union but is contained in its image, must have all image paths $f^t(e)$ disjoint from this union, thus contradicting the above minimality assumption (iii)). But this contradicts the expansiveness of f. Thus τ_0 consists only of edges satisfying statement (2). From statement (1) it follows $\tau_0 = \tau$. □

3.2. **Gates, turns, train tracks.**

Definition 3.9. (a) For any expanding self-map of graphs $f : \tau \to \tau$ and any vertex v of τ one partitions the edges with initial vertex v into equivalence classes, called *gates* $\mathfrak{g}_i$, by the following rule: Two edges e_1 and e_2, both with initial vertex v, belong to a common gate $\mathfrak{g}_i$ if and only if there is an exponent $t \geq 1$, such that $Df^t(e_1) = Df^t(e_2)$. To be specific, in case of a *loop edge* e at v (i.e. the initial and the terminal vertex of e both coincide with v) the edges e and $\overline{e}$ count as distinct edges with initial vertex v, which can or cannot belong to the same gate at v.

(b) This definition is sometimes extended to points P in the interior of an edge e, at which there are precisely two gates, one containing precisely the terminal segment e'' of e which starts at P, and the other one containing precisely the edge segment $\overline{e}'$, where e' is the initial segment of e which terminates at P.

Remark 3.10. For any vertex (or point P) of τ the map f induces via the map Df a map $f^P_{\mathfrak{G}}$ from the gates at P to the gates at $f(P)$. It follows directly from Definition 3.9 (a) that this map $f^P_{\mathfrak{G}}$ is injective, and hence, if P is a periodic point, that $f^P_{\mathfrak{G}}$ is bijective.

Definition 3.11. (a) A pair of edges of τ forms a *turn* (e, e') at a vertex v if and only if both, e and e', have v as initial vertex. The turn is said to be *degenerated* if $e = e'$. Otherwise it is called *non-degenerate*.

(b) A path γ *crosses over a turn* (e, e') (or *contains the turn* (e, e')) if γ contains the edge path $\overline{e} \circ e'$ as subpath (or a path $\overline{e}_0 \circ e'_0$ for non-degenerate initial segments e_0 of e and e'_0 of e').

Note that a path γ is reduced if and only of it doesn't cross over any degenerate turn.

Definition 3.12. Let $f : \tau \to \tau$ be an expanding self-map of a graph τ.

(1) The map f induces canonically a map $D^2 f$ on turns, by setting $D^2((e, e')) = (Df(e), Df(e'))$.
(2) A turn (e, e') is called *legal* if for all $t \geq 0$ the image turn $D^2 f^t((e, e'))$ is non-degenerate. Otherwise (e, e') is called *illegal*. In particular, all degenerate turns are illegal.
   A path (or loop) γ in τ is called *legal* if it crosses only over legal turns. Otherwise it is called *illegal*.
   (From Definition 3.9 it follows directly that the turn (e, e') is legal if and only if e and e' belong to distinct gates at their common initial vertex.)
(3) The map f is called a *train track map* if for every edge e of τ the edge path $f(e)$ is legal. (It is easy to see that this condition is equivalent to the requirement that f maps any legal path γ to a legal path $f(\gamma)$.)

Remark 3.13. (a) It follows directly from the Definition 3.12 (2) that $D^2 f$ maps legal turns to legal turns, and illegal turns to illegal turns.

(b) If $f : \tau \to \tau$ is not expanding, then one can use the following definition of *legal paths*, which in the expanding case is equivalent to the one given in Definition 3.12 (2):

A path γ in τ is said to be legal if the maps f^t are locally injective along γ, for all $t \geq 1$.

(c) In some circumstances it can be useful to consider more general "train track maps", i.e. self-maps of graphs which have the *train track property* that legal maps are mapped to legal paths (where "legal paths" are defined as in part (b) above). However, in this paper we insist on the assumption that a train track map $f : \tau \to \tau$ is always expanding !

We now state a crucial property of train track maps, which follows directly from the Definition 3.12:

Remark 3.14. For any train track map $f : \tau \to \tau$ and any path γ in τ the number of illegal turns which are crossed over by γ, denoted by $\mathrm{ILT}(\gamma)$, satisfies:

$$\mathrm{ILT}([f(\gamma)]) \leq \mathrm{ILT}(f(\gamma)) = \mathrm{ILT}(\gamma)$$

Of course, a path γ is legal if and only if $\mathrm{ILT}(\gamma) = 0$.

Lemma 3.15. *If $f : \tau \to \tau$ is a train track map which represents an iwip automorphism of F_N, then at every vertex of τ there are at least 2 gates.*

Proof. Let v be any vertex of τ, and let e be an edge with initial vertex v. Since f is a train track map, the path $f^t(e)$ must be legal, for any $t \geq 0$. By Lemma 3.8 (2) for some value of t the path $f^t(e)$ contains e (or $\overline{e}$) as interior edge. Let e' be the edge adjacent to this occurrence of e (or $\overline{e}$) on $f^t(e)$, so that $e'e$ or $\overline{e}\,\overline{e}'$ is a subpath of $f^t(e)$. By Definition 3.12 (2) the edges $\overline{e}'$ and e belong to distinct gates (at the vertex v). $\square$

The following is one of the fundamental results for automorphisms of free groups (slightly adapted to the language specified in this section):

Theorem 3.16 ([3])**.** *For every iwip automorphism α of F_N there exists a train track map $f : \tau \to \tau$ that represents α.*

3.3. **Eigenrays.**

Throughout this subsection we assume that $f : \tau \to \tau$ is a train track map.

Definition 3.17. (a) A *ray* ρ is an infinite path in τ which doesn't necessarily start at a vertex. The path ρ can alternatively be thought of as locally injective map $\mathbb{R}_{\geq 0} \to \tau$ which crosses infinitely often over vertices, or as reduced infinite edge path $e'_1 e_2 e_3 \ldots$, where the e_i with $i \geq 2$ are edges of τ, and e'_1 is a non-degenerate terminal segment of some edge e_1 of τ (which includes the possibility $e'_1 = e_1$). The point $P \in \tau$ (not necessarily a vertex !) which is the inital point of e'_1 (from the second viewpoint) or the image of $0 \in \mathbb{R}_{\geq 0}$ (from the first viewpoint) is called the *starting point* of ρ. If P is a vertex, then we also call it the *initial vertex* of ρ.

(b) A ray ρ is called an *eigenray* if one has $f^t(\rho) = \rho$ for some integer $t \geq 1$. (Note that the condition $f^t(\rho) = \rho$ is stronger than requiring just $[f^t(\rho)] = \rho$!) In particular, it follows that the starting point P of ρ is f-periodic, and that ρ is a legal path.

Lemma 3.18. *Let γ be a non-trivial finite path in τ, and assume that for some $t \geq 1$ the map f^t maps γ to a path $f^t(\gamma)$ which contains γ as initial subpath. Then there is precisely one eigenray ρ in τ which has γ as initial subpath.*

Proof. It is easy to see that the union of the $f^{kt}(\gamma)$ for all $k \geq 1$ form an eigenray. Conversely, any eigenray ρ which contains γ as initial subpath must also contain any $f^{kt}(\gamma)$ as initial subpath. □

We see from Lemma 3.18 that an eigenray can never bifurcate in the forward direction to give rise two eigenrays with same initial point. However, we will see later that, in the negative direction, an eigenray may well bifurcate, giving rise to an INP (compare Lemma 3.26).

Proposition 3.19. *Let $f : \tau \to \tau$ be a train track map, and let P be a periodic vertex (or interior point) of τ. Then every gate $\mathfrak{g}_i$ at P is mapped by f periodically. A gate $\mathfrak{g}_i$ contains precisely one edge e_i on which Df acts periodically (an "eigen edge"), and there is precisely one eigenray ρ which starts "from $\mathfrak{g}_i$", i.e. which starts with an edge (or edge segment) that belongs to $\mathfrak{g}$. This edge is precisely the eigen edge e_i.*

Proof. By Definition 3.9 for every gate $\mathfrak{g}$ there is an exponent $t \geq 0$ such that Df^t maps every edge of $\mathfrak{g}$ to a single edge. Recall from Remark 3.10 that for each periodic point P of τ the map f induces a periodic map $f^P_{\mathfrak{G}}$ on the gates at P. Hence each gate $\mathfrak{g}_i$ at P contains precisely one edge e_i (the "eigen edge") which is periodic under the induced map Df, say with period $k(i) \in \mathbb{N}$.

Thus $f^{k(i)}(e_i)$ is an edge path with initial subpath e_i. Hence Lemma 3.18 gives us a well defined eigenray which starts from the gate $\mathfrak{g}_i$, with initial edge e_i.

Since the initial edge (or edge segment) of an eigenray is necessarily periodic under the map Df, and two eigenrays with same initial edge (or edge segment) must agree (see Lemma 3.18), it follows that from every gate only one eigenray can start. □

From Proposition 3.19 we obtain directly the following:

Corollary 3.20. *For every periodic point P of τ there is a canonical bijection between (i) the gates at P, (ii) the eigen edges with initial vertex P, and (iii) the eigenrays with starting point P.* □

We also need later on the following property:

Lemma 3.21. *Let $f : \tau \to \tau$ be a train track map which represents an automorphism of F_N, and let ρ be an eigenray in τ. Then ρ can not be an eventually periodic path, i.e. a path of the form $\gamma_0 \circ \gamma_1 \circ \gamma_1 \circ \gamma_1 \circ \ldots$*

Proof. By our convention from Definition 3.12 (3) the map f is expanding, which contradicts the fact that the loop γ_1 would have to be mapped to itself, given that by hypothesis f represents an automorphisms of F_N. □

3.4. **INPs.**

As before, assume that throughout this subsection $f : \tau \to \tau$ is a train track map. In addition, we assume in this section that f induces an automorphism on $\pi_1(\tau)$.

Definition 3.22. A path η in τ which crosses over precisely one illegal turn is called a *periodic indivisible Nielsen path* (or *INP*, for short), if for some exponent $t \geq 1$ one has $[f^t(\eta)] = \eta$, (where $[\gamma]$ denotes as before the path obtained via reduction from a possibly unreduced path γ).

The illegal turn on $\eta = \gamma' \circ \overline{\gamma}$ is called the *tip* of η, while the two maximal initial legal subpaths γ' and γ, of η and $\overline{\eta}$ respectively, are called the *branches* of η.

Remark 3.23. Let $\eta = \gamma' \circ \overline{\gamma}$ be an INP of the train track map $f : \tau \to \tau$.

(1) The two endpoints of η are fixed points (but not necessarily vertices !) of f^t for some $t \geq 1$. The branches γ and γ' of η are initial segments of eigenrays ρ and ρ' respectively, defined as unions of the nested sequences of paths $(f^{kt}(\gamma))_{k\in\mathbb{N}}$ and $(f^{kt}(\gamma'))_{k\in\mathbb{N}}$ (compare Lemma 3.18). The rays ρ and ρ' coincide up to the initial segments γ and γ' respectively: indeed the common terminal segment of ρ and ρ' is precisely the union of all the paths crossed back and forth by the backtracking subpaths (see Convention-Definition 3.3 (b)) of the unreduced paths $f^{kt}(\eta)$, at the tip of the illegal turn of η, for all $k \geq 1$.

(2) Two eigenrays ρ and ρ', which coincide up to initial segments γ and γ' respectively, define always an INP $\eta = \gamma' \circ \overline{\gamma}$ as above in (1). This follows directly from the definitions.

Remark 3.24. For every train track map $f : \tau \to \tau$ there are only finitely many INPs in τ, if one assumes (as done throughout this section) that (i) f is expanding, (ii) τ is finite, and (iii) f induces an automorphism on $\pi_1(\tau)$.

This is a consequence of the fact that f induces a quasi-isometry on the universal covering space (with respect to the simplicial metric L), so that for any geodesic path in τ the length of any backtracking subpath in the image path is bounded above by a constant depending only on f. Since f is expanding, this gives a bound to the maximal length of the legal branches γ and γ' of any INP $\eta = \gamma' \circ \overline{\gamma}$ in τ, since the only backtracking subpath in $f(\eta)$ is the subpath cancelled at the tip when $f(\eta)$ is reduced to $[f(\eta)]$.

(If the reader wants to fill in the details, we recommend restricting to the case where f induces an iwip automorphism on $\pi_1(\tau)$ and working with the PF-metric L^{PF} introduced in subsection 3.7 rather than with the simplicial metric L; by Remark 3.39 the two metrics define a quasi-isometry for $\widetilde{\tau}$.)

Convention 3.25. As pointed out in Remark 3.23 (1) the endpoints of an INP may a priori well lie in the interior of an edge. However, since these points are f-periodic, and since by Remark 3.24 there are only finitely many INPs in τ, we can assume from now on that the edges of τ have been subdivided accordingly, so that all endpoints of INPs are vertices.

(Of course, such a subdivision must be followed potentially by the procedure lined out in Remark 3.7, in order to make sure that after subdivision the train track map f is still expanding.)

Lemma 3.26. *There exists an exponent $r_1 \geq 1$ with the following property: Let γ be a path in τ, and assume that it contains precisely two illegal turns, each being the tip of an INP-subpath η_1 and η_2 of γ. Assume that η_1 and η_2 overlap in a non-degenerate subpath. Then $f^{r_1}(\gamma)$ reduces to a path $[f^{r_1}(\gamma)]$ which is legal.*

Proof. By hypothesis, $\eta_1 = \gamma_1' \circ \overline{\gamma}_1$ and $\eta_2 = \gamma_2' \circ \overline{\gamma}_2$ intersect in a non-degenerate path, which by Convention 3.25 must be an edge path γ' with vertices as initial and terminal point.

Each INP contains only one illegal turn, and by assumption γ contains two. Hence the path γ' cannot contain either, so that it must be a legal subpath of both, the branch γ_2' of η_2, and the inverse of the branch γ_1 of η_1.

Since by assumption γ' is non-degenerate, it is expanded under iteration of f to a legal edge path of arbitrary big length.

Assume that an exponent $r \geq 0$ is chosen such that $L(f^r(\gamma'))$ is strictly bigger than the sum $L(\gamma_1) + L(\gamma_2')$. In this case it follows directly that, in reducing $f^r(\gamma)$, both illegal turns disappear into the backtracking subpaths, so that the resulting path $[f^r(\gamma)]$ is legal.

Since f is expanding and since (see Remark 3.24) there are only finitely many INPs in τ, it is easy to find a bound r_1 as required in the claim. $\square$

The following is one of the crucial properties of train track maps of graphs. It goes back to the first paper [3] on the subject. An alternative proof is given in [18].

Proposition 3.27. *For every train track map $f : \tau \to \tau$ there exists a constant $r_2 = r_2(\tau) \geq 0$ such that every path γ with precisely 1 illegal turn satisfies:*

Either γ contains an INP as subpath, or else $[f^{r_2}(\gamma)]$ is legal. □

Recall from Remark 3.14 that $\mathrm{ILT}(\gamma)$ denotes the number of illegal turns in a path γ.

Proposition 3.28. *There exists an exponent $r = r(f) \geq 0$ such that every finite path γ in τ with $ILT(\gamma) \geq 2$ satisfies*

$$ILT([f^r(\gamma)]) < ILT(\gamma)\,,$$

unless every illegal turn on γ is the tip of an INP-subpath η_i of γ, where any two η_i are either disjoint subpaths on γ, or they overlap precisely in a common endpoint (for this terminology compare Convention-Definition 3.3 (d)).

Proof. It suffices to take r to be the maximum of r_1 from Lemma 3.26 and r_2 from Proposition 3.27. We then apply the latter to any maximal subpath with precisely one illegal turn. If each such subpath contains an INP-subpath, we apply Lemma 3.26. □

Bestvina-Handel [3] proved the following result which is very useful in many contexts:

Proposition 3.29. *Every iwip automorphism φ is represented by a train track map $f : \tau \to \tau$ such that τ contains at most one INP η. The path η is closed if and only if φ is induced by a homeomorphism of some surface with at least one boundary component or puncture.*

Such special train track maps have been termed *stable* in [3].

3.5. **Used turns.**

As before, let $f : \tau \to \tau$ be a train track map which is fixed throughout this subsection.

Definition 3.30. A turn (e, e') in τ is called *used* if there is an edge e'' in τ with the property that for some $t \geq 1$ the image path $f^t(e'')$ crosses over this turn. Otherwise the turn (e, e') is called *unused*.

Remark 3.31. From Definition 3.30 we derive directly the following facts:

(1) Every used turn is legal. The converse is in general wrong.
(2) The image turn (under the map $D^2 f$) of any used turn is also used.
(3) The image of an unused turn (e, e') may well be used. There is, however, a constant $s \geq 0$ (which only depends on the total number of turns in τ) such that either $D^2 f^s(e, e')$ is used, or else all forward iterates of (e, e') under $D^2 f$ are unused.

Lemma 3.32. *If the train track map $f : \tau \to \tau$ represents an iwip automorphism, then for every edge e of τ there is an edge e' which has the same initial vertex as e, such that $\overline{e}e'$ is a used legal turn.*

Proof. This is a direct consequence of Lemma 3.8 (2). $\square$

Lemma 3.33. *Let $\mathfrak{g}_1$ and $\mathfrak{g}_2$ be two gates at the same periodic vertex v of τ, and assume that some turn (e'_1, e'_2) with $e'_i \in \mathfrak{g}_i$ is used. Then the turn (e_1, e_2) is also used, where e_i is the initial edge of the eigenray ρ_i that starts from the gate $\mathfrak{g}_i$, for $i = 1$ and $i = 2$. (By Proposition 3.19 this is equivalent to stating that each edge e_i is the eigen edge of the gate $\mathfrak{g}_i$.)*

Proof. From the definition of a gate it follows directly that for every periodic gate $\mathfrak{g}_i$ there is an exponent $t_i \geq 1$ such that Df^{t_i} maps every edge e' in $\mathfrak{g}_i$ to the eigen edge e_i of $\mathfrak{g}_i$.

Now, if the turn (e'_1, e'_2) is used, then for some edge e and some $t \geq 1$ the edge path $f^t(e)$ crosses over the turn (e'_1, e'_2). It follows that the path $f^{tt_1t_2}(e)$ crosses over the turn (e_1, e_2). $\square$

3.6. BFH's "stable" lamination.

The initials BFH in the follow definition refer to Bestvina-Feighn-Handel, who introduced and studied the following lamination in [2]; in particular they showed a special attraction property of this lamination, for train track maps which represent iwip automorphisms.

Definition 3.34. For any train track map $f : \tau \to \tau$ we define the *BFH-attracting lamination* $L^2_{BFH}(f)$ as the lamination which is generated by the family of paths $f^t(e)$, for any edge e of τ and any exponent $t \geq 0$.

It is easily seen that, if f represents an automorphism α of F_N, that L^2_{BFH} is invariant under the natural map induced by α on $\partial^2 F_N$.

The following is an easy consequence of the definition of "used turn" in the previous section, and of Lemma 3.32.

Lemma 3.35. *Let $f : \tau \to \tau$ be a train track map which represents an iwip automorphism. Then any leaf (X, Y) of the lamination $L^2_{BFH}(f)$ is represented by biinfinite path γ in τ which only crosses over used turns.* □

The following proposition is due to Bestvina-Feighn-Handel. We note here that an alternative proof of part (1) follows as direct consequence of our Theorem 1.2, since $L^2(T_+)$ depends only on α and not on the particular train track representative $f : \tau \to \tau$.

We would also recommend to the reader to try proving part (2) as an exercise, as it is quite doable and will enhance the reader's understanding of the main subjects treated in this paper.

Proposition 3.36 ([2])**.** *Let $f : \tau \to \tau$ be a train track representative of an iwip automorphism α of F_N.*

(1) The lamination $L^2_{BFH}(f)$ depends only on $\alpha \in Out(F_N)$ and not on the particular choice of the train track representative f.

(2) The lamination$L^2_{BFH}(f)$ is minimal (see Definition 2.4). □

3.7. **The limit tree and BBT.**

In this subsection we freely use the terminology and notation about $\mathbb{R}$-trees introduced previously in section 2.

Let T be the forward limit tree of the iwip automorphism α (i.e. $T\alpha = \lambda T$ for $\lambda := \lambda_+(\alpha) > 1$). Let $f : \tau \to \tau$ be any train track representative of α.

It is well known (see [10]) that for the universal covering $\widetilde{\tau}$ of τ there exists an F_N-equivariant map $i : \widetilde{\tau} \to T$ which is injective on legal paths (where a turn in $\widetilde{\tau}$ is legal if and only if its image turn in τ is legal). If one uses this map to pull back the metric on T to define an edge length $L^{PF}(e)$ for every edge e of $\widetilde{\tau}$ (an thus, by the F_N-equivariance of i) also for the edges of τ, then one obtains a *Perron-Frobenius column eigenvector* $\vec{v} = (L^{PF}(e))_{e \in \mathrm{Edges}(\tau)}$ (with eigenvalue λ) of the non-negative primitive transition matrix $M(f)$, see for example [19]. One obtains:

Remark 3.37. The map $i : \widetilde{\tau} \to T$ has the property that, with respect to the PF-metric L^{PF}, every finite legal path γ in $\widetilde{\tau}$ is mapped isometrically to the geodesic segment $i(\gamma)$.

The map i satisfies the *bounded backtracking property (BBT)*: There is a constant $BBT(i) \geq 0$ such that for any two points $x, y \in \widetilde{\tau}$ the geodesic segment $[x, y] \subset \widetilde{\tau}$ is mapped by i into the $BBT(i)$-neighborhood of the geodesic segment $[i(x), i(y)]$:

$$i([x,y]) \subset \mathcal{N}_{BBT(i)}([i(x), i(y)])$$

Based on results of [2] it has been shown in [10] that:

$$BBT(i) \leq \mathrm{vol}^{PF}(\tau) := \sum_{e \in \mathrm{Edges}(\tau)} L^{PF}(e)$$

Lemma 3.38. *Consider any element $w \in F_N$, with translation length $0 \leq ||w||_T < c$ on T, where c is smallest PF-length of any loop in τ, and let $\widehat{\gamma}$ be a reduced loop in τ which represents the conjugacy class $[w] \subset F_N$. Let γ' be a legal subpath of $\widehat{\gamma}$. The PF-length of γ' satisfies:*

$$L^{PF}(\gamma') \leq ||w||_T + 4BBT(i)$$

(Indeed, one can also show the stronger inequality $L^{PF}(\gamma') \leq ||w||_T + 2\ BBT(i)$)

Proof. We can lift $\widehat{\gamma}$ to a biinfinite geodesic $\widetilde{\gamma}$ in the universal covering $\widetilde{\tau}$, which is mapped by $i : \widetilde{\tau} \to T$ to a (possibly non-reduced) path $i(\widetilde{\gamma})$ in T that covers the axis $\mathrm{Ax}(w)$ in T. The action of w on $\widetilde{\tau}$ and on T fixes $\widetilde{\gamma}$ and $\mathrm{Ax}(w)$ and translates each of them (by the amount of $||w||_T$, for $\mathrm{Ax}(w)$).

From the assumption $||w||_T < c$ it follows (by Remark 3.37) that $\widehat{\gamma}$ is not a legal loop, so that in particular the legal subpath γ' can not wrap around all of $\widehat{\gamma}$. Let P be a point on $\widehat{\gamma}$ which is not contained in γ', and let $\widetilde{P}$ be a lift of P to $\widetilde{\gamma}$. Hence the point $P_T := i(\widetilde{P})$ lies in the BBT(i)-neighborhood of $\mathrm{Ax}(w)$. In particular it follows that $d(P_T, wP_T) \leq ||w||_T + 2\mathrm{BBT}(i)$.

We can now lift γ' to a legal path $\widetilde{\gamma}'$ which is contained in the geodesic segment $[\widetilde{P}, w\widetilde{P}] \subset \widetilde{\gamma}$. Hence the endpoints of $i(\widetilde{\gamma}')$ must lie in the BBT(i)-neighborhood of $[P_T, wP_T]$. Thus they have distance at most $||w||_T + 4\mathrm{BBT}(i)$. But $\widetilde{\gamma}'$ is legal and hence (Lemma 3.37) mapped isometrically to its i-image. Thus $\widetilde{\gamma}'$ and hence also γ' have PF-length bounded above by $||w||_T + 4\mathrm{BBT}(i)$. ∎

Remark 3.39. (1) Note that with respect to the Perron-Frobenius length L^{PF} any legal path γ satisfies $L^{PF}(f(\gamma)) = \lambda L^{PF}(\gamma)$. In particular, it follows that no edge e can have $L^{PF}(e)$ equal to 0, since otherwise the subgraph $\tau_0 \subset \tau$ defined by all such edges would contradict Lemma 3.8 (1).

(2) As a consequence we observe that the simplicial edge length L and the Perron-Frobenius edge length L^{PF} give rise to quasi-isometric metrics on the universal covering $\widetilde{\tau}$.

4. Steps 1 and 2

The material in this section is reminiscent to some of the techniques used previously by the second author in [20].

Let T_+ be the forward limit tree of the atoroidal iwip automorphism φ (i.e. $T_+\varphi = \lambda_+ T$ for $\lambda_+ > 1$). Let $f_+ : \tau_+ \to \tau_+$ be a train track representative of φ. By Proposition 3.29 we can assume the following:

Hypothesis 4.1. *There is at most one INP η in τ_+, and the two endpoints of η are distinct.*

Definition 4.2. For any constant $C \geq 1$ a (possibly infinite or bi-infinite) path γ in τ_+ is called *totally C-illegal* if every legal subpath γ' of γ has length

$$L(\gamma') \leq C\,.$$

Remark 4.3. The Definition 4.2 is motivated by Lemma 4.4 below. Another, rather useful property of any finite totally C-illegal path γ, which follows directly from the definition, is given by the second of following inequalities (while the first one doesn't require any hypotheses):

$$\mathrm{ILT}(\gamma) + 1 \leq L(\gamma) \leq C(\mathrm{ILT}(\gamma) + 1)$$

Lemma 4.4. *There exists a constant $C \geq 0$ such that for every pair (X, Y) in the dual lamination $L^2(T_+) \subset \partial^2 F_N$ the reduced biinfinite path $\gamma = \gamma_{\tau_+}(X, Y)$ in τ_+ (the "geodesic realization" of the pair (X, Y), see Definition 2.7) is totally C-illegal.*

Proof. By Remark 2.9 a finite geodesic path γ' is a subpath of the geodesic realization $\gamma_{\tau_+}(X, Y)$ for some $(X, Y) \in L^2(T_+)$ if and only if for every $\varepsilon > 0$ there is an element $w \in F_N$ with translation length on T_+ of seize $||w||_{T_+} \leq \varepsilon$, such that the conjugacy class of w in F_N is represented by a geodesic loop $\widehat{\gamma}$ which contains γ' as subpath. Hence the desired inequality is a direct consequence of Lemma 3.38, where the constant C can be calculated from the value $4\mathrm{BBT}(i)$ for the map $\widetilde{\tau}_+ \to T_+$ explained in subsection 3.7 and the quasi-isometry constants between the length functions L and L^{PF}, see Remark 3.39 (2). $\square$

Below we use the following terminology: If γ is a (possibly infinite or biinfinite) path in τ_+, and γ_1 a subpath of γ, then a boundary subpath γ' of $[f_+^t(\gamma_1)]$ (for some $t \geq 1$) is *cancelled by the reduction of $f_+^t(\gamma)$ to $[f_+^t(\gamma)]$* if, for $\gamma = \gamma_0 \circ \gamma_1 \circ \gamma_2$, when $[f_+^t(\gamma_0)] \circ [f_+^t(\gamma_1)] \circ [f_+^t(\gamma_2)]$ is reduced to give $[f_+^t(\gamma)]$, then γ' is entirely contained in one of backtracking subpaths at the concatenation points of the product path $[f_+^t(\gamma_0)] \circ [f_+^t(\gamma_1)] \circ [f_+^t(\gamma_2)]$.

Lemma 4.5. *Let $C \geq 1$ be any constant, and let $(X, Y) \in \partial^2 F_N$ be such that, for every integer $t \geq 0$, the geodesic realization in τ_+ of the pair $(\partial\varphi^t(X), \partial\varphi^t(Y))$ by a reduced biinfinite path $\gamma_t = \gamma_{\tau_+}(\partial\varphi^t(X), \partial\varphi^t(Y))$ is totally C-illegal. Then there exists an integer $s \geq 0$ which has the property that for any any finite subpath γ' of γ_0 with ILT$(\gamma') \geq 2$ one has either*

$$ILT([f_+^s(\gamma')]) < ILT(\gamma'),$$

or else a boundary subpath of $[f_+^s(\gamma')]$ which crosses over at least one illegal turn is completely cancelled by the reduction of $f_+^s(\gamma)$ to $[f_+^s(\gamma)] = \gamma_s$.

Proof. From Proposition 3.28 we know that for any subpath γ' of γ_0 with 2 or more illegal turns one has $\mathrm{ILT}([f_+^r(\gamma')]) < \mathrm{ILT}(\gamma')$ (for r as given in Proposition 3.28), unless every illegal turn of γ' is the tip of an INP entirely contained as subpath in γ', such that any two such INP-subpaths which are adjacent on γ' either (i) are disjoint, or (ii) intersect precisely in a common endpoint.

The second case (ii) is ruled out by our Hypothesis 4.1, since γ' is a geodesic and the only INP η contained in τ_+ is not closed.

In order to rule out the first case (i), we observe that the subpath γ'' of γ', which connects the two endpoints of the adjacent INP-subpaths of γ', has to run over at least one edge, and that γ'' is legal. Hence, by the expansiveness of f_+, for some $t \geq 1$ the path $[f_+^t(\gamma')]$ contains a legal subpath $f_+^t(\gamma'')$ of length bigger than the constant C. This contradicts the hypothesis that $[f_+^t(\gamma)] = \gamma_t$ is totally C-illegal, unless one of the two INP-subpaths of $[f_+^t(\gamma')]$ which are adjacent to the subpath $f_+^t(\gamma'')$ is completely contained in a boundary subpath of $[f_+^t(\gamma')]$ which is cancelled by the reduction of $f_+^t(\gamma)$ to $[f_+^t(\gamma)] = \gamma_t$.

Hence it suffices to take $s \geq r$ big enough so that $L(f^s(e)) > C$ for any edge e of τ_+. □

Corollary 4.6. *Let $C \geq 0, (X, Y) \in \partial^2 F_N, \gamma'$ and $s \geq 1$ be as in Lemma 4.5. Let us furthermore assume that ILT$(\gamma') \geq 5$. Then one has*

$$ILT([f_+^s(\gamma')]) \leq \frac{1}{4} ILT(\gamma')$$

Proof. By the assumption $\mathrm{ILT}(\gamma') \geq 5$ we can subdivide $\gamma' = \gamma_1 \circ \ldots \circ \gamma_q$ as concatenation of subpaths γ_i, where each γ_i satisfies $2 \leq \mathrm{ILT}(\gamma_i) \leq 3$, and where at each concatenation vertex the path γ' crosses over an illegal turn. We now apply Lemma 4.5 to each of the subpaths γ_i and obtain, for $s \geq 1$ as specified there, that either $\mathrm{ILT}([f_+^s(\gamma_i)] < \mathrm{ILT}(\gamma_i)$, or else one of the illegal turns in $[f_+^s(\gamma_i)]$ is cancelled when $f_+^s(\gamma')$ is reduced to $[f_+^s(\gamma')]$. Thus one obtains, for

$$[f_+^s(\gamma')] = \gamma'_1 \circ \ldots \circ \gamma'_q\,,$$

where each γ'_i is the (possibly trivial) subpath of $[f_+^s(\gamma_i)]$ which is left-over when the concatenation $[f_+^s(\gamma_1)] \circ \ldots \circ [f_+^s(\gamma_q]$ is reduced, that each of the paths γ'_i satisfies $\mathrm{ILT}(\gamma'_i) \leq \mathrm{ILT}(\gamma_i) - 1$. This gives $\mathrm{ILT}([f_+^s(\gamma')] < \frac{1}{4}\mathrm{ILT}(\gamma')$, as claimed. □

Proposition 4.7. *For any $\lambda > 1$ there is a constant $c > 0$ and an exponent $t \geq 0$ with the following property:*

Consider $(X, Y) \in L^2(T_+) \subset \partial^2 F_N$, and let $\gamma = \gamma_{\tau_+}(X, Y)$ be a geodesic realization as reduced biinfinite path in τ_+. Let γ' be a finite subpath of γ of simplicial length $L(\gamma') \geq c$. Then the reduced image path $[f_+^t(\gamma')]$ satisfies:

$$L([f_+^t(\gamma')]) \leq \frac{1}{\lambda} L(\gamma')$$

Proof. By Lemma 4.4 there exists a constant $C \geq 0$ such that for any $t \in \mathbb{Z}$ the geodesic realization in τ_+ of the pair $(\partial\varphi^t(X), \partial\varphi^t(Y)) \in L^2(T_+)$ by a reduced path $\gamma_t = \gamma_{\tau_+}(\partial\varphi^t(X), \partial\varphi^t(Y))$ is totally C-illegal. Hence any finite subpath γ' of γ_0 with sufficiently large $\mathrm{ILT}(\gamma')$ satisfies by Corollary 4.6 the inequality

$$\mathrm{ILT}([f_+^{ks}(\gamma')]) \leq \frac{1}{4^k}\mathrm{ILT}(\gamma')$$

for some large $k \geq 1$.

Now, one can decompose $[f_+^{ks}(\gamma')]$ as concatenation $[f_+^{ks}(\gamma')] = \delta_0 \circ [f_+^{ks}(\gamma')]_{\gamma_{ks}} \circ \delta_1$, where $[f_+^{ks}(\gamma')]_{\gamma_{ks}}$ is the maximal subpath of $[f_+^{ks}(\gamma')]$ which is also subpath of γ_{ks}, while δ_0 and δ_1 are boundary subpaths that are cancelled when $f_+^{ks}(\gamma_0)$ is reduced to $[f_+^{ks}(\gamma_0)] = \gamma_{ks}$. In particular, it follows that for each δ_i there is a backtracking subpath δ'_i of $f_+^{ks}(\gamma_0)$, such that δ_i is obtained from a subpath of δ'_i by canceling (in the process of reducing the subpath $f_+^{ks}(\gamma')$ of $f_+^{ks}(\gamma_0)$ to $[f_+^{ks}(\gamma')]$) certain further backtracking subpaths. Hence (compare Remark 3.5) $L(\delta'_i)$ and thus $L(\delta_i)$ is bounded above by a constant $K \geq 0$ which depends only on f_+^{ks}.

As subpath of the totally C-illegal path γ_{ks} the path $[f_+^{ks}(\gamma')]_{\gamma_{ks}}$ must be itself totally C-illegal. Hence we deduce, using Remark 4.3, that

$$\begin{array}{rcl} L([f_+^{ks}(\gamma')]) & \leq & L([f_+^{ks}(\gamma')]_{\gamma_{ks}}) + 2K \\ & \leq & C(\mathrm{ILT}([f_+^{ks}(\gamma')]_{\gamma_{ks}}) + 1) + 2K \\ & \leq & C(\mathrm{ILT}([f_+^{ks}(\gamma')]) + 1) + 2K \\ & \leq & C(\frac{1}{4^k}\mathrm{ILT}(\gamma')) + 1) + 2K \\ & \leq & C(\frac{1}{4^k}L(\gamma)) + 1) + 2K \end{array}$$

It is now easy to determine first k and thus t large enough (in dependence of the given value of λ) and then (using again the inequality $L(\gamma) \leq C(\mathrm{ILT}(\gamma)+1)$ from Lemma 4.3) a sufficiently large constant c which give the desired inequality. $\square$

Remark 4.8. It follows easily from standard arguments of geometric group theory that the strong uniform contraction property of the automorphism φ along the leaves of the dual lamination $L^2(T_+)$ of its forward limit tree T_+, which is stated in Proposition 4.7 above in terms of the train track representative f_+ of φ, is in fact an intrinsic property, which can be similarly expressed with respect to any representative of φ as a self-map of some marked graph, or, more algebraically, directly for the automorphisms φ expressed as traditionally by means of a basis $\mathcal{A}$ of F_N and the image words of the generators $a_i \in \mathcal{A}$. Any φ-invariant subset of $\partial^2 F_N$ which satisfies this property is called *uniformly expanding.*

It is easy to see that, passing over to φ^{-1}, there is a similar way to define the same sets by an analogous property as *uniformly* φ^{-1}-*expanding* subsets of $\partial^2 F_N$.

5. STEPS 3 AND 4

Throughout this section we assume that $f_- : \tau_- \to \tau_-$ is a train track map that represents the automorphism φ^{-1} of F_N. Since φ is assumed to be iwip, it follows (as immediate consequence of Definition 2.11) that φ^{-1} is also iwip. Just as in the previous section for $f_+ : \tau_+ \to \tau_+$, we can assume here:

Hypothesis 5.1. *There is at most one INP η in τ_-, and the two endpoints of η are distinct.*

We consider the lamination $L^2(T_+)$ from the previous section, and we shorten here the notation to $L^2_- := L^2(T_+)$. From the previous section (see Remark 4.8) we deduce:

Lemma 5.2. *For any $\lambda > 1$ there is a constant $c > 0$ and an exponent $t \geq 1$ such that for any pair $(X, Y) \in L^2_-$ and any subpath γ' of the geodesic realization $\gamma = \gamma_{\tau_-}(X, Y)$ of (X, Y) in τ_-, with simplicial length $L(\gamma') \geq c$, one has:*

$$L(f_-^t(\gamma')) \geq \lambda L(\gamma')$$

$\square$

For any $K \geq 0$ we define the set $\mathcal{L}_K = \mathcal{L}_K(L^2_-, \tau_-)$ to be the set of subpaths γ of length $L(\gamma) = K$ of any biinfinite path $\gamma_{\tau_-}(X, Y)$ in τ_- which is a geodesic realization any pair $(X, Y) \in L^2_-$.

For any constant $C \geq 0$ and any path γ in τ_- we denote by $\gamma\dagger_C$ the (possibly trivial) subpath of γ which is left after erasing from γ the two boundary subpaths of length C.

Recall from Remark 3.5 that there is a constant $C(f_-) \geq 0$ which bounds the length of any backtracking subpath in the (unreduced) image $f_-(\gamma)$ of any reduced path γ in τ_-.

Remark 5.3. It follows that for any reduced path γ in τ_- and any subpath γ' of γ the path $[f_-(\gamma')]\dagger_{C(f_-)}$ is a subpath of the reduced path $[f_-(\gamma)]$. Note that this statement is not necessarily true if $[f_-(\gamma')]\dagger_{C(f_-)}$ is replaced by the path $[f_-(\gamma')]$.

Lemma 5.4. *There exists a constant $C_0 > 0$ such that, for any $C \geq C_0$ and for $t \geq 1$ as in Lemma 5.2, the following holds:*

For any path $\gamma \in \mathcal{L}_C$ there exists a path $\gamma' \in \mathcal{L}_C$ with the property that γ is a subpath of $[f_-^t(\gamma')]\dagger_{C(f_-)}$.

Proof. This is a direct consequence of the above Lemma 5.2. □

Recall from Convention 3.25 that we can assume without loss of generality that every endpoint of an INP in τ_- is a vertex of τ_-.

Lemma 5.5. *(a) For C as in Lemma 5.4 there is at most one illegal turn on any path $\gamma \in \mathcal{L}_C$. This illegal turn must be the tip of some INP which is contained as subpath in γ.*

(b) Furthermore, if γ is legal, then there is at most one unused turn on any path $\gamma \in \mathcal{L}_C$. If γ contains an INP, then all turns outside the INP are used turns. (To be specific: At the initial and terminal point of the INP the path γ may well cross over an unused (but legal) turn.)

Proof. (a) For any path $\gamma =: \gamma_0 \in \mathcal{L}_C$ we can assume by Lemma 5.4 that (after possibly replacing f by a positive power f^t) there is an infinite family of paths $\gamma_n \in \mathcal{L}_C$, for any integer $n \leq 0$, such that each γ_n is a subpath of $[f_-(\gamma_{n-1})]\dagger_{C(\varphi)}$. We note (compare Remark 3.14) that $\mathrm{ILT}(\gamma_n) \geq \mathrm{ILT}(\gamma_m)$, for any $n \leq m \leq 0$.

Since all γ_n have length bounded by C, it follows from Remark 3.14 that for sufficiently negative $n \leq m$ one has $\mathrm{ILT}(\gamma_n) = \mathrm{ILT}(\gamma_m)$. Thus it follows from Proposition 3.28, for $-m$ sufficiently large, that every illegal turn in γ_{m-r} is the tip of some INP which is entirely contained as subpath in γ_{m-r}, and adjacent such INP-subpaths can only overlap in a common endpoint. The same statement must be true for all $f^s(\gamma_{m-r})$ and hence for all γ_k with $k \geq m - r$, and thus also for $\gamma = \gamma_0$.

But the same argument applies to any of the γ_n. It follows that any maximal legal subpath γ' of γ_n, between two adjacent INP-subpaths, is the f_--image of the maximal legal subpath between adjacent INP-subpaths in γ_{n-1}. Since the endpoints of INPs are vertices, so that such non-degenerate maximal legal subpaths can not become arbitrary small. It follows from the assumption that f_- is expanding that each of the legal paths γ' between two adjacent INP-subpaths of any γ_n must has length 0. But by Hypotheses 5.1 there is only one INP in τ_-, and its endpoints are distinct. Thus no reduced path γ in τ_- can contain two subsequent INPs which have a common endpoint. This shows statement (a).

(b) Recall from Remark 3.31 (2) that used turns are mapped by $D^2 f_-$ to used turns. However, by Remark 3.31 (3), any unused legal turn must be mapped similarly to an unused legal turn, as one can take its (necessarily unused legal) preimage arbitrarily far back.

We first assume that γ (and hence any of the γ_n as in part (a)) is legal. From the last paragraph it follows that any maximal legal subpath γ' of γ_n, between two adjacent unused turns, is the f'-image of the maximal legal subpath between adjacent unused turns in γ_{n-1}. But the unused turns do only occur at vertices, so that such non-degenerate maximal legal subpaths can not become arbitrary small. It follows from the assumption that f' is expanding that each of the legal paths γ' between two adjacent unused turns of any γ_n must have length 0; in other words: there is only one such unused turn on any of the γ_n.

The same argument applies to the maximal legal subpath between an unused turn and the boundary point of an INP, in case some γ_n would contain both. It follows again that this maximal legal subpath must have length 0, so that the only unused turns, on any γ_n which contains an INP, can occur at the two endpoints of the INP. This proves claim (b). $\square$

As a direct consequence of Lemma 5.5 we obtain:

Proposition 5.6. *(a) For any $(X, Y) \in L^2_-$ there is on any of the geodesic realizations $\gamma := \gamma_{\tau_-}(X, Y)$ in τ_- at most one "singularity": This can either be an unused legal turn, or an illegal turn at the tip of an INP, or an INP with one or two unused legal turns at its boundary points. The remainder of γ is legal and crosses only over used turns.*

(b) If $\gamma = \gamma_{\tau_-}(X, Y)$ contains a non-used legal turn, then so does every $\gamma_t = \gamma_{\tau_-}(\partial\varphi^t(X), \partial\varphi^t(Y)) \in L^2_-$, for any $t \in \mathbb{Z}$. The analogous statement holds if γ crosses over an INP with k unused legal turns at its boundary points, for $k \in \{0, 1, 2\}$.

Proof. (a) It suffices to choose C in Lemma 5.5 large enough, so that the two singularities would be both contained in a path of $\mathcal{L}_C$, in contradiction to this lemma.
(b) This is a direct consequence of part (a), of the fact that the image of an INP is (after reduction) again an INP, and of the fact (explained in the first paragraph of the proof of part (b) of Lemma 5.5) that under the map $D^2 f_-$ used turns are mapped to used turns, and any unused legal turn on γ_t is mapped to an unused legal turn on γ_{t+1}. □

In the following we call a (possibly infinite or biinfinite) path γ *used legal* if it is legal and if it only crosses over used turns.

Corollary 5.7. *(a) For any $(X, Y) \in L^2_-$ every finite used legal subpath γ' of the geodesic realization $\gamma_{\tau_-}(X, Y)$ in τ_- is also contained as subpath in $f^t_-(e)$, for some $t \geq 0$ and some edge e of τ_-.*
(b) Assume that for $(X, Y) \in L^2_-$ the geodesic realization $\gamma_{\tau_-}(X, Y)$ is a used legal path. Then (X, Y) belongs to the BFH-attracting lamination $L^2_{BFH}(f_-)$ of $f_- : \tau_- \to \tau_-$ (see Definition 3.34).

Proof. (a) From the expansiveness of f_- and the legality of γ' we can use Proposition 5.6 (b) to "iterate" f_- backwards until we find a path γ'' which runs over at most two adjacent edges e, e' in τ_- such that $f^t_-(\gamma'')$ contains γ' as subpath, for some $t \geq 0$. But since all turns of γ' are used, the same must be true for all of its preimages, so that the turn between e and e' must be used. From Definition 3.30 it follows directly that this proves the claim (a).
(b) This is a direct consequence of part (a), by the definition of $L^2_{BFH}(f_-)$ in Definition 3.34. □

6. STEPS 5 - 7

Throughout this section we use the same conventions as in section 5, in particular Hypothesis 5.1 and the notation from the paragraph preceding it.

Lemma 6.1. *Let $(\rho_t)_{t \in \mathbb{Z}}$ be a family of infinite geodesics rays in τ_-, and assume that $[f_-(\rho_t)] = \rho_{t+1}$ for all $t \in \mathbb{Z}$. Assume furthermore that each ρ_t is legal, and that each ρ_t starts at a vertex v_t of τ_-. Then each ρ_t is an eigenray for the train track map $f_- : \tau_- \to \tau_-$.*

Proof. Since each ρ_t is legal, one has always $[f_-(\rho)] = f_-(\rho) = \rho_{t+1}$. It follows that f_- maps the starting point v_t of ρ_t to the starting point v_{t+1} of ρ_{t+1}. Since the starting point of each ρ_t is (by hypothesis) a vertex, it follows from the finiteness of the vertex set of τ_- that each v_t is a periodic vertex of τ_-.

Similarly, since all ρ_t are legal, the initial edge e_{t+1} of ρ_{t+1} must always be equal to the initial edge of the (legal) edge path $f_-(e_t)$, where e_t denotes the initial edge of ρ_t. Again by the finiteness of the graph τ_-, we see that for any $t \in \mathbb{Z}$ the initial edge e_t of ρ_t lies on a periodic orbit, with respect to the map Df_- (see Definition 3.6 (b)). In other words: each e_t is the eigen edge of the gate to which it belongs (compare subsection 3.3).

We now argue by induction: Suppose that all ρ_{kt} for some fixed integer $k \neq 0$ start with the same initial subpath γ_n of simplicial length $n \geq 1$. Then it follows from the expansiveness of f_- (and from the fact that all paths are legal) that their images have a common initial subpath of simplicial length $\geq n+1$. Thus the assumption $f_-^k(\rho_t) = \rho_{t+k}$ implies that the maximal common initial subpath of the ρ_t must have infinite length, so that we obtain $\rho_t = \rho_{t+k}$ for all $t \in \mathbb{Z}$. But this is the defining equality for eigenrays. $\square$

Lemma 6.2. *(a) For each eigenray ρ which starts at a vertex of τ_- there is an eigenray ρ', with same initial vertex, such that $\overline{\rho} \circ \rho'$ is a biinfinite used legal path.*

(b) Any path of type $\overline{\rho} \circ \rho'$, with used legal turn at the concatenation point, realizes a pair $(X, Y) \in \partial^2 F_N$ which belongs to the BFH-attracting lamination $L^2_{BFH}(f_-)$.

Proof. (a) We first derive from Lemma 3.32 that there is an edge e with same initial vertex v as ρ such that $\overline{e} \circ \rho$ is used legal. In particular e doesn't lie in the gate from which ρ starts. By Proposition 3.19 there is an eigenray ρ' which starts from the gate that contains e, and by Lemma 3.33 the turn from ρ' to ρ is used. Furthermore, every eigenray is used legal, since every finite subpath if contained in some f_--iterate of its initial edge (compare Proposition 3.19 and its proof). This proves the claim.

(b) By part (a) the turn at the concatenation point between $\overline{\rho}$ and ρ' is used: There is an edge e of τ_- and an exponent $t \geq 1$ such that the edge path $f_-^t(e)$ contains a subpath which consists of non-trivial initial segments of both, $\overline{\rho}$ and ρ'. It follows from the expansiveness of f_- and the definition of an eigenray that every path which consists of non-trivial initial segments of both, $\overline{\rho}$ and ρ', is a a subpath of $f_-^{nt}(e)$ for some integer $n \geq 0$. By Definiton 3.34 it follows that (X, Y) is an element $L^2_{BFH}(f_-)$. $\square$

Recall from subsection 2.2 that for any $(X, Y) \in \partial^2 F_N$ the geodesic realization in τ_- satisfies $[f_-(\gamma_{\tau_-}(X, Y))] = \gamma_-(\partial\varphi^{-1}(X), \partial\varphi^{-1}(Y))$.

As in section 5 we use the for the lamination dual to the forward limit tree T_+ the abbreviation $L^2_- = L^2(T_+)$.

Proposition 6.3. *Let $(X, Y) \in \partial^2 F_N$, and let $(\gamma_t)_{t\in\mathbb{Z}}$ be a family of biinfinite geodesics $\gamma_t = \gamma_{\tau_-}(\partial\varphi^{-t}(X), \partial\varphi^{-t}(Y))$ which realize the pair $(\partial\varphi^{-t}(X), \partial\varphi^{-t}(Y)) \in \partial^2 F_N$ in τ_-. In particular we assume that $[f_-(\gamma_t)] = \gamma_{t+1}$ for all $t \in \mathbb{Z}$.*

(a) Assume that each γ_t is legal, and that on each γ_t there is precisely one unused legal turn. Then there are eigenrays ρ and ρ', such that $\gamma_t = \overline{\rho} \circ \rho'$, where the unused turn is situated at the concatenation point.

(b) Assume that on each γ_t there is precisely one INP subpath η_t, and that every turn of γ_t is legal except for the turn at the tip of η_t. Then there are eigenrays ρ and ρ', such that $\gamma_t = \overline{\rho} \circ \eta \circ \rho'$.

Proof. Note that in case (a) each unused turn in γ_t takes place at a vertex v_0^t of τ_-. Similarly, in case (b) the endpoints v_1^t and v_2^t of any INP in γ_t are vertices, by Convention 3.25. Notice also that in case (a) one has $[f_-(\gamma_t)] = f_-(\gamma_t)$, and in case (b) the only backtracking subpath of $f_-(\gamma_t)$ is situated at the tip of η_{t+1}. Hence f_- maps the "singular" vertices v_i^t in γ_t to the singular vertices v_i^{t+1} in γ_{t+1}. Thus we obtain the statements of (a) and (b) is a direct consequence of Lemma 6.1. ◻

Definition 6.4. Let v be a periodic vertex v of τ_-. For any two eigenrays ρ, ρ' with initial vertex v we write

$$\rho \sim \rho'$$

if in the biinfinite path $\overline{\rho} \circ \rho'$ the turn at the concatenation point is used. We consider the equivalence relation which is generated by this relation, which we also denote by $\sim$.

It follows directly from the definitions that the equivalence relation $\sim$ is modeled precisely after the definition of the diagonal closure, so that one observes:

Remark 6.5. Let v be a periodic vertex of τ_-, and assume that the BFH-attracting lamination $L^2_{BFH}(f_-)$ contains pairs (X_1, Y_1) and (X_2, Y_2). Assume that the geodesic realizations $\gamma_1 = \gamma_{\tau_-}(X_1, Y_1)$ and $\gamma_2 = \gamma_{\tau_-}(X_2, Y_2)$ are both concatenations of eigenrays, $\gamma_1 = \overline{\rho}_1 \circ \rho'_1$ and $\gamma_2 = \overline{\rho}_2 \circ \rho'_2$, with concatenation vertex equal to v. Then $\rho_1 \sim \rho_2$ implies that (X_1, X_2) is contained in the diagonal closure of $L^2_{BFH}(f_-)$.

Note that the converse implication is also true, but it is a little bit less immediate: one needs to employ Lemma 6.1 and Lemma 6.2 (b).

We will now use the equivalence relation $\sim$ to derive from τ_- and f_- a new graph τ_-^* and a new self-map f_-^*. This will be defined in several steps as follows:

We assume that at some periodic vertex v of τ_- there is more than one $\sim$-equivalence class of eigenrays. We first assume that v is fixed by f_-.

We then split the vertex into several new vertices, called *subvertices* of v, precisely one for each equivalence class, such that all eigenrays starting at the same subvertex are equivalent, and conversely.

We now attach any other edge e of τ_- with initial vertex v to one of the new subvertices, according to which eigenray started in τ_- from the same gate as e. Recall from Proposition 3.19 that for each gate $\mathfrak{g}_i$ at a periodic vertex there is precisely one eigenray starting from $\mathfrak{g}_i$.

We connect the $k \geq 2$ subvertices of v by a k-pod $\mathcal{P}$ to obtain a new graph which transforms back to τ_-, if $\mathcal{P}$ is contracted to a single point.

We now consider any vertex v' of τ_- which is mapped eventually by f_- to v, and we do the analogous construction there, where the subdivision of v' into subvertices is simply lifted from v to v'. (Indeed, since v' is not periodic, we can not use again the relation $\sim$ for the subdivision of v'.) The graph obtained by this *blow-up* of vertices is denoted by τ_-^*.

We observe that the train track map f_- induces naturally a map $f_-^* : \tau_-^* \to \tau_-^*$ which maps $\mathcal{P}$ homeomorphically onto itself and leaves also $\tau_-^* \smallsetminus \mathring{\mathcal{P}}$ invariant. Since f_- is expanding, it follows that $\tau_-^* \smallsetminus \mathring{\mathcal{P}}$ has non-trivial fundamental group. Hence $\pi_1(\tau_-^* \smallsetminus \mathring{\mathcal{P}})$ is a non-trivial proper φ-invariant free factor of F_N, contradicting the assumption that f_- represents an iwip automorphisms.

Now, if v is not fixed by f_-, one has to do the same blow-up procedure at the whole f_--orbit of v, but the argument remains in this case precisely the same as in the simpler case just considered.

This shows:

Lemma 6.6. *At every periodic vertex v of τ_- there is only one $\sim$-equivalence class of eigenrays.* □

Proposition 6.7. *Let $(X, Y) \in L_-^2$ be such that its geodesic realization $\gamma_{\tau_-}(X, Y)$ in τ_- is not used legal. Then (X, Y) belongs to the diagonal closure of the BFH-attracting sublamination $L_{BFH}^2(f_-)$ of the train track map f_-.*

Proof. By Proposition 5.6 we know that $\gamma := \gamma_{\tau_-}(X, Y)$ (and hence any γ_t which realizes $(\partial\varphi^t(X), \partial\varphi^t(Y))$, for arbitrary $t \in \mathbb{Z}$) contains either precisely one non-used legal turn, or else it runs precisely once over an INP-subpath η of τ_-. Thus one deduces from Proposition 6.3 that $\gamma = \overline{\rho} \circ \rho'$ or $\gamma = \overline{\rho} \circ \eta \circ \rho'$, for eigenrays ρ and ρ'.

In the first case the unused turn occurs at the concatenation vertex v of the two eigenrays, and it follows directly from Lemma 6.2 that there are eigenrays ρ_1 and ρ_2 at v such that the biinfinite paths $\gamma_1 = \overline{\rho} \circ \rho_1$ and $\gamma_2 = \overline{\rho}' \circ \rho_2$ are used legal, and that they are geodesic realizations $\gamma_1 = \gamma_{\tau_-}(X, Z_1)$ and $\gamma_2 = \gamma_{\tau_-}(Y, Z_2)$ of pairs (X, Z_1) and (Y, Z_2) that both belong to to the BFH-attracting lamination $L^2_{BFH}(f_-)$. Hence it follows directly from Lemma 6.6 and Remark 6.5 that (X, Y) belongs to the diagonal closure of $L^2_{BFH}(f_-)$.

In the second case, where $\gamma = \overline{\rho} \circ \eta \circ \rho'$, we consider the eigenrays ρ_1 and ρ_2 which have the two legal branches of η as initial subpaths and agree along an infinite legal subray, see Remark 3.23. By the case treated before, the biinfinite paths $\overline{\rho} \circ \rho_1$ and $\overline{\rho}_2 \circ \rho'$ realize elements (X, Z) and (Z, Y) which belong to the diagonal closure of L^2_{BFH}. Thus (X, Y) also belongs to this diagonal closure. $\square$

We have now assembled all ingredients necessary to prove the main result of this paper:

Proof of Theorem 1.2. In Corollary 5.7 (b) together with Proposition 6.7 it is shown that $L^2(T_+)$ is contained in $\overline{\text{diag}}(L^2_{BFH}(f_-))$. More precisely, a pair $(X, Y) \in L^2(T_+)$ is contained in $\overline{\text{diag}}(L^2_{BFH}(f_-)) \smallsetminus L^2_{BFH}(f_-)$ if and only if its geodesic realization in τ_- is a concatenation of eigenrays at either an unused legal turn, or else at an INP. But since eigenrays in τ_- are in 1-1 relation with periodic gates of train track map $f_- : \tau_- \to \tau_-$ (see Corollary 3.20), there exist only finitely many distinct eigenrays. Since also the number of unused legal turns is finite, as well as the number of INPs in τ_-, it follows that $L^2(T_+)$ contains only finitely many F_N-orbits of pairs (X, Y) which are contained in $\overline{\text{diag}}(L^2_{BFH}(f_-)) \smallsetminus L^2_{BFH}(f_-)$. This establishes claim (4) of Theorem 1.2.

Furthermore, eigenrays of expanding maps are never periodic or eventually periodic (see Lemma 3.21), so that by Remark 2.3 (b) any minimal sublamination (see Definition 2.4) of $L^2(T_+)$ must intersect $L^2_{BFH}(f_-)$ and hence be contained in $L^2_{BFH}(f_-)$. Note that by Lemma 4.2 of [6] every algebraic lamination contains a minimal sublamination. But $L^2_{BFH}(f_-)$ is itself minimal (see Proposition 3.36), so that it must be contained in $L^2(T_+)$, as unique minimal sublamination. This proves parts (1) and (2) of Theorem 1.2.

By Proposition 2.10 the lamination $L^2(T_+)$ is diagonally closed, so that it contains with $L^2_{BFH}(f_-)$ also its diagonal closure. This inclusion, together with the converse inclusion stated in the first sentence of this proof, gives the equality claimed in part (3) of Theorem 1.2. $\square$

7. Discussion

Throughout this section T denotes always an $\mathbb{R}$-tree from $\overline{\text{cv}}_N$, and μ a current from $\text{Curr}(F_N)$. Recall that such a pair (T, μ) is called *perpendicular* if $\langle T, \mu \rangle = 0$. Recall also that following [13] this is equivalent to the statement $\text{supp}(\mu) \subset L^2(T)$, which in turn is equivalent to $\text{diag}(\text{supp}(\mu)) \subset L^2(T)$. The purpose of this section is to give (partial) answers to the following question:

Under which circumstances does this last inclusion actually improve to

$$\text{diag}(\text{supp}(\mu)) = L^2(T)\,?$$

To shape the discussion a bit, we propose:

Definition 7.1. (a) A pair $(T, \mu) \in \overline{\text{cv}}_N \times \text{Curr}(F_N)$ is said to be *diagonally equal* if $L^2(T) = \text{diag}(\text{supp}(\mu))$. In this case we also say that T *is diagonally equal to* μ, or μ *is diagonally equal to* T.

(b) A tree T is called *diagonally equalizable (DE)* if there exists a current μ such that (T, μ) is diagonally equal. Similarly, a current μ is called *diagonally equalizable (DE)* if there exists a tree T such that (T, μ) is diagonally equal.

(c) A tree T is called *totally diagonally equalizable (TDE)* if every perpendicular current μ is diagonally equal to T. Similarly, a current μ is called *totally diagonally equalizable (TDE)* if every perpendicular tree T is diagonally equal to μ.

(d) A tree T is called *uniquely diagonally equalizable (UDE)* if among all perpendicular currents there is (up to scalar multiples) precisely one current μ which is diagonally equal to T. Similarly, a current μ is called *uniquely diagonally equalizable (UDE)* if among all perpendicular trees there is (up to scalar multiples) precisely one tree T which is diagonally equal to μ.

All of these definitions descend directly to the projectivized objects, so that one can speak for example of a "uniquely diagonally equalizable" $[T] \in \overline{\text{CV}}_N$. The result of Theorem 1.1 can be rephrased by stating that for any atoroidal iwip $\varphi \in \text{Out}(F_N)$ the attracting fixed point $[T_+] \in \overline{\text{CV}}_N$ and the repelling fixed point $[\mu_-] \in \mathbb{P}\text{Curr}(F_N)$ constitute a pair which is diagonally equal.

By the above stated result of [13] we know that every diagonally equal pair $([T], [\mu])$ must be perpendicular. The set I_0 of such perpendicular pairs contains the uniquely determined minimal set $\mathcal{M}^2 \subset \overline{\mathrm{CV}}_N \times \mathbb{P}\mathrm{Curr}(F_N)$ with respect to the $\mathrm{Out}(F_N)$-action (see [11]), which is also contained in the cartesian product $\mathcal{M}^{cv} \times \mathcal{M}^{curr}$ of the (much better understood) minimal sets $\mathcal{M}^{cv} \subset \overline{\mathrm{CV}}_N$ and $\mathcal{M}^{curr} \subset \mathbb{P}\mathrm{Curr}(F_N)$. However, we only know these inclusions; it is open whether $\mathcal{M}^2$ is equal to $I_0 \cap (\mathcal{M}^{cv} \times \mathcal{M}^{curr})$.

Remark 7.2. We will give below examples of perpendicular pairs which are not diagonally equal. The set $\mathcal{DE}$ of diagonally equal pairs $([T], [\mu])$ is by definition $\mathrm{Out}(F_N)$-invariant and non-empty, but it will follow from the results presented below that $\mathcal{DE}$ is not closed in $\overline{\mathrm{CV}}_N \times \mathbb{P}\mathrm{Curr}(F_N)$.

Remark 7.3. The set $\mathcal{DE}$ defines also an $\mathrm{Out}(F_N)$-invariant subgraph of the intersection graph $\mathcal{I}(F_N)$ defined in [12], and it seems worthwhile to think which kind of a subgraph this is. Since all “limit pairs” $([T_+], [\mu_-])$ for atoroidal iwips are contained in $\mathcal{DE}$, it must consist of many distinct connected components. It seems likely that even the subgraph $\mathcal{DE}_0$ defined by all $\mathcal{DE}$-pairs which are contained in the main component $\mathcal{I}_0(F_N)$ of $\mathcal{I}(F_N)$ is non-connected.

We will now turn to the question of necessary and sufficient conditions for a pair (T, μ) to be diagonally equal, and for T or μ to be DE, TDE or UDE. We first recall some facts:

Facts 7.4. (1) For every algebraic lamination L^2 over F_N there exist a current $\mu \neq 0$ with $\mathrm{supp}(\mu) \subset L^2$ (see [8]). In particular, for every tree $T \in \partial\mathrm{cv}_N := \overline{\mathrm{cv}}_N \smallsetminus \mathrm{cv}_N$ there exists a current $\mu \in \mathrm{Curr}(F_N)$ which is perpendicular to T.

The “dual” statement is wrong, if we accept the subspace $\mathrm{Curr}(F_N)_+ \subset \mathrm{Curr}(F_N)$ of currents $\mu \in \mathrm{Curr}(F_N)$ with *full support* (i.e. $\mathrm{supp}(\mu) = \partial^2 F_N$) as “dual” of cv_N. (Note that cv_N can be characterized also as the subspace in $\overline{\mathrm{cv}}_N$ that consists of all $\mathbb{R}$-trees T with $L^2(T) = \emptyset$). There exist currents $\mu \in \mathrm{Curr}(F_N) \smallsetminus \mathrm{Curr}(F_N)_+$ which are not perpendicular to any tree in $\overline{\mathrm{cv}}_N$. Examples of such *filling* currents are given in [14].

(2) If $T, T' \in \partial\mathrm{cv}_N$, and if there is a length decreasing F_N-equivariant map $T \to T'$, then one has $L^2(T) \subset L^2(T')$. This applies in particular to the case where T' results from T by contracting an F_N-invariant forest T_0: Here each connected component of T_0 is mapped to a distinct point in T'. If T_0 contains connected components with infinite diameter, then the inclusion $L^2(T) \subset L^2(T')$ is proper: one has $L^2(T) \neq L^2(T')$.

(3) Let $f : \tau \to \tau$ is an expanding train track map which represents some $\psi \in \mathrm{Aut}(F_N)$, but contrary to the case considered in the main body of this paper, we assume that the non-negative transition matrix $M(f)$ (assumed to be in normal form for non-negative matrices) has precisely two primitive diagonal submatrices, with eigenvalues λ_{top} and λ_{bottom} respectively, and that the upper right off-diagonal block is non-zero (while, according to the normal form, the lower left off-diagonal block must be the zero matrix). For such exponential *two-strata* train track maps one has to distinguish two cases:

If $\lambda_{top} < \lambda_{bottom}$, then both eigenvalues possess non-negative row eigenvectors $\vec{v}^*_{top}$ and $\vec{v}^*_{bottom}$ respectively, and each of them determines a *Perron-Frobenius-tree* $T_{top} \in \partial \mathrm{cv}_N$ and $T_{bottom} \in \partial \mathrm{cv}_N$ respectively. Both projective classes, $[T_{top}] \in \partial \mathrm{CV}_N$ and $[T_{bottom}] \in \partial \mathrm{CV}_N$ are fixed points of the induced action of the automorphism $\psi \in \mathrm{Out}(F_N)$ which is represented by this train track map f.

If $\lambda_{top} \geq \lambda_{bottom}$, then only $\vec{v}^*_{top}$ is non-negative, so that only $T_{top} \in \partial \mathrm{cv}_N$ exists (and is projectively fixed by ψ), but no T_{bottom} as in the other case.

On the other hand, if $\lambda_{top} > \lambda_{bottom}$, then there are two non-negative column eigenvectors $\vec{v}_{top}$ and $\vec{v}_{bottom}$ respectively, and both determine projectively ψ-invariant currents $\mu_{top}, \mu_{bottom} \in \mathrm{Curr}(F_N)$. If $\lambda_{top} \leq \lambda_{bottom}$, then there is only one non-negative column eigenvector $\vec{v}_{bottom}$ and only one projectively ψ-invariant current $\mu_{bottom} \in \mathrm{Curr}(F_N)$.

It turns out that it is easier to prove results about trees:

Proposition 7.5. *Let $T \in \partial cv_N = \overline{cv}_N \smallsetminus cv_N$.*

(1) *Not every $T \in \partial cv_N$ is DE. A counterexample is given, for any 2-strata exponential train track map (which represents some $\psi \in Aut(F_N)$) with $\lambda_{top} < \lambda_{bottom}$, by the Perron-Frobenius tree T_{bottom}.*
(2) *The tree T is TDE if and only if the dual lamination $L^2(T)$ is the diagonal closure of a (uniquely determined) minimal sublamination $L^2_0(T)$.*
(3) *Every $\mathbb{R}$-tree T which is UDE must also TDE. A TDE-tree T is UDE if and only if the minimal sublamination $L^2_0(T) \subset L^2(T)$ is uniquely ergodic.*

Proof. (1) The dual lamination of T_{bottom} consists, from the bottom up, of a minimal sublamination L^2_{min} which is carried by the bottom stratum and has geodesic realizations by used legal paths. This lamination can be completed (by finitely many additional leaves) to $\mathrm{diag}(L^2_{min})$, which is still carried by the bottom stratum. There is another (non-minimal) lamination L^2_{used} which is also carried by used legal paths, but those may run through both strata of τ. The lamination $L^2(T_{bottom})$ is equal to the diagonal closure of L^2_{used}.

The current μ_{bottom} has as support the lamination L^2_{min}. Any other current μ' with support in $L^2(T_{bottom})$ can be decomposed as convex linear combination of μ_{bottom} and a current μ'_0 with support in $L^2(T_{bottom})$ where μ'_0 is *extremal* (i.e. no non-zero multiple of μ_{bottom} can be subtracted). The set of such extremal currents (all carried by $L^2(T_{bottom})$) is ψ-invariant, and by construction does not contain μ_{bottom}.

We consider the column vectors defined by these extremal currents and consider the action of f, which commutes with the action of ψ on the currents. But the vector space of these column vectors is finite dimensional (which isn't necessarily true for the space of currents carried by $L^2(T_{bottom})$, as $\mathbb{P}\mathrm{Curr}(F_N)$ is infinite dimensional !), so that any invariant proper subspace would lead to a second eigenvector, which we know does not exist. Thus there must be an extremal current which projects to the same vector as μ_{bottom}. This contradicts our results from [16].

(2) The "if" direction is immediate. In order to prove the "only if" direction we consider a minimal sublamination L'^2 which has strictly smaller diagonal closure than $L^2(T_{bottom})$. By Fact 7.4 (1) there is a current μ' with support in L'^2, which hence can not be diagonally equal to T_{bottom}.

(3) If T is not TDE, then by (2) in $L^2(T_{bottom})$ there is a sublamination L'^2 with strictly smaller diagonal closure than $L^2(T_{bottom})$, which by (1) carries some current μ'. In particular μ' is different from any current μ which is diagonally equal to T_{bottom}. Hence with any such μ every $\mu + \varepsilon\mu'$ is also diagonally equal to T_{bottom}, for any $\varepsilon > 0$.

The statement about unique ergodicity is really only a definitory statement. □

For the dual setting we get analogous, but in part weaker results:

Proposition 7.6. *Let $\mu \in Curr(F_N) \smallsetminus Curr(F_N)_+$, and assume that μ is not filling.*

1. *Not every non-filling $\mu \in Curr(F_N) \smallsetminus Curr(F_N)_+$ is DE. Counterexamples are given, for any 2-strata exponential train track map with $\lambda_{top} > \lambda_{bottom}$, by the Perron-Frobenius current μ_{top}.*
2. *A current μ is TDE if its support $supp(\mu)$ has diagonal closure which is a maximal element (with respect to the inclusion) of $\Lambda^2(F_N) \smallsetminus \{\partial^2 F_N\}$ ($diag(supp(\mu))$ is "maximal").*

(3) *A TDE-current μ with maximal diagonal closure of supp(μ) is UDE if and only if diag(supp(μ)) is uniquely ergometric.*

Proof. The proof is given by a dualization of the above proof of the corresponding parts of Proposition 7.5. The proof of part (1) becomes easier, since $\overline{\mathrm{CV}}_N$ is finite dimensional. □

References

[1] M. Bestvina and M. Feighn, *Outer Limits*, preprint, 1993; http://andromeda.rutgers.edu/~feighn/papers/outer.pdf

[2] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups.* Geom. Funct. Anal. **7** (1997), no. 2, 215–244

[3] M. Bestvina, and M. Handel, *Train tracks and automorphisms of free groups.* Ann. of Math. (2) **135** (1992), no. 1, 1–51

[4] M. Cohen and M. Lustig, *Very small group actions on R-trees and Dehn twist automorphisms.* Topology **34** (1995), no. 3, 575–617

[5] D. Cooper, *Automorphisms of free groups have finitely generated fixed point sets.* J. Algebra **111** (1987), 453–456

[6] T. Coulbois, A. Hilion, and M. Lustig, $\mathbb{R}$*-trees and laminations for free groups I: Algebraic laminations,* J. Lond. Math. Soc. (2) **78** (2008), no. 3, 723–736

[7] T. Coulbois, A. Hilion, and M. Lustig, $\mathbb{R}$*-trees and laminations for free groups II: The dual lamination of an* $\mathbb{R}$*-tree,* J. Lond. Math. Soc. (2) **78** (2008), no. 3, 737–754

[8] T. Coulbois, A. Hilion, and M. Lustig, $\mathbb{R}$*-trees and laminations for free groups III: Currents and dual* $\mathbb{R}$*-tree metrics,* J. Lond. Math. Soc. (2) **78** (2008), no. 3, 755–766

[9] M. Culler, K. Vogtmann, *Moduli of graphs and automorphisms of free groups.* Invent. Math. **84** (1986), no. 1, 91–119

[10] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups.* Duke Math. J. **93** (1998), no. 3, 425–452

[11] I. Kapovich and M. Lustig, *The actions of* $Out(F_k)$ *on the boundary of outer space and on the space of currents: minimal sets and equivariant incompatibility.* Ergodic Theory Dynam. Systems **27** (2007), no. 3, 827–847

[12] I. Kapovich and M. Lustig, *Geometric Intersection Number and analogues of the Curve Complex for free groups*, Geometry & Topology **13** (2009), 1805–1833

[13] I. Kapovich and M. Lustig, *Intersection form, laminations and currents on free groups*, Geom. Funct. Analysis (GAFA), to appear; http://www.arxiv.org/abs/0711.4337

[14] I. Kapovich and M. Lustig, *Ping-pong and Outer space*, preprint, 2009; arXiv:0902.4017

[15] I. Kapovich and M. Lustig, *Stabilizers of* $\mathbb{R}$*-trees with free isometric actions of* F_N, to appear in J. Group Theory; arXiv:0904.1881

[16] I. Kapovich and M. Lustig, work in progress

[17] G. Levitt and M. Lustig, *Irreducible automorphisms of F_n have North-South dynamics on compactified outer space.* J. Inst. Math. Jussieu **2** (2003), no. 1, 59–72

[18] M. Lustig. *Structure and conjugacy for automorphisms of free groups I.* MPI-Bonn preprint series 2000, No. 241; http://www.mpim-bonn.mpg.de/preprints

[19] M. Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, in "Geometric Group Theory", Trends in Mathematics, 197–224. Birkhäuser Verlag, Basel, 2007

[20] M. Lustig, *The F_N-action on the product of the two limit trees for an iwip automorphism*, in "Geometric Group Theory: Dortmund and Carleton conferences", Ed. O. Bogopolski, I. Bumagin, O. Kharlampovich, E. Ventura, Trends in Mathematics, Birkhäuser

Department of Mathematics, University of Illinois at Urbana-Champaign, 1409 West Green Street, Urbana, IL 61801, USA
http://www.math.uiuc.edu/~kapovich/

E-mail address: kapovich@math.uiuc.edu

Mathématiques (LATP), Université Paul Cézanne -Aix Marseille III, 52 av. Escadrille Normandie-Niémen, 13397 Marseille 20, France

E-mail address: Martin.Lustig@univ-cezanne.fr