Journal of Topology 3 (2010) 280–310
doi:10.1112/jtopol/jtq008

Growth of intersection numbers for free group automorphisms

Jason Behrstock, Mladen Bestvina and Matt Clay

Abstract

For a fully irreducible automorphism ϕ of the free group F_k we compute the asymptotics of the intersection number $n \mapsto i(T, T'\phi^n)$ for the trees T and T' in Outer space. We also obtain qualitative information about the geometry of the Guirardel core for the trees T and $T'\phi^n$ for n large.

Introduction

Parallels between $\mathrm{GL}_n(\mathbb{Z})$, the mapping class group $\mathcal{MCG}(\Sigma)$ and the outer automorphism group of a free group $\mathrm{Out}(F_k)$ drive much of the current research of these groups and is particularly fruitful in the case of $\mathrm{Out}(F_k)$. The paper [**7**] lists many similarities between these groups and uses known results in one category to generate questions in another. A significant example of this pedagogy is the question of the existence of a complex useful for studying the large-scale geometry of $\mathrm{Out}(F_k)$ analogous to the spherical Tits building for $\mathrm{GL}_n(\mathbb{Z})$ or Harvey's curve complex [**17**] for the mapping class group.

The curve complex is a simplicial complex which has vertices that correspond to homotopy classes of essential simple closed curves and which have simplices that encode when curves can be realized disjointly on the surface. The curve complex has played a large role in the study of the mapping class group; one of the first major results was the computation of its homotopy type and its consequences for homological stability, dimension and duality properties of the mapping class group [**15**, **16**]. Another fundamental result is that the automorphism group of the curve complex is the (full) mapping class group (except for some small complexity cases) [**19**, **22**, **24**]. The curve complex has also played an important role in understanding the large-scale geometry of the mapping class group [**1**]; a key property there is that the curve complex is Gromov hyperbolic [**25**].

The situation with $\mathrm{Out}(F_k)$ seems to be much more complicated and an emphasis on a particular feature of the curve complex leads to a different analog. A discussion of some of these analogs and their basic properties is provided in [**21**]. Without much doubt, a construction of such a complex and a proof of its hyperbolicity is the central question in the study of $\mathrm{Out}(F_k)$ today. In this introduction we feature three candidate complexes.

Recall that an (outer) automorphism of F_k is *fully irreducible* (sometimes called *irreducible with irreducible powers* (*iwip*)) if no conjugacy class of a proper free factor is periodic. These automorphisms are analogous to pseudo-Anosov homeomorphisms in mapping class groups. Further recall that Culler and Vogtmann's Outer space $\mathcal{CV}_k$ is the space of minimal free simplicial (metric) F_k-trees normalized such that the volume of the quotient graph is 1 (see [**10**]). We consider the unprojectivized version cv_k as well.

The complex of free factors of a free group. An n-simplex in this complex is a chain $FF_0 < FF_1 < \ldots < FF_n$ of nontrivial proper free factors in F_k modulo simultaneous conjugacy.

Received 8 July 2008; revised 10 December 2009; published online 30 April 2010.

2000 *Mathematics Subject Classification* 20E36, 20F65, 20F36.

The research was partially supported by NSF grants DMS-0812513 and DMS-0502441.

Hatcher and Vogtmann showed that this complex has the homotopy type of a bouquet of spheres, a result that is analogous to that of the spherical Tits building [**18**] and of the curve complex [**15**]. By analogy with the curve complex situation, where pseudo-Anosov homeomorphisms have unbounded orbits and other homeomorphisms have bounded orbits, Kapovich and Lustig have shown that fully irreducible automorphisms act with unbounded orbits and other automorphisms with bounded orbits [**21**]. Kapovich and Lustig proved their result via a notion of intersection number using geodesic currents analogous to a construction of Bonahon in the surface case [**5**].

The complex of (connected) subgraphs. A vertex is a nontrivial proper free factor of F_k modulo conjugacy. A collection FF_i of such free factors spans a simplex if they are compatible: there is a filtered graph $G_0 \subset G_1 \subset \ldots \subset G_m = G$ representing F_k so that each FF_i is represented by a connected component of some G_i. Just like the collection of very short curves in a hyperbolic surface (if nonempty) determines a simplex in the curve complex, leading to the homotopy equivalence between the thin part of Teichmüller space and the curve complex, so does the core of the union of very short edges of a marked metric graph (if nonempty) determine a simplex in the complex of subgraphs, leading to the homotopy equivalence between this complex and the thin part of the Outer space. There is a natural embedding of the free factor complex into the subgraph complex (the vertex sets are equal) and this embedding is a quasi-isometry.

The splitting complex. This complex is a refinement of the complex of free factors, where one also provides a complementary free factor. More precisely, this is the simplicial complex which has a vertex set that is the set of free product decompositions $F_k = A * B$ (modulo simultaneous conjugation and switching of the factors), where $n+1$ free product decompositions span an n-simplex if they are pairwise *compatible*. Two free product decompositions $A * B$ and $A' * B'$ are compatible if there is a two-edge graph of group decomposition $F_k = X * Y * Z$ such that collapsing one edge yields the decomposition $A * B$ and collapsing the other edge yields the decomposition $A' * B'$. The motivation for studying this complex comes from the observation that an essential simple closed curve on a surface determines a splitting of the fundamental group over $\mathbb{Z}$ (and not just one of the factors). Moreover, as we explain below, there is a hope that a proof of hyperbolicity of the curve complex generalizes to the splitting complex.

Scott and Swarup have shown that compatibility of $A * B$ and $A' * B'$ can be interpreted as the vanishing of an intersection number $i(A * B, A' * B')$ (see [**27**]). This number $i(-,-)$ is defined for any two splittings, either as an amalgamated free product or as an HNN-extension, of a finitely generated group. When the group is the fundamental group of a surface and the splittings arise from simple closed curves on the surface, this intersection number agrees with the geometric intersection between the two curves.

Guirardel, incorporating the work of [**11**], generalized Scott's intersection number to the setting of G-trees [**12**]. More importantly, given two G-trees T and T', Guirardel constructed a *core* $\mathcal{C}(T \times T') \subset T \times T'$. This core specifies a geometry for the pair of splittings in the sense that it is CAT(0) and it is equipped with two patterns, respectively, representing the splittings such that the geometric intersection number between the patterns in the quotient $\mathcal{C}(T \times T')/G$ is the intersection number $i(T, T')$. When the group is the fundamental group of a closed surface and the splittings arise from filling simple closed curves, the quotient $\mathcal{C}(T \times T')/G$ is the same surface endowed with the singular euclidean structure used by Bowditch in his proof of hyperbolicity of the curve complex [**6**]. It seems promising that following a careful understanding of the geometry of the core for free product decompositions of F_k, Bowditch's proof of the hyperbolicity of the curve complex may also show hyperbolicity of the splitting complex. In fact, the goal of this paper is much more modest. Instead of attempting to understand the geometry of the core for a pair of free splittings, we restrict ourselves to two points in Outer space, and the two points differ by a high power of a fully irreducible automorphism.

One of the main differences between the mapping class group and $\mathrm{Out}(F_k)$ is the inherent asymmetry present in $\mathrm{Out}(F_k)$. This difference that arises as a mapping class is represented by a homeomorphism of a surface, a symmetric object, whereas in general an outer automorphism of F_k is not represented by a homeomorphism of a graph but merely a homotopy equivalence, which has no symmetry imposed. The asymmetry is most easily seen in *expansion factors*. A pseudo-Anosov homeomorphism of a surface has two measured foliations, a stable and unstable foliation, and an expansion factor λ such that the homeomorphism multiplies the measure on the unstable foliation by the factor λ and the measure on the stable foliation by the factor λ^{-1} (see [**29**]). For the inverse of the homeomorphism, the roles of the foliations change place and the expansion factor λ remains the same. For a fully irreducible automorphism there is a homotopy equivalence of metric graphs $\sigma : \Gamma \to \Gamma$ representing the automorphism that maps vertices to vertices and linearly expands each edge of Γ by the same number λ, known as the *expansion factor*, and all positive powers of σ are locally injective on the interior of every edge (such maps are called *train-track maps* [**4**]). However, unlike the surface case, there is no reason why the expansion factor for an automorphism should equal the expansion factor of its inverse. Indeed the following automorphism and it inverse provide an example where the two expansion factors are not equal:

$$\begin{array}{rr} a \longmapsto b, & a \longmapsto cA, \\ \phi : b \longmapsto c, & \phi^{-1} : b \longmapsto a, \\ c \longmapsto ab, & c \longmapsto b. \end{array}$$

The expansion factor for ϕ is approximately 1.32 and the expansion factor for ϕ^{-1} is approximately 1.46. We remark that Handel and Mosher have shown that the ratio between the logarithms of the expansion factors is bounded by a constant depending on k (see [**13**]).

Intersection numbers for G-trees, as intersection numbers for curves on a surface, are symmetric (first proved by Scott [**26**], obvious for Guirardel's construction). They are also invariant under automorphisms: $i(T, T') = i(T\phi, T'\phi)$ for any G-trees T and T' and an automorphism ϕ of G (compare to $i(\alpha, \beta) = i(\psi(\alpha), \psi(\beta))$ for any curves α and β on a surface and a mapping class ψ). This imposes a symmetry on free group automorphisms. In particular, for any F_k-tree T and $\phi \in \mathrm{Out}(F_k)$ one has $i(T, T\phi^n) = i(T\phi^{-n}, T) = i(T, T\phi^{-n})$ for all n. This naturally leads one to inquire about the asymptotic behavior of the function $n \mapsto i(T, T'\phi^n)$.

In the surface setting, for a pseudo-Anosov homeomorphism ψ and any curves α and β the function $n \mapsto i(\alpha, \psi^n(\beta))$ behaves like $n \mapsto \lambda^n$, where λ is the expansion factor of ψ (see [**29**]). This leads one to first guess that for a fully irreducible automorphism ϕ and any two F_k-trees T and T' the function $i(T, T'\phi^n)$ also behaves like λ^n, where λ is the expansion factor of ϕ. However, as stated above, this cannot possibly be true in general since the asymptotics of $i(T, T'\phi^n)$ and $i(T, T'\phi^{-n})$ are the same but the expansion factors of ϕ and ϕ^{-1} need not be the same.

To state the correct result we remind the reader about the *stable tree* associated to a fully irreducible automorphism. For a fully irreducible automorphism $\phi \in \mathrm{Out}(F_k)$ and any tree $T \in \mathcal{CV}_k$, the sequence $T\phi^n$ has a well-defined limit in the compactification of the Outer space, called the stable tree [**3**]. The stable tree is a nonsimplicial (projectivized) $\mathbb{R}$-tree. It is *geometric* if it is dual to a measured foliation on a 2-complex [**2**, **23**]. Geometricity of the stable tree is characterized by a train-track representative for ϕ (see [**2**]); this characterization appears as Theorem 5.1. For two functions $f, g : \mathbb{R} \to \mathbb{R}$ we write $f \sim g$ to mean that there are constants K and C such that $(1/K)f(x) - C \leqslant g(x) \leqslant Kf(x) + C$.

THEOREM 1. *Suppose that $\phi \in \mathrm{Out}(F_k)$ is fully irreducible with an expansion factor λ and $T, T' \in \mathrm{cv}_k$. Let T^+ be the stable tree for ϕ and let μ be the expansion factor of ϕ^{-1}. Then we have either of the following.*
(1) *If T^+ is geometric, then $i(T, T'\phi^n) \sim \lambda^n$.*
(2) *If T^+ is nongeometric, then $i(T, T'\phi^n) \sim \lambda^n + \lambda^{n-1}\mu + \ldots + \lambda\mu^{n-1} + \mu^n$.*

REMARK 2. Case (2) in Theorem 1 can be simplified in two ways depending on whether or not $\lambda = \mu$. If $\lambda \neq \mu$, then $i(T, T'\phi^n) \sim \max\{\lambda, \mu\}^n$. If $\lambda = \mu$, then $i(T, T'\phi^n) \sim n\lambda^n$.

As a corollary we get a statement about the expansion factors for certain fully irreducible automorphisms. This corollary was first proved by Handel and Mosher.

COROLLARY 3 [**14**]. *Suppose that $\phi \in \mathrm{Out}(F_k)$ is fully irreducible with an expansion factor λ. Let T^+ and T^- be the stable trees for ϕ and ϕ^{-1}, respectively, and let μ be the expansion factor of ϕ^{-1}. If T^+ is geometric and T^- is nongeometric, then $\lambda > \mu$.*

Proof. If ϕ is as in the hypotheses of the corollary, then applying Theorem 1 we get $\lambda^n \sim \lambda^n + \lambda^{n-1}\mu + \ldots + \lambda\mu^{n-1} + \mu^n$. Therefore $\lambda > \mu$. □

By [**12**, Corollary 9.3] or [**14**, Corollary 3], if both T^+ and T^- are geometric, then ϕ is realized by a homeomorphism of a punctured surface. Therefore, the automorphism ϕ appearing in Corollary 3 is a *parageometric* automorphism, that is, an automorphism with a geometric stable tree that is not realized by a homeomorphism of a punctured surface.

REMARK 4. In view of this discussion, Theorem 1 can be more simply stated as follows:
(1) $i(T, T'\phi^n) \sim \max\{\lambda, \mu\}^n$ if $\lambda \neq \mu$ or if ϕ is realized on a surface;
(2) $i(T, T'\phi^n) \sim n\lambda^n$ if $\lambda = \mu$ and ϕ is not realized on a surface.

We briefly outline the rest of the paper. In Section 1 we collect all of the necessary properties of Guirardel's core that we need. Our first step toward Theorem 1 is showing that the asymptotics of $i(T, T'\phi^n)$ do not depend on the trees $T, T' \in \mathrm{cv}_k$ in Section 2. Following this, in Section 3 we derive an algorithm for computing Guirardel's core in terms of the *map on ends* associated to an automorphism and present an example. In Section 4, we further refine our algorithm to show that $i(T, T'\phi^n)$ is comparable to the volume of the subtree spanned by $(f^n)^{-1}(p)$, where $f : T \to T$ is a lift of a special representative for ϕ called a *train-track representative* (Proposition 4.3). The appearance of the inverse of the train-track map explains the appearance of both expansion factors in the second case of Theorem 1. Further in Section 4 we give a quick proof showing that $i(T, T'\phi^n)$ is asymptotically bounded below by $\max\{\lambda, \mu\}^n$.
The rest of the proof of Theorem 1 follows from the analysis in Section 5 of the subtree T_e^n in Proposition 4.3. There are two cases depending on whether the stable tree T^+ is geometric (Propositions 5.18 and 5.22). Putting together the three aforementioned propositions, we get Theorem 1. Finally we present examples of the core for both an automorphism with a geometric stable tree and an automorphism with a nongeometric stable tree highlighting the difference.
We conclude our introduction by recalling some standard definitions.

$\mathbb{R}$-trees: An *$\mathbb{R}$-tree* T (or simply *tree*) is a metric space such that any two points $x, y \in T$ are the endpoints of a unique arc and this arc is isometric to an interval of $\mathbb{R}$ of length $d_T(x, y)$. In particular every $\mathbb{R}$-tree has a Lebesgue measure. A point p is called a *branch point* if the

number of connected components of $T \setminus \{p\}$ is greater than 2. An $\mathbb{R}$-tree is a simplicial tree if the set of branch points is discrete. In this case, we denote the originating and terminating vertices of an oriented edge $e \subset T$ by $o(e)$ and $t(e)$, respectively, and the edge with reverse orientation by $\bar{e}$. A G-tree is a tree with an isometric action of G. We identify two G-trees if there is a G-equivariant isometry between them. A G-tree is *trivial* if there exists a global fixed point. We always assume that G is finitely generated and that G-trees are minimal (no proper invariant subtree).

We shall briefly recall the definition of the boundary of a tree T. A *ray* is an isometry $R : [0, \infty) \to T$. Two rays are equivalent if their images lie in a bounded neighborhood of one another; an equivalence class of rays is called an *end*. The equivalence class of the ray R is denoted by R_∞. A *geodesic* is an isometry $\rho : (-\infty, \infty) \to T$ to which are associated two ends denoted by ρ_∞ and $\rho_{-\infty}$. We often identify a ray or geodesic with its image in T. The *boundary* of T, denoted by ∂T, is the set of ends. If a basepoint $p \in T$ is fixed, the set of ends can be identified with the set of rays that originate at p. This enables us to topologize the boundary; two rays are close if they have a large intersection. As such, ∂T is homeomorphic to a Cantor set.

For a G-tree T it is well known that every element $g \in G$ either fixes a point in T (elliptic) or else it has an axis A_g (a set isometric to $\mathbb{R}$) on which it acts by translation (hyperbolic). The set of rays $R \subset A_g$ such that $gR \subset R$ defines a unique point $\omega_T(g) \in \partial T$.

Morphism in cv_k. For trees $T, T' \in \mathrm{cv}_k$, by a *morphism* $f : T \to T'$ we mean an equivariant cellular map that linearly expands every edge of T over a tight edge path in T'. This definition of morphism differs from the usual definition for $\mathbb{R}$-trees. The notion of *bounded cancellation* arises frequently when dealing with automorphisms of free groups. There are many statements of bounded cancellation, the one of importance to us is that any morphism $f : T \to T'$ is a *quasi-isometry*. In particular, there exist positive constants K and C such that $(1/K)d_T(x, y) - C \leqslant d_{T'}(f(x), f(y)) \leqslant K d_T(x, y) + C$. The original statement of bounded cancellation, along with a proof, can be found in [**8**]. A morphism $f : T \to T'$ descends to a (linear) Lipschitz homotopy equivalence $\sigma : \Gamma \to \Gamma'$, where $\Gamma = T/F_k$ and $\Gamma' = T'/F_k$. Also, a morphism $f : T \to T'$ induces an equivariant homeomorphism $f_\infty : \partial T \to \partial T'$ called the *map on ends*.

Notation. The edge path obtained by tightening an edge path α relative to its endpoints is denoted by $[\alpha]$ and the concatenation of two edge paths α and β is denoted by $\alpha \cdot \beta$.

Train-track representatives. We recall the basics of train-tracks; see [**4**] for proofs. For a metric graph Γ, a cellular homotopy equivalence $\sigma : \Gamma \to \Gamma$ that is linear on edges is a *train-track map* if there is a collection $\mathcal{L}$ of unordered pairs of distinct germs of adjacent edges such that:

(1) $\mathcal{L}$ is closed under the action of σ;
(2) for an edge $e \subset \Gamma$, any pair of germs crossed by $\sigma(e)$ is in $\mathcal{L}$.

The unordered pairs in $\mathcal{L}$ are called *legal turns*; an unordered pair of distinct germs of adjacent edges not in $\mathcal{L}$ is called an *illegal turn*. An edge path is *legal* if it only crosses legal turns. There is a metric on Γ such that σ linearly expands each edge of Γ by the same factor λ, called the *expansion factor*. This factor is the Perron–Frobenious eigenvalue of the transition matrix for σ, a positive eigenvector for this eigenvalue specifies the metric on Γ. Bounded cancellation implies that there is a bound on the amount of cancellation when tightening $\sigma(\alpha \cdot \beta)$, where α and β are legal paths. We denote the optimal constant by $\mathrm{BCC}(\sigma)$. As such, when α, β and γ are legal paths, if $\mathrm{length}(\beta) > 2\mathrm{BCC}(\sigma)/(\lambda - 1)$, then the length of $[\sigma^n(\alpha \cdot \beta \cdot \gamma)]$ goes to infinity as $n \to \infty$. The number $2\mathrm{BCC}(\sigma)/(\lambda - 1)$ is called the *critical constant* for the map σ.

Nielsen paths. A *Nielsen path* γ is a tight path with $[\sigma(\gamma)] = \gamma$. A tight path γ is a *periodic Nielsen path* if $[\sigma^n(\gamma)] = \gamma$ for some $n > 0$ and $[\sigma^i(\gamma)]$, with $i = 0, 1, \ldots, n-1$, is an *orbit* of (periodic) Nielsen paths. In this paper we always consider periodic Nielsen paths and for convenience usually omit the adjective 'periodic'. A Nielsen path is *indivisible* if it is not a concatenation of nontrivial Nielsen subpaths, and similarly for orbits of Nielsen paths. An indivisible Nielsen path (or iNp) always has a unique illegal turn and the two legal pieces have equal lengths.

1. *The Guirardel core*

In this section we describe Guirardel's construction of a core $\mathcal{C}(T \times T') \subset T \times T'$ for two G-trees T and T'. In addition, we introduce some notation and state some of the basic properties of cores needed for the following. Roughly speaking, the core is the essential part of $T \times T'$ in terms of the diagonal action of G on the product $T \times T'$. The following definitions and remarks appear in [**12**].

Definition 1.1. Let T be a tree and let p be a point in T. A *direction at* p is a connected component of $T - \{p\}$. If T is simplicial and e is an oriented edge of T, then we use the notation δ_e to denote the direction at $o(e)$ that contains e. Given two trees T and T', a *quadrant* is a product $\delta \times \delta'$, where $\delta \subset T$ and $\delta' \subset T'$ are directions.

For a direction $\delta \subset T$, let $\delta_\infty \subset \partial T$ denote the set of ends determined by rays contained in δ.

Definition 1.2. Let T and T' be G-trees and fix a basepoint $(p, p') \in T \times T'$. A quadrant $\delta \times \delta' \subset T \times T'$ is *heavy* if there exists a sequence of elements $g_i \in G$ such that we have:
(1) $g_i(p, p') \in \delta \times \delta'$;
(2) $d_T(p, g_i p) \to \infty$ and $d_{T'}(p', g_i p') \to \infty$ as $i \to \infty$.
If the quadrant is not heavy, it is *light*.

Remark 1.3. The choice of a basepoint is irrelevant. In particular, if the intersection of a quadrant $\delta \times \delta'$ with the orbit of any point is a bounded set (or more generally has a bounded projection), then $\delta \times \delta'$ is light.

Remark 1.4. Suppose that we have an inclusion of quadrants $\delta \times \delta' \subseteq \eta \times \eta' \subset T \times T'$. If $\delta \times \delta'$ is heavy, then $g_i(p, p') \in \delta \times \delta' \subseteq \eta \times \eta'$ for some sequence of elements $g_i \in G$. As the second condition of Definition 1.2 does not depend on the quadrant, we see that the quadrant $\eta \times \eta'$ is heavy.

We can now define Guirardel's core. It is the part of $T \times T'$ that is not in any light quadrant.

Definition 1.5 *The Guirardel core*. Let T and T' be minimal G-trees. Let $\mathcal{L}(T \times T')$ be the collection of light quadrants in the product $T \times T'$. The *core* $\mathcal{C}(T \times T')$ is defined as

$$\mathcal{C}(T \times T') = T \times T' - \bigcup_{\delta \times \delta' \in \mathcal{L}(T \times T')} \delta \times \delta'.$$

Since the collection $\mathcal{L}(T \times T')$ of light quadrants is invariant with respect to the diagonal action of G on $T \times T'$, the group G acts on the core $\mathcal{C}(T \times T')$. Guirardel defines the *intersection number* between two G-trees T and T', denoted by $i(T, T')$, as the volume of $\mathcal{C}(T \times T')/G$.

The measure on $\mathcal{C}(T \times T')$ is induced from the product Lebesgue measure on $T \times T'$. If T and T' are simplicial, then the intersection number $i(T, T')$ is the sum of the areas of the 2-cells in $\mathcal{C}(T \times T')/G$.

Guirardel shows that this intersection number agrees with the usual notion of intersection number for simple closed curves on a surface when T and T' are the Bass–Serre trees for the splitting of the surface group associated to the curves. Also, Guirardel shows that this intersection number agrees with Scott's definition of intersection number for splittings [**26**] and relates the core $\mathcal{C}(T \times T')$ to various other constructions; see [**12**] for references and further motivation for this definition.

We now state several properties of the core proved by Guirardel that we need for the following.

Proposition 1.6 [**12**, Proposition 2.6]. *Let T and T' be G-trees. The core $\mathcal{C}(T \times T')$ has convex fibers, that is, the sets $\mathcal{C}(T \times T') \cap \{x\} \times T'$ and $\mathcal{C}(T \times T') \cap T \times \{x'\}$ are connected (possibly empty) for any $x \in T$ and $x' \in T'$.*

Recall that a G-tree is *irreducible* if there exist two hyperbolic elements which have axes that intersect in a compact set [**9**]. In particular any tree in cv_k is irreducible.

Proposition 1.7 [**12**, Proposition 3.1]. *Let T and T' be G-trees. If either T or T' is irreducible, then $\mathcal{C}(T \times T') \neq \emptyset$. In particular if $T, T' \in \mathrm{cv}_k$, then $\mathcal{C}(T \times T') \neq \emptyset$.*

In [**12**], the above proposition is more general, but for our purposes, the above version is sufficient. Whenever the core is nonempty, there is a stronger condition for heavy quadrants.

Proposition 1.8 [**12**, Corollary 3.8]. *Let T and T' be G-trees and suppose that $\mathcal{C}(T \times T') \neq \emptyset$. Then a quadrant $\delta \times \delta'$ is heavy if and only if there is an element $g \in G$ that is hyperbolic for both T and T' such that $\omega_T(g) \in \partial\delta$ and $\omega_{T'}(g) \in \partial\delta'$.*

It may happen that the core is not connected. In this case there is a procedure of adding 'diagonal' edges to the core, resulting in the *augmented core* $\hat{\mathcal{C}}$. Adding edges to $\mathcal{C}$ does not affect the volume of $\mathcal{C}/G$ and hence does not affect the intersection number.

2. *The core $\mathcal{C}(T \times T')$ for trees in* cv_k

In this section we give a method for computing the intersection number between two trees in cv_k and show that the asymptotics of $n \mapsto i(T, T'\phi^n)$ do not depend on the trees T and T'.

Lemma 2.1. *Let $f : T \to T'$ be a morphism between $T, T' \in \mathrm{cv}_k$. For the two directions $\delta \subset T$ and $\delta' \subset T'$ the quadrant $\delta \times \delta'$ is heavy if and only if $f_\infty(\delta_\infty) \cap \delta'_\infty \neq \emptyset$.*

Proof. Suppose that $\delta \times \delta'$ is heavy. By Proposition 1.7 the core $\mathcal{C}(T \times T')$ is nonempty, and hence by Proposition 1.8 there is an element $x \in F_k$ with $\omega_T(x) \in \delta_\infty$ and $\omega_{T'}(x) \in \delta'_\infty$. As $f_\infty(\omega_T(x)) = \omega_{T'}(x)$, we see that $f_\infty(\delta_\infty) \cap \delta'_\infty \neq \emptyset$.

Conversely, suppose that $f_\infty(\delta_\infty) \cap \delta'_\infty \neq \emptyset$. Let $R \subset \delta$ be a ray, which has an equivalence class that is mapped by f_∞ into δ'_∞. We can assume that $f(R) \subset \delta'$. As T/F_k is a finite graph, there is a point $R_0 \in R$ and elements $x_i \in F_k$ such that $x_i R_0 \in R$ and $d_T(R_0, x_i R_0) \to \infty$ as $i \to$

∞. Now $x_i f(R_0) = f(x_i R_0) \in f(R) \subset \delta'$ and by bounded cancellation $d_{T'}(f(R_0), x_i f(R_0)) \to \infty$ as $i \to \infty$. Thus the point $(R_0, f(R_0)) \in T \times T'$ and elements $x_i \in F_k$ witness $\delta \times \delta'$ as a heavy quadrant. □

Using the above lemma, we can determine which rectangles of $T \times T'$ are in the core $\mathcal{C}(T \times T')$. This is enough to compute the intersection number $i(T, T')$. The following definition is closely related to the notion of a (one-sided) *cylinder* in [**20**].

Definition 2.2. For an oriented edge $e \subset T$, the clopen subset of ∂T consisting of equivalence classes of geodesic rays originating at $o(e)$ and containing e is called a *box*, which we denote by $[\![e]\!]$. In other words, $[\![e]\!] = (\delta_e)_\infty$.

Lemma 2.3. *Let $f : T \to T'$ be a morphism between $T, T' \in \mathrm{cv}_k$. Given two edges $e \subset T$ and $e' \subset T'$, the rectangle $e \times e' \subset T \times T'$ is in the core $\mathcal{C}(T \times T')$ if and only if each of the subsets $f_\infty([\![e]\!]) \cap [\![e']\!]$, $f_\infty([\![\bar{e}]\!]) \cap [\![e']\!]$, $f_\infty([\![e]\!]) \cap [\![\bar{e}']\!]$ and $f_\infty([\![\bar{e}]\!]) \cap [\![\bar{e}']\!]$ is nonempty.*

Proof. Let (p, p') be an interior point in the rectangle $e \times e'$. There are two directions at each of p and p'. These four directions combine to give us four quadrants Q_i, with $i = 1, 2, 3, 4$. The important feature of these four quadrants is that any quadrant that contains the point (p, p') must also contain one of the Q_i. Thus by Remark 1.4, the Q_i quadrants are heavy if and only if (p, p') is in $\mathcal{C}(T \times T')$. As any of the four directions at p and p' lie in a bounded neighborhood of one of $\delta_e, \delta_{\bar{e}}, \delta_{e'}$ or $\delta_{\bar{e}'}$, the Q_i are heavy if and only if the quadrants $\delta_e \times \delta_{e'}, \delta_{\bar{e}} \times \delta_{e'}, \delta_e \times \delta_{\bar{e}'}$ and $\delta_{\bar{e}} \times \delta_{\bar{e}'}$ are heavy. Therefore by Lemma 2.1, each of the four quadrants above are heavy and hence $(p, p') \in \mathcal{C}(T \times T')$ if and only if the sets in the statement of the lemma are nonempty. As this is true for any point in the rectangle $e \times e'$, this rectangle is in the core if and only if each of these four sets is nonempty. □

Remark 2.4. In a similar manner, we can also determine exactly when a vertex $(v \times v')$, a vertical edge $(v \times e')$, a horizontal edge $(e \times v')$ or a 'diagonal edge' (in the augmented core replacing twice light rectangles; see [**12**] for details) is in the core. The conditions are not as simple to state for vertices or horizontal and vertical edges as there can be several directions at a vertex. It is easy to see that $e \times e'$ is a twice-light rectangle if and only if $f_\infty([\![e]\!]) = [\![e']\!]$.

Remark 2.5. Another way to phrase Lemma 2.3 is given two edges $e \subset T$ and $e' \subset T'$, the rectangle $e \times e'$ is in the core $\mathcal{C}(T \times T')$ if and only if there exist two geodesics $\rho^+, \rho^- : (-\infty, \infty) \to T$ which have an image that contains e such that $f_\infty(\rho^+_\infty), f_\infty(\rho^-_{-\infty}) \in [\![e']\!]$ and $f_\infty(\rho^+_{-\infty}), f_\infty(\rho^-_\infty) \in [\![\bar{e}']\!]$.

Definition 2.6. Let e be an edge of T. The *slice* of the core $\mathcal{C}(T \times T')$ above e is the set

$$\mathcal{C}_e = \{e' \in T' \mid e \times e' \subset \mathcal{C}(T \times T')\}.$$

Similarly define the slice $\mathcal{C}_{e'} \subset T$ for an edge e' of T'.

By Proposition 1.6, the slice $\mathcal{C}_e$ is a subtree of T'. Clearly, every rectangle in the core belongs to one of the sets $e \times \mathcal{C}_e$. For any interior point $p_e \in e$, the tree $\{p_e\} \times \mathcal{C}_e$ embeds in the quotient $\mathcal{C}(T \times T')/F_k$, as the stabilizer of any edge in T is trivial. Therefore the intersection number

$i(T,T')$ can be expressed as the sum

$$i(T,T') = \sum_{e \subset T/F_k} l_T(e)\,\mathrm{vol}(\mathcal{C}_e), \tag{2.1}$$

where $l_T(e)$ is the length of the edge $e \subset T$ and $\mathrm{vol}(\mathcal{C}_e)$ is the sum of the lengths of the edges in $\mathcal{C}_e$. We are therefore interested in finding the slices $\mathcal{C}_e$ for a set of representative edges for T/F_k.

Using the above characterization of the core, we can show that the asymptotics of $n \mapsto i(T, T'\phi^n)$ only depend on the automorphism ϕ and not on the trees $T, T' \in \mathrm{cv}_k$.

LEMMA 2.7. *For any $T, T' \in \mathrm{cv}_k$ there are constants $K \geqslant 1$ and $C \geqslant 0$ such that, for any $T'' \in \mathrm{cv}_k$ and any edge $e'' \subset T''$, we have*

$$\frac{1}{K}\,\mathrm{vol}(\mathcal{C}'_{e''}) - C \leqslant \mathrm{vol}(\mathcal{C}_{e''}) \leqslant K\,\mathrm{vol}(\mathcal{C}'_{e''}) + C,$$

where $\mathcal{C}_{e''}$ and $\mathcal{C}'_{e''}$ are the respective slices of the cores $\mathcal{C}(T \times T'')$ and $\mathcal{C}(T' \times T'')$ above the edge e''.

Proof. Fix morphisms $f : T \to T''$ and $g : T' \to T$. Then $h = f \circ g$ is a morphism from T' to T'' after tightening relative to the vertices of T'. This tightening does not affect the map on ends h_∞. Fix an edge $e'' \subset T''$. First we will show that there are constants $K_0 \geqslant 1$ and $C_0 \geqslant 0$ that only depend on the morphism $g : T' \to T$ such that $(1/K_0)\,\mathrm{vol}(\mathcal{C}'_{e''}) - C_0 \leqslant \mathrm{vol}(\mathcal{C}_{e''})$.

As $g : T' \to T$ is a quasi-isometry, there are constants K_1 and K_2 such that the following conditions hold.

(1) If α is a subpath of length at least K_1 contained in the intersection of two geodesics ρ and ρ' in T', then $[g(\alpha)] \cap [g(\rho)] \cap [g(\rho')]$ contains a segment of length 1.

(2) If the distance between geodesic segments α and α' in T' is at least K_2, then $[g(\alpha)] \cap [g(\alpha')] = \emptyset$.

For a finite subtree $X \subset T'$, let $N(X)$ be equal to the maximum cardinality of collection of segments of length K_1 in X that are pairwise distance at least K_2 apart. There exist constants $K_0 \geqslant 1$ and $C_0 \geqslant 0$ such that $(1/K_0)\,\mathrm{vol}(X) - C_0 \leqslant N(X)$.

Let $\mathcal{A}$ denote a maximal collection of segments of length K_1 in $\mathcal{C}'_{e''}$ that are pairwise distance at least K_2 apart. Given a segment $\alpha \in \mathcal{A}$, as $\alpha \subseteq \mathcal{C}'_{e''}$, by Lemma 2.3 (see Remark 2.5) we can find two geodesics $\rho^+ : \mathbb{R} \to T'$ and $\rho^- : \mathbb{R} \to T'$ that contain the segment α such that $h_\infty(\rho^+_\infty), h_\infty(\rho^-_{-\infty}) \in [\![e'']\!]$ and $h_\infty(\rho^+_{-\infty}), h_\infty(\rho^-_\infty) \in [\![\bar{e}'']\!]$, respectively.

Now let $\beta \subseteq T$ be a segment of length 1 contained in the intersection of $[g(\alpha)] \cap [g(\rho^+)] \cap [g(\rho^-)]$. Such a segment exists by condition (1) above. Thus $g_\infty(\rho^+_\infty), g_\infty(\rho^-_\infty) \in [\![\beta]\!]$. As $h_\infty = f_\infty \circ g_\infty$, we see that $f_\infty([\![\beta]\!]) \cap [\![e'']\!] \neq \emptyset$ and similarly for the three other sets in Lemma 2.3, and hence $\beta \in \mathcal{C}_{e''}$. Repeat the same for all other segments in $\mathcal{A}$. By condition (2) above, the segments in T corresponding to the $N(\mathcal{C}'_{e''})$ segments of $\mathcal{A}$ are disjoint, and therefore $N(\mathcal{C}'_{e''}) \leqslant \mathrm{vol}(\mathcal{C}_{e''})$. Hence $(1/K_0)\,\mathrm{vol}(\mathcal{C}'_{e''}) - C_0 \leqslant \mathrm{vol}(\mathcal{C}_{e''})$.

A similar argument using the morphisms $f' : T' \to T''$ and $g' : T \to T'$ shows that there are constants $K'_0 \geqslant 1$ and $C'_0 \geqslant 0$ only depending on the morphism $g' : T \to T'$ such that $(1/K'_0)\,\mathrm{vol}(\mathcal{C}_{e''}) - C'_0 \leqslant \mathrm{vol}(\mathcal{C}'_{e''})$. This proves the lemma. □

From this lemma, we can show the independence of the asymptotics of $i(T, T'\phi^n)$ in a special case.

COROLLARY 2.8. *For any $T, T' \in \mathrm{cv}_k$, and any automorphism $\phi \in \mathrm{Out}(F_k)$, we have $i(T, T\phi^n) \sim i(T', T\phi^n)$.*

Proof. This follows directly from (2.1) and Lemma 2.7 as the constants from the lemma only depend on T and T' and not on $T'' = T\phi^n$. □

From this special case we easily derive the full independence using equivariance and symmetry.

PROPOSITION 2.9. *For any $T, T', T'' \in \mathrm{cv}_k$, and any automorphism $\phi \in \mathrm{Out}(F_k)$, we have $i(T, T\phi^n) \sim i(T', T''\phi^n)$.*

Proof. Applying Corollary 2.8 along with equivariance and symmetry of intersection numbers, it follows that $i(T', T\phi^n) = i(T'\phi^{-n}, T) \sim i(T'\phi^{-n}, T'') = i(T', T''\phi^n)$. Combining this equivalence with the equivalence in Corollary 2.8, we get $i(T, T\phi^n) \sim i(T', T\phi^n) \sim i(T', T''\phi^n)$, as desired. □

3. *Slices of the core $\mathcal{C}(T \times T')$*

3.1. *The map on ends for a free group automorphism*

Let $f : T \to T'$ be a morphism between $T, T' \in \mathrm{cv}_k$. This descends to a Lipschitz linear homotopy equivalence $\sigma : \Gamma \to \Gamma'$, where $\Gamma = T/F_k$ and $\Gamma' = T'/F_k$. Fix a morphism $f' : T' \to T$ such that the induced map $\sigma' : \Gamma' \to \Gamma$ is a homotopy inverse to σ. Homotoping σ' if necessary, we assume that the image of any small open neighborhood of any vertex of Γ' is not contained in an edge of Γ. Fix basepoints $*' \in T'$ and $* \in T$ such that $f'(*') = *$. For simplicity, we denote the images of these basepoints in Γ and Γ' by $*$ and $*'$, respectively. We use the map $\sigma' : \Gamma' \to \Gamma$ to find the map on ends $f_\infty : \partial T \to \partial T'$. This is the inverse of the homeomorphism $f'_\infty : \partial T' \to \partial T$.

Let e be an (oriented) edge of Γ. Subdivide e into e_+e_- and denote the subdivision point by p_e. Fix a tight edge path $\alpha_e \subset \Gamma$ from $*$ to p_e which has a final edge e_+. The path α_e corresponds to choosing a representative lift for e in T. For simplicity, we call this lift e and the lift of p_e contained in this edge p_e. This lift e is oriented 'away' from $*$, that is, $o(e)$ separates p_e from $*$. Consider the set of points $\Sigma_e = (\sigma')^{-1}(p_e) \subset \Gamma'$. There is one point in this set for each edge of Γ' which has an image under σ' that crosses either e or $\bar{e}$ counted with multiplicities.

For $q \in \Sigma_e$ there is a tight path $\gamma_q \subset \Gamma'$ from $*'$ to q such that up to homotopy $\alpha_e = [\sigma'(\gamma_q)]$. Further, the path γ_q is unique as σ' is a homotopy equivalence. If, before tightening the path $\sigma'(\gamma_q)$, the final edge is e_+ then we assign q a '+' sign; Else, the final edge is $\bar{e}_-$ and we assign q a '−' sign.

Let $\widetilde{\gamma}_q$ be the lift of γ_q to T' that originates at $*'$. Denote the terminal point of $\widetilde{\gamma}_q$ by $\tilde{q}$. Let $\widetilde{\Sigma}_e \subset T'$ be the collection of the point $\tilde{q}$. In other words, $\widetilde{\Sigma}_e = (f')^{-1}(p_e)$. This set decomposes into $\widetilde{\Sigma}_e^+$ ('+' points) and $\widetilde{\Sigma}_e^-$ ('−' points), depending on the sign of the images of the points in Σ_e. Since f' is locally injective on the interior of edges, for any edge $e' \subset T'$, the set $\widetilde{\Sigma}_e \cap e'$ is at most a single point. Also, since f' is cellular, it follows that $\widetilde{\Sigma}_e$ does not contain any vertices of T'.

REMARK 3.1. Our hypothesis that the image of any small open neighborhood of any vertex in Γ' by σ' is not contained in an edge of Γ implies that every component of $T' - \widetilde{\Sigma}_e$ is unbounded.

EXAMPLE 3.2. Let $T \in \mathrm{cv}_2$ be the tree with all edge lengths equal to 1 and such that $\Gamma = T/F_2$ is the 2-rose which has petals that are marked a and b. We use capital letters to

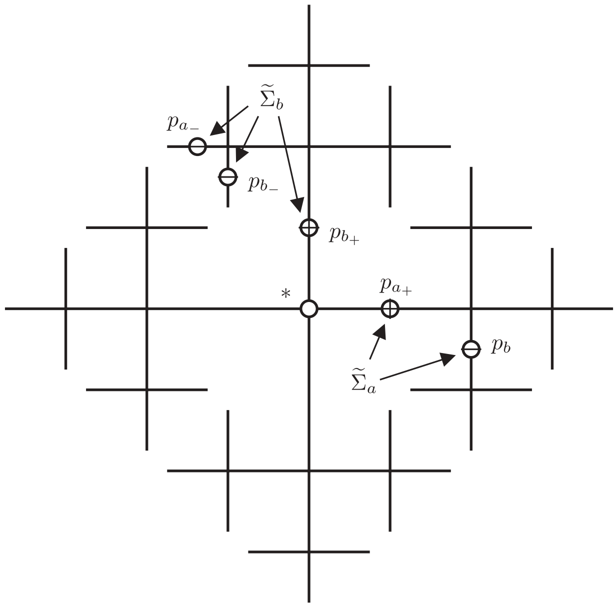

FIGURE 1. *The sets* $\widetilde{\Sigma}_a$ *and* $\widetilde{\Sigma}_b$ *in Example* 3.2. *The '+' points are displayed by* $\oplus$ *and the '−' points are displayed by* $\ominus$.

denote the edges with opposite orientation. Consider the automorphism $\phi \in \mathrm{Out}(F_2)$ given by $\phi(a) = ab$ and $\phi(b) = bab$. There is an obvious homotopy equivalence $\sigma' : \Gamma \to \Gamma$ representing the automorphism ϕ. We carefully look at the above construction for the map σ'.

Subdivide the edges of Γ as described above, creating the points p_a and p_b. Fix the preferred paths $\alpha_a = a_+$ and $\alpha_b = b_+$. Also subdivide $a = a_1a_2a_3$ with the subdivision points p_{a_+}, p_{a_-} and $b_+ = b_1b_2, b_- = b_3b_4$ with the subdivision points p_{b_+}, p_{b_-}. We may assume that $\sigma'(a_1) = a_+, \sigma'(a_2) = a_-b_+, \sigma'(a_3) = b_-$ and $\sigma'(b_1) = b_+, \sigma'(b_2) = b_-a_+, \sigma'(b_3) = a_-b_+, \sigma'(b_4) = b_-$.

Therefore $\Sigma_a = \{p_{a_+}, p_b\}$ and $\Sigma_b = \{p_{a_-}, p_{b_+}, p_{b_-}\}$. Then $[\sigma'(a_1)] = [a_+] = \alpha_a$ and $[\sigma'(aB_4B_3)] = [abBA_-] = a_+ = \alpha_a$. Hence p_{a_+} is a '+' point and p_b is a '−' point. Further $[\sigma'(bAA_3)] = [babBAB_-] = b_+ = \alpha_b$, $[\sigma'(b_1)] = [b_+] = \alpha_b$ and $[\sigma'(bAB_4)] = [babBAB_-] = b_+ = \alpha_b$. Therefore p_{b_+} is a '+' point and p_{a_-} and p_{b_-} are '−' points. Figure 1 shows the sets $\widetilde{\Sigma}_a$ and $\widetilde{\Sigma}_b$ in the tree T.

For $\tilde{q} \in \widetilde{\Sigma}_e$, let $e'_{\tilde{q}}$ be the oriented edge in T' that contains $\tilde{q}$ and such that the initial segment of $e'_{\tilde{q}}$ is contained in $\tilde{\gamma}_q$, that is, $e'_{\tilde{q}}$ is oriented away from $*'$. A ray containing $e'_{\tilde{q}}$ represents an end in the box $[\![e'_{\tilde{q}}]\!]$. A ray originating from $*'$ is mapped by f'_∞ into $[\![e]\!]$ if and only if it intersects $\widetilde{\Sigma}_e$ and if the final point in $\widetilde{\Sigma}_e$ it intersects is a '+' point. Likewise, a ray originating from $*'$ is mapped by f'_∞ into the complement of $[\![\bar{e}]\!]$ if and only if it does not intersect $\widetilde{\Sigma}_e$ or if the final point of $\widetilde{\Sigma}_e$ it intersects is a '−' point. Therefore, as f_∞ is the inverse of f'_∞, we can express $f_\infty([\![e]\!])$ as a union and difference of the boxes $[\![e'_{\tilde{q}}]\!]$ for $\tilde{q} \in \widetilde{\Sigma}_e$; see Examples 3.9 and 3.10.

A component $X \subset T' - \widetilde{\Sigma}_e$ is assigned a sign '+' if $f'(X)$ is contained in the direction $\delta_e \subset T$; else, the image $f'(X)$ is contained in the direction $\delta_{\bar{e}} \subset T$ and we assign it a '−' sign. In the following lemma we show that along any ray, the signs of the components of $T' - \widetilde{\Sigma}_e$ that the ray intersects alternate.

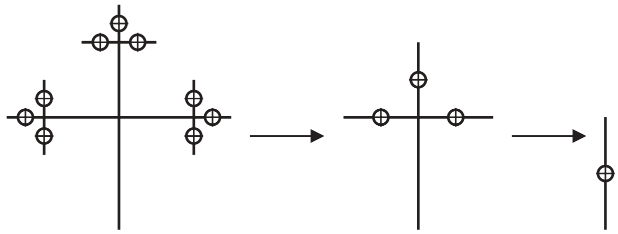

FIGURE 2. *Consolidating vertices in* T'_e.

LEMMA 3.3. *With the above notation, let* $R : [0,\infty) \to T'$ *be a ray that originates at a vertex of* T' *and suppose that there are two different components* $X_0, X_1 \subset T' - \widetilde{\Sigma}_e$ *which both have the same sign and for some* $0 \leqslant t_0 < t_1$ *we have* $R(t_0) \in X_0$ *and* $R(t_1) \in X_1$. *Then there is a component* $\widehat{X} \subset T' - \widetilde{\Sigma}_e$ *and* $\hat{t}$, *with* $t_0 < \hat{t} < t_1$, *such that* $R(\hat{t}) \in \widehat{X}$ *and the sign of* $\widehat{X}$ *is opposite to that of* X_0 *and* X_1.

Proof. Let $x_0 = R(t_0)$ and $x_1 = R(t_1)$ and consider the tight segment γ connecting x_0 to x_1; this is a subsegment of the ray $R([0,\infty))$. As X_0 and X_1 are different components of $T' - \widetilde{\Sigma}_e$, for some interior point $\hat{x}$ of γ necessarily $\hat{x} \in \widetilde{\Sigma_e}$. As f' maps vertices of T' to vertices of T, the point $\hat{x}$ is in the interior of some edge $\hat{e} \subset T'$. As f' is injective on the interior of edges, there are points $\hat{x}_+$ and $\hat{x}_-$ in $\hat{e}$ close to $\hat{x}$ that are in the components of $T' - \widetilde{\Sigma}_e$ with opposite sign. One of these points gives $R(\hat{t})$. □

Let $T'_e \subset T'$ be the subtree spanned by the points $\widetilde{\Sigma_e}$. If e' is an edge of T' that is not contained in T'_e and is oriented away from T'_e, then $f'_\infty([\![e']\!])$ is contained in either $[\![e]\!]$ or $[\![\bar{e}]\!]$. Thus, either $f'_\infty([\![e']\!]) \cap [\![e]\!]$ or $f'_\infty([\![e']\!]) \cap [\![\bar{e}]\!]$ is empty and therefore by Lemma 2.3, the edge e' is not in the slice $\mathcal{C}_e$. Hence, the slice $\mathcal{C}_e$ is contained in the subtree of interior edges of T'_e. We will show that the difference between the volume of T'_e and $\mathcal{C}_e$ is bounded, at least when the morphism $f' : T' \to T$ is the lift of a train-track representative.

First, let us consider a situation where an interior edge of T'_e is not in $\mathcal{C}_e$. Without loss of generality, this implies that any geodesic ray that contains this edge is eventually contained in a component of $T' - \widetilde{\Sigma}_e$ that has a '+' sign. This can happen if there is a vertex of T'_e such that all but one of its adjacent edges contains a point of $\widetilde{\Sigma}_e$, all of which are terminal vertices of T'_e. In this case we would like to remove a neighborhood of this vertex and *consolidate* the set of points, repeating if necessary; see Figure 2.

We now define this operation in detail and show that it terminates in the slice of the core $\mathcal{C}_e$. Let $V_t(T'_e)$ denote the set of terminal vertices of T'_e. A vertex $v \in T'_e$ is *full* if the valence of v in T'_e equals the valence of v in T'.

DEFINITION 3.4. Let $v \in T'_e$ be a full vertex with adjacent edges $\hat{e}, e_1, \ldots, e_m$ (that is, $o(\hat{e}) = o(e_1) = \ldots = o(e_m) = v$ and these are the only such edges with the originating vertex v). We say v is *removable* if $t(e_i) \in V_t(T'_e)$ for $i = 1, \ldots, m$ but $t(\hat{e}) \notin V_t(T'_e)$.

At a removable vertex v, subdivide the edge $\hat{e}$ into $\hat{e}_0\hat{e}_1$ and denote the subdivision vertex by $p_{\hat{e}}$. Since the image of any small open neighborhood of v is not contained in a single edge of T, the edge $\hat{e}$ does not contain any points in $\widetilde{\Sigma}_e$. Remove the $(m+1)$-pod $\hat{e}_0 \cup e_1 \cup \ldots \cup e_m$, leaving the vertex $p_{\hat{e}}$. Delete the points $\{t(e_1), \ldots, t(e_m)\}$ from $\widetilde{\Sigma}_e$ and add the point $p_{\hat{e}}$. In the new subtree spanned by $\widetilde{\Sigma}_e$, we have that $p_{\hat{e}}$ is a terminal vertex. This changes the components of $T' - \widetilde{\Sigma}_e$. The components of $T' - \{t(e_1), \ldots, t(e_m)\}$ that do not contain v are combined with the $(m+1)$-pod $\hat{e}_0 \cup e_1 \cup \ldots \cup e_m$. The signs of all of these components were the same before the operation, and we assign this sign to the new component.

This process may create removable vertices, specifically the vertex $t(\hat{e})$. Note that for this new collection of points $\widetilde{\Sigma}_e$ and components of $T' - \widetilde{\Sigma}_e$, Remark 3.1 and Lemma 3.3 still hold.

Repeat this process until there are no removable vertices. As T'_e is a finite subtree and at each step we remove a finite subtree that does not disconnect the new space, this process will terminate with a subtree Y'_e. In the following lemma we collect some elementary properties of the subtrees that are removed.

Lemma 3.5. *With the above notation, suppose that A is a connected component of $T'_e - Y'_e$. Then:*

(1) *all but one of the terminal vertices of A is a terminal vertex in T'_e and belongs to the set $\widetilde{\Sigma}_e$;*
(2) $\widetilde{\Sigma}_e \cap A \subset V_t(A)$.

Proof. If A has two terminal vertices that are not terminal vertices of T'_e, then Y'_e is not connected, which is a contradiction. As all of the terminal vertices of T'_e belong to $\widetilde{\Sigma}_e$, we see that (1) holds.

Suppose that $e' \subset A$ contains a point of $\widetilde{\Sigma}_e$ that is not a terminal vertex of T'_e. Orient e' to point away from the terminal vertex of A which is not a terminal vertex of T'_e. Then the component of $T'_e - \widetilde{\Sigma}_e$ that contains the vertex $t(e')$ is bounded, which is a contradiction, and hence (2) holds. □

Let Z'_e denote the subtree of edges in Y'_e that are not adjacent to valence one vertices of Y'_e, that is, Z'_e is the subtree consisting of the edges of Y'_e that are edges in T'.

Lemma 3.6. *For any oriented edge $e' \subset Z'_e$ there are rays $R^+ : [0,\infty) \to T'$ and $R^- : [0,\infty) \to T'$ containing e' with the specified orientation, such that $f'_\infty(R^+_\infty) \in [\![e]\!]$ and $f'_\infty(R^-_\infty) \in [\![\bar{e}]\!]$.*

Proof. Let $e' \subset T'_e$ and assume that, for every ray R containing e', we have $f'_\infty(R_\infty) \in [\![\bar{e}]\!]$. Let A be the component of $T'_e - o(e')$ that contains e'. We will show that A is removed from T'_e in the construction of Z'_e and hence e' is not in Z'_e.

Let v be the vertex of T' in A that is furthest away from $o(e')$. We claim that v is removable. If not, then by choice of v there is an edge e_1 in T' with $o(e_1) = v$ such that $\delta_{e_1} \cap \widetilde{\Sigma}_e = \emptyset$. Thus v is in a '$-$' component. However, again by choice of v, there is an edge e_2 in T' with $o(e_2) = v$ such that δ_{e_2} intersects $\widetilde{\Sigma}_e$ in a single point x, and this point x lies in the interior of e_2. By the hypothesis on e', the component of $T' - \{x\}$ that contains $t(e_2)$ is a '$-$' component. However, this contradicts Lemma 3.3. Thus v is removable. Remove the $(m+1)$-pod about v in T'_e to get a new subtree in which the hypothesis on e' still holds. Repeating in this fashion, we see that A is removed.

Hence, if $e' \subset Z'_e$, not every ray containing e' is mapped by f'_∞ to $[\![\bar{e}]\!]$. Similarly, we see that, for $e' \subset Z'_e$, not every ray that contains e' is mapped by f'_∞ to $[\![e]\!]$. □

Lemma 3.7. *For any edge $e \subset T$ the slice $\mathcal{C}_e$ of the core $\mathcal{C}(T \times T')$ is the subtree $Z'_e \subset T'$.*

Proof. We have already seen that $\mathcal{C}_e \subseteq T'_e$. By Lemma 3.5 we see that, for any edge $e' \subset T'_e - Z'_e$, one of the sets $f_\infty([\![e']\!])$ or $f_\infty([\![\bar{e}']\!])$ is contained in either $[\![e]\!]$ or $[\![\bar{e}]\!]$. Therefore by Lemma 2.3 we have $\mathcal{C}_e \subseteq Z'_e$. By Lemmas 2.3 and 3.6 we have $Z'_e \subseteq \mathcal{C}_e$. Hence $Z'_e = \mathcal{C}_e$. □

3.2. *Examples of* $\mathcal{C}(T \times T')$

We now present some examples of computing the map on ends and building the core for some trees $T, T' \in \mathrm{cv}_k$.

NOTATION 3.8. We adopt the convention of using '$[\![e]\!] + [\![e']\!]$' to denote the union of the two boxes $[\![e]\!], [\![e']\!]$. Also, when $[\![e']\!]$ is contained within $[\![e]\!]$, we use '$[\![e]\!] - [\![e']\!]$' to denote the set of ends contained within $[\![e]\!]$ but not $[\![e']\!]$.

EXAMPLE 3.9. Let $T \in \mathrm{cv}_2$ and $\phi \in \mathrm{Aut}(F_2)$ be as in Example 3.2. We can identify T with the Cayley graph for F_2 and the ends of T with right-infinite words in F_2. From Figure 1 we see the map on ends $f_\infty : \partial T\phi \to \partial T$ is given by

$$\begin{aligned} f_\infty : [\![a]\!] &\longmapsto [\![a]\!] - [\![aB]\!] \\ [\![b]\!] &\longmapsto [\![b]\!] - [\![bAB]\!] - [\![bAA]\!] \end{aligned}$$

In terms of right-infinite words, $f_\infty([\![a]\!]) = [\![a]\!] - [\![aB]\!]$ is interpreted as saying that the right-infinite words starting with a are taken homeomorphically by ϕ^{-1} to the set of right-infinite words starting with a, except those that start with aB.

Also, from Figure 1, we see that the subtrees T'_a and T'_b do not have any removable vertices, and therefore the slices $\mathcal{C}_a$ and $\mathcal{C}_b$ of the core $\mathcal{C}(T\phi \times T)$ are the respective subtrees with the 'half' edges removed. Thus the slice $\mathcal{C}_b$ is the edge bA and the slice $\mathcal{C}_a$ is the vertex $a\tilde{*}$. This is not the complete description of the core in this case; there are some vertical edges $(v \times e)$ that can be found by examining a homotopy inverse to σ'. This does however give us all of the 2-cells in the core, and hence we see $i(T\phi, T) = 1$.

EXAMPLE 3.10. Let $T \in \mathrm{cv}_3$ be the tree with all edge lengths equal to 1 and such that $\Gamma = T/F_3$ is the 3-rose which has petals that are labeled a, b and c. In this example we use the above algorithm to find the map on ends $f_\infty : \partial T \to \partial T\phi$ and build the slices of the core $\mathcal{C}(T \times T\phi)$ for the automorphism $\phi \in \mathrm{Aut}(F_3)$ given by $\phi(a) = baC, \phi(b) = cA, \phi(c) = a$, where $F_3 = \langle a, b, c\rangle$; here we use capital letters to denote the inverses of the generators. The inverse of ϕ is the map $\phi^{-1}(a) = c, \phi^{-1}(b) = ab, \phi^{-1}(c) = bc$. There is an obvious homotopy equivalence $\sigma' : \Gamma \to \Gamma$ representing ϕ^{-1}. Subdivide the edges of Γ as in Example 3.2. We can assume that $\sigma'(p_a) = p_c$. Further subdivide $b_+ = b_1b_2$ and $b_- = b_3b_4$ with the subdivision points p_{b_+}, p_{b_-} such that $\sigma'(p_{b_+}) = p_a$ and $\sigma'(p_{b_-}) = p_b$. Likewise subdivide both $c_+ = c_1c_2$ and $c_- = c_3c_4$ with the subdivision points p_{c_+} and p_{c_-} such that $\sigma'(p_{c_+}) = p_b$ and $\sigma'(p_{c_-}) = p_c$. Our preferred paths are $\alpha_a = a_+, \alpha_b = b_+$ and $\alpha_c = c_+$.

The preimages of p_a, p_b and p_c are $\Sigma_a = \{p_{b_+}\}$, $\Sigma_b = \{p_{b_-}, p_{c_+}\}$ and $\Sigma_c = \{p_a, p_{c_-}\}$, respectively. As $[\sigma'(b_1)] = [a_+] = \alpha_a$, we see that $f_\infty([\![a]\!]) = [\![b]\!]$. Then $[\sigma'(cAB_4)] = [bcCB_-] = b_+ = \alpha_b$ and $[\sigma'(c_1)] = [b_+] = \alpha_b$, and hence $f_\infty([\![b]\!]) = [\![c]\!] - [\![cAB]\!]$. Finally we see that $[\sigma'(a_+)] = [c_+] = \alpha_c$ and $[\sigma(aC_4)] = [cC_-] = c_+ = \alpha_c$, and hence $f_\infty([\![c]\!]) = [\![a]\!] - [\![aC]\!]$.

Now F_k-equivariance can be used to find the image of any other box, for example, we compute $f_\infty([\![B]\!])$. First note that $[\![B]\!] = B(\neg[\![b]\!])$, where $\neg[\![b]\!]$ denotes the complement of $[\![b]\!]$ in ∂T. Therefore $f_\infty([\![B]\!]) = f_\infty(B(\neg[\![b]\!])) = \phi(B)f_\infty(\neg[\![b]\!]) = aC([\![a]\!] + [\![A]\!] + [\![b]\!] + [\![B]\!] + [\![cAB]\!] + [\![C]\!]) = [\![aC]\!] + [\![B]\!]$.

From f_∞ we see that the slice $\mathcal{C}_a$ is empty ($a \times c$ is a twice-light rectangle), $\mathcal{C}_c$ is a single vertex and $\mathcal{C}_b$ is a single edge cA. Hence $i(T, T\phi) = 1$.

To get a more complicated example of the core, we look at $\mathcal{C}(T \times T\phi^3)$. Explicitly, ϕ^3 is the automorphism as follows:

$$\begin{aligned} a &\longmapsto acABcAbaCAcAB, \\ b &\longmapsto baCacABaC, \\ c &\longmapsto cAbaCA. \end{aligned}$$

The third power of the above map f_∞ is the map on ends for ϕ^3. To find this map, we can either use the algorithm with the automorphism ϕ^{-3}, or we can formally iterate f_∞; for example, to find $f_\infty^2([\![b]\!]) = f_\infty([\![c]\!] - [\![cAB]\!]) = f_\infty([\![c]\!]) - \phi(cA)f_\infty([\![B]\!]) = [\![a]\!] - [\![aC]\!] - acAB([\![aC]\!] + [\![B]\!]) = [\![a]\!] - [\![aC]\!] - [\![acaBaC]\!] - [\![acABB]\!]$. From either procedure we find

$$\begin{aligned} [\![a]\!] &\longmapsto [\![a]\!] - [\![aC]\!] - [\![acABB]\!] - [\![acABaC]\!], \\ f_\infty^3 : [\![b]\!] &\longmapsto [\![b]\!] - [\![baCA]\!] - [\![baCC]\!] - [\![baCacABaCB]\!] - [\![baCacABaCaC]\!], \\ &\qquad - [\![baCacABaCbaCC]\!] - [\![baCacABaCbaCA]\!], \\ [\![c]\!] &\longmapsto [\![c]\!] - [\![cAB]\!] - [\![cAbaCAC]\!] - [\![cAbaCAA]\!] - [\![cAbaCAcAB]\!]. \end{aligned}$$

Using this map and Lemma 3.7, we can build the slices of the core $\mathcal{C}(T \times T\phi^3)$. The sets $e \times \mathcal{C}_e$ for $e = a, b, c$ are displayed in Figure 3. The labeling of the edges and vertices is as follows: for each set $e \times \mathcal{C}_e$, the bottom left vertex is shown. For instance, for the set $a \times \mathcal{C}_a$, this vertex is $(*, a)$, where $*$ denotes the basepoint and a denotes the image of the basepoint by a. The rest of the edges are labeled by their image in the rose T/F_3. Hence the two vertices adjacent to $(*, a)$ are (a, a) and $(*, ac)$. Reading upward, the remaining vertices above $(*, a)$ are $(*, ac)$, $(*, acA)$, $(*, acAB)$ and $(*, acABa)$. Looking at the image for $f_\infty([\![a]\!])$, we see that the span of these vertices is the slice $\mathcal{C}_a$. Adding up the number of squares, we see that $i(T, T\phi^3) = 23$.

The colors and arrows denote the identifications; the thick black lines are free edges. For an example of the identifications let us look at the vertex on the bottom left; this vertex is $(*, a)$. Then $c(*, a) = (c, \phi^3(c)a) = (c, cAbaC)$, which is the fifth vertex from the bottom on the right. The identifications for the other vertices are found similarly. The observant reader will note that $\chi(\mathcal{C}(T \times T\phi^3)) = -2$, as expected.

4. *Bounding consolidation*

When the morphism $f : T' \to T$ used in Section 3 is the lift of a train-track representative, we are able to control the difference in volume between T'_e and $\mathcal{C}_e$ by bounding the amount of consolidation that occurs when removing removable vertices. Any train-track map satisfies the hypothesis used in Subsection 3.1 that the image of a small open neighborhood of a vertex is not contained in an edge. We begin by showing that there is a bound on the depth of consolidation.

LEMMA 4.1. *Let $\sigma : \Gamma \to \Gamma$ be a train-track map for a fully irreducible automorphism $\phi \in \mathrm{Out}(F_k)$ and let $f : T \to T$ be a lift. Let e be an edge of T, let p_e be an interior point of e, let T_e^n be the subtree spanned by $(f^n)^{-1}(p_e)$ and let Y_e^n be the subtree found by iteratively removing removable vertices. There is a constant $C \geqslant 0$, independent of n, such that any subtree A of T_e^n that is removed in the formation of Y_e^n has volume less than C.*

Proof. Without loss of generality, we can assume that A is a maximal subtree that is removed. We will show that the diameter of A is bounded; since the length of an edge in T

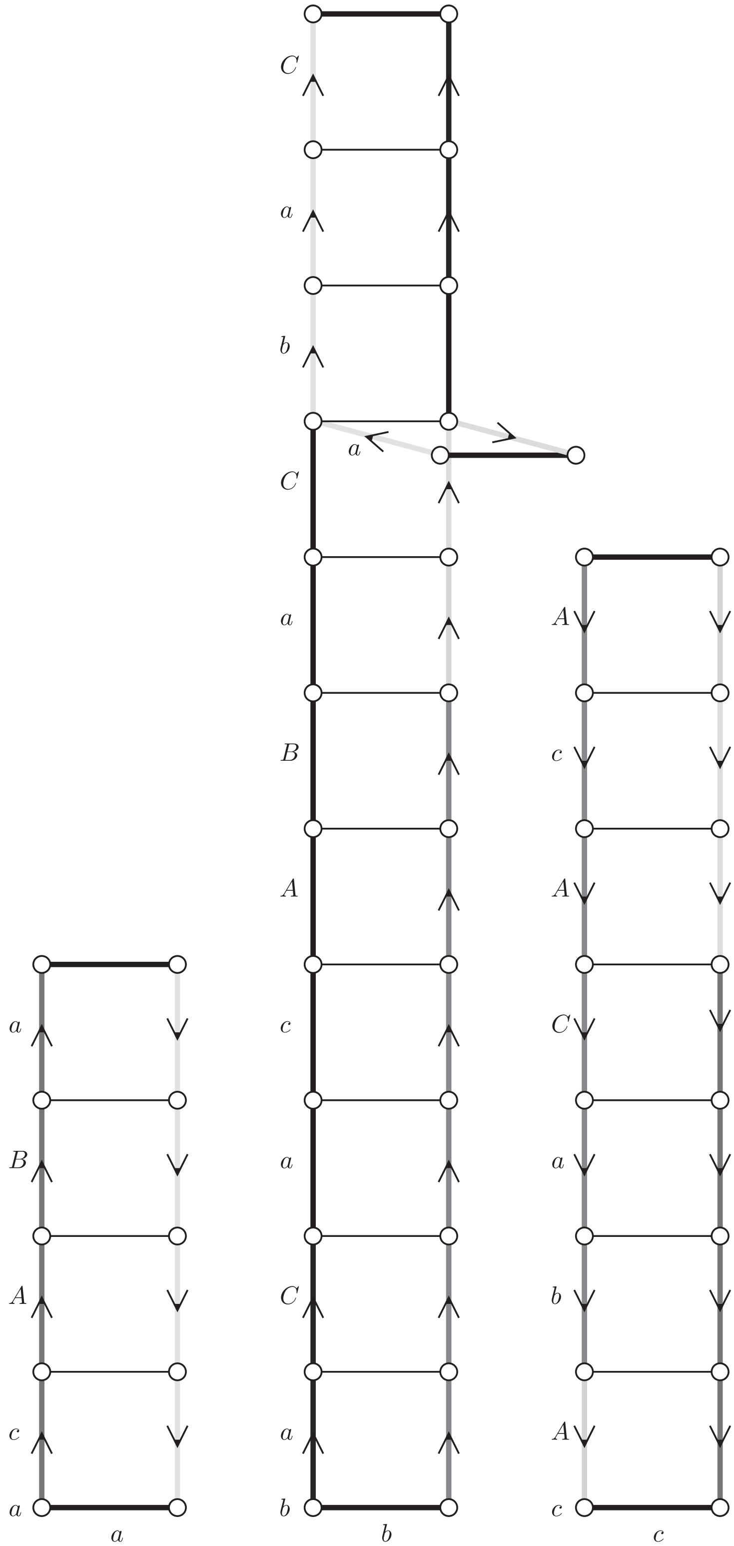

FIGURE 3. *The core* $\mathcal{C}(T \times T\phi^3)$ *from Example* 3.10.

is bounded and the valence of any vertex in T is bounded, this implies that the volume of A is bounded.

For any interior vertex $v \in A$, label the edges incident to v by $\hat{e}, e_1, \dots, e_k$, where $\hat{e}$ is the edge that leads to the unique terminal vertex of A which is not a terminal vertex of T_e^n. Any geodesic ray originating at v and containing one of the e_i contains exactly one point from $(f^n)^{-1}(p_e)$ by Lemma 3.5.

We claim that the only legal turns are contained in the edge paths $\bar{e}_i\hat{e}$ for $i = 1, \dots, k$. A turn contained in $\bar{e}_i e_j$ is illegal since f^n identifies the germs of the edges e_i and e_j. Since every edge must be in some legal turn, we see that the turns in $\bar{e}_i\hat{e}$ are legal.

Therefore, the geodesic containing any two terminal vertices of A has a single illegal turn. Write this geodesic as $\beta \cdot \gamma$, where β and γ are legal paths. If the length of $\beta \cdot \gamma$ is more than twice the critical constant for f, then at least one of β or γ has length bounded below by the critical constant. Therefore the length of $[f^m(\beta \cdot \gamma)]$ goes to infinity. But since $[f^m(\beta \cdot \gamma)]$ is a point for $m \geqslant n$, we see that the diameter of A is bounded by twice the critical constant. □

PROPOSITION 4.2. *Let $\sigma : \Gamma \to \Gamma$ be a train-track map for a fully irreducible automorphism $\phi \in \mathrm{Out}(F_k)$ and let $f : T \to T$ be a lift. Let e be an edge of T, let p_e an interior point of e, let T_e^n be the subtree spanned by $(f^n)^{-1}(p_e)$ and let $\mathcal{C}_e^n$ be the slice of the core $\mathcal{C}(T \times T\phi^n)$ above e. Then there exist constants $K \geqslant 1$ and $C \geqslant 0$ such that, for any $n \geqslant 0$, we have*

$$\mathrm{vol}(\mathcal{C}_e^n) \leqslant \mathrm{vol}(T_e^n) \leqslant K\,\mathrm{vol}(\mathcal{C}_e^n) + C.$$

Proof. As before, denote by Y_e^n the subtree of T_e^n obtained by iteratively removing removable vertices and by Z_e^n the subtree of interior edges of Y_e^n. Then, by Lemma 2.3, we have that $Z_e^n = \mathcal{C}_e^n$. Let M and m denote the length of the longest and shortest edges of Γ, respectively, and let b denote the maximum valence of any vertex in Γ. Then, by adding at most b edges of length M to each vertex of Z_e^n, we can cover Y_e^n. Since Z_e^n has at most $\mathrm{vol}(Z_e^n)/m + 1$ vertices, this shows that there are constants $K_1 \geqslant 1$ and $C_1 \geqslant 0$ such that

$$\mathrm{vol}(\mathcal{C}_e^n) \leqslant \mathrm{vol}(Y_e^n) \leqslant K_1\,\mathrm{vol}(\mathcal{C}_e^n) + C_1.$$

Therefore we only need to show that we can find constants K_2 and C_2 such that $\mathrm{vol}(T_e^n) \leqslant K_2\,\mathrm{vol}(Y_e^n) + C_2$.

By Lemma 4.1, we see that Y_e^n is obtained from T_e^n by removing disjoint subtrees, all of which have volumes that are bounded by a constant K_3. Since the maximum number of these trees is $b(\mathrm{vol}(Y_e^n)/m + 1)$ we see that

$$\mathrm{vol}(T_e^n) \leqslant \mathrm{vol}(Y_e^n) + K_3 b\left(\frac{\mathrm{vol}(Y_e^n)}{m} + 1\right) = \left(\frac{K_3 b}{m} + 1\right)\mathrm{vol}(Y_e^n) + K_3 b.$$

Hence the proposition follows. □

Combining the above results, we get a way to estimate the growth rate of the intersection number for a fully irreducible automorphism.

PROPOSITION 4.3. *Let ϕ be a fully irreducible automorphism and let $f : T \to T$ be a lift of a train-track representative for ϕ. Then, for any edge $e \subset T$ and trees $T', T'' \in \mathrm{cv}_k$, we have*

$$i(T', T''\phi^n) \sim \mathrm{vol}(T_e^n),$$

where p_e is the midpoint of e and $T_e^n \subset T$ is the subtree spanned by the points in $(f^n)^{-1}(p_e)$.

Proof. By Proposition 4.2, we have $\mathrm{vol}(T_e^n) \sim \mathrm{vol}(\mathcal{C}_e^n)$ where $\mathcal{C}_e^n$ is the slice of $\mathcal{C}(T \times T\phi^n)$ above e. Since ϕ is fully irreducible, it follows that $\mathrm{vol}(T_e^n) \sim \mathrm{vol}(T_{e'}^n)$ for any edges $e, e' \subset T$. (This becomes clear in Section 5, where we compute $\mathrm{vol}(T_e^n)$ up to $\sim$ equivalence independent of the edge e.) Thus $\mathrm{vol}(T_e^n) \sim \sum_{e' \subset T/F_k} l_T(e')\,\mathrm{vol}(\mathcal{C}_{e'}^n) = i(T, T\phi^n) \sim i(T', T''\phi^n)$. The final equivalence is from Proposition 2.9. □

Using the techniques developed thus far, we can get a lower bound on the asymptotics of $n \mapsto i(T, T'\phi^n)$.

COROLLARY 4.4. *Suppose that ϕ is a fully irreducible automorphism and let λ and μ denote the growth rates of ϕ and ϕ^{-1}, respectively. Then, for any $T, T' \in \mathrm{cv}_k$, there exist constants $K \geqslant 1$ and $C \geqslant 0$ such that $(1/K)\max\{\lambda, \mu\}^n - C \leqslant i(T, T'\phi^n)$.*

Proof. By Proposition 4.3 we only need to find a lower bound for $\mathrm{vol}(T_e^n)$, where T_e^n is the span of $(f^n)^{-1}(p_e)$ for a lift of a train-track map $f : T \to T$. Since ϕ is fully irreducible and f uniformly expands edges by λ, it follows that the cardinality of $(f^n)^{-1}(p_e)$ is $\sim \lambda^n$. Since each point of $(f^n)^{-1}(p_e)$ lies in a unique edge of T_e^n, there exist constants $K \geqslant 1$ and $C \geqslant 0$ such that $(1/K)\lambda^n - C \leqslant i(T, T\phi^n)$. Repeating this argument for $f' : T' \to T'$, a lift of a train-track map representing ϕ^{-1}, we see that there are constants $K' \geqslant 1$ and $C' \geqslant 0$ such that $(1/K')\mu^n - C' \leqslant i(T', T'\phi^{-n})$. Since $i(T', T'\phi^n) \sim i(T, T\phi^{-n})$, the result follows. □

5. *Growth rates*

In this section we compute the growth rate of $n \mapsto \mathrm{vol}(T_e^n)$, where T_e^n is as in Proposition 4.3. There are two cases depending on whether or not T^+, the stable tree of ϕ, is geometric. Every fully irreducible automorphism admits a train-track representative (called *stable*) with at most one orbit of indivisible Nielsen paths (this was proved in [**4**, Lemma 3.2] in the case of period 1 but the proof in general is identical).

Moreover, the existence of an indivisible Nielsen path characterizes whether T^+ is geometric.

THEOREM 5.1 [**2**, Theorem 3.2]. *Let $\phi \in \mathrm{Out}(F_k)$ be a fully irreducible automorphism, let $\sigma : \Gamma \to \Gamma$ be a stable train-track map and let T^+ be the stable tree for ϕ. Then T^+ is geometric if and only if $\sigma : \Gamma \to \Gamma$ contains an indivisible orbit of Nielsen paths.*

We shall not make use of any of the additional properties that stableness of a train-track map guarantees, with the exception of Proposition 5.20. These properties are mentioned within this proposition; see [**4**] for a definition. The following definition is of central importance to understanding the tree T_e^n from Proposition 4.3.

DEFINITION 5.2. An *i-vanishing path for $\sigma : \Gamma \to \Gamma$* $(i \geqslant 0)$ is an immersion $\iota : [0,1] \to \Gamma$ such that the image $\sigma^i\iota([0,1])$ is homotopic to a point relative to the endpoints. We always assume that i is minimal. If $f : T \to T$ is a lift of $\sigma : \Gamma \to \Gamma$, then an *$i$-vanishing path* is an embedding $\iota : [0,1] \to T$ such that $f^i\iota(0) = f^i\iota(1)$. Clearly any i-vanishing path in T projects to a i-vanishing path in Γ and vice versa. A *vanishing path* is an i-vanishing path for some i.

The importance of vanishing paths is given by the following remark.

REMARK 5.3. Suppose that $\sigma : \Gamma \to \Gamma$ is a map and $f : T \to T$ is a lift. Let T_e^n be the subtree spanned by $(f^n)^{-1}(p_e)$ for any edge $e \subset T$. Then T_e^n can be expressed as a union of i-vanishing paths, where $i \leqslant n$. In fact, the path joining any two points in $(f^n)^{-1}(p_e)$ is an i-vanishing path for some $i \leqslant n$.

Therefore we are interested in examining the lengths of vanishing paths. As we demonstrate in Subsection 5.2, every indivisible Nielsen path γ contains subpaths γ_ϵ that are i-vanishing, for arbitrarily large i. Note that all such vanishing paths have uniformly bounded length. In Subsection 5.1 we will prove a converse to this observation in the absence of indivisible Nielsen paths. Specifically, we show the following conditions.

(1) If T^+ is geometric, then every vanishing path is a composition of suitable γ_ϵ; see Proposition 5.22.

(2) If T^+ is nongeometric, then there exist arbitrarily long vanishing paths that are not compositions of shorter vanishing paths. In fact, any i-vanishing path has length approximately μ^i, where μ is the expansion factor for ϕ^{-1}; see Proposition 5.15.

Example 5.24 shows these different behaviors. We begin by examining the case when T^+ is nongeometric.

5.1. *T^+ nongeometric*

CONVENTION 5.4. We fix some notation for use in the rest of this section: ϕ is a fully irreducible automorphism; $\sigma : \Gamma \to \Gamma$ is a train-track map for ϕ with an expansion factor λ; $\sigma' : \Gamma' \to \Gamma'$ is a train-track map for ϕ^{-1} with an expansion factor μ, and $\tau : \Gamma \to \Gamma'$ and $\tau' : \Gamma' \to \Gamma$ are Lipschitz homotopy equivalences representing the change in marking. This is summarized in the following commutative (up to homotopy) diagram.

$$\begin{array}{ccccccccc} \cdots & \xrightarrow{\sigma} & \Gamma & \xrightarrow{\sigma} & \Gamma & \xrightarrow{\sigma} & \Gamma & \xrightarrow{\sigma} & \cdots \\ & & \tau \downarrow \uparrow \tau' & & \tau \downarrow \uparrow \tau' & & \tau \downarrow \uparrow \tau' & & \\ \cdots & \xleftarrow{\sigma'} & \Gamma' & \xleftarrow{\sigma'} & \Gamma' & \xleftarrow{\sigma'} & \Gamma' & \xleftarrow{\sigma'} & \cdots \end{array} \tag{5.1}$$

As stated above, we will show that when Γ does not contain an indivisible orbit of periodic Nielsen paths, the length of an i-vanishing path for σ is approximately μ^i. First we give an upper bound on the length of an i-vanishing path. The following proposition does not depend on the absence of Nielsen paths.

PROPOSITION 5.5. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4. There exists a constant $K \geqslant 0$ such that, for every i-vanishing path γ for $\sigma : \Gamma \to \Gamma$, we have* $\text{length}(\gamma) \leqslant K\mu^i$.

Proof. Since the composition $\sigma'\tau\sigma$ is homotopic to τ and the composition $\tau'\tau$ is homotopic to the identity map, there are constants $K_1, K_2 \geqslant 1$, such that any path γ in Γ satisfies

$$|\text{length}([\tau(\gamma)]) - \text{length}([\sigma'\tau\sigma(\gamma)])| \leqslant K_1, \tag{5.2}$$

$$|\text{length}(\gamma) - \text{length}([\tau'\tau(\gamma)])| \leqslant K_2. \tag{5.3}$$

For a 1-vanishing path γ, using the inequality (5.2) we have

$$\text{length}([\tau(\gamma)]) \leqslant \text{length}([\sigma'\tau\sigma(\gamma)]) + K_1 = K_1$$

as $\text{length}([\sigma(\gamma)]) = 0$.

We now proceed by induction. Suppose that the image under τ of any $(i-1)$-vanishing path has length at most $K_1 \sum_{j=1}^{i-1} \mu^{j-1}$; note that the previous paragraph verified the base case, a 1-vanishing path. Now consider an i-vanishing path γ. Then $\sigma(\gamma)$ is an $(i-1)$-vanishing path, and thus the inductive hypothesis yields $\mathrm{length}([\tau\sigma(\gamma)]) \leqslant K_1 \sum_{j=1}^{i-1} \mu^{j-1}$. Since σ' is a train-track map, this implies that $\mathrm{length}([\sigma'\tau\sigma(\gamma)]) \leqslant \mu K_1 \sum_{j=1}^{i-1} \mu^{j-1}$. Then (5.2) implies that $\mathrm{length}([\tau(\gamma)]) \leqslant \mathrm{length}([\sigma'\tau\sigma(\gamma)]) + K_1$, thus yielding the following which completes our induction:

$$\mathrm{length}([\tau(\gamma)]) \leqslant \mu K_1 \sum_{j=1}^{i-1} \mu^{j-1} + K_1 = K_1 \sum_{j=1}^{i} \mu^{j-1}.$$

Therefore $\mathrm{length}([\tau'\tau(\gamma)]) \leqslant K_3 K_1 \sum_{j=1}^{i} \mu^{j-1}$, where K_3 is the Lipschitz constant for τ'. As before, using (5.3) we have

$$\mathrm{length}(\gamma) \leqslant K_3 K_1 \sum_{j=1}^{i} \mu^{j-1} + K_2.$$

Setting $K = K_3 K_1(1/\mu - 1) + K_2$ completes the proof. □

Using Proposition 5.5, we can estimate the difference between the homotopic maps τ and $\sigma'^n\tau\sigma^n$. Again, the following lemma does not depend on the absence of Nielsen paths.

LEMMA 5.6. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention* 5.4. *Then there is a constant $K \geqslant 0$ such that, for any path $\gamma \subset \Gamma$ and $n \geqslant 0$, we have*

$$\mathrm{length}([\tau(\gamma)]) \geqslant \mathrm{length}([\sigma'^n\tau\sigma^n(\gamma)]) - K\mu^n.$$

Proof. Since τ is homotopic to $\sigma'^n\tau\sigma^n$, if γ is a closed loop, then $\mathrm{length}([\tau(\gamma)]) = \mathrm{length}([\sigma'^n\tau\sigma^n(\gamma)])$.

If γ is not a loop, then we can add a segment α of bounded length to $[\sigma^n(\gamma)]$ to get a closed loop γ_1 such that $[\sigma^n(\gamma)]$ is a subpath of γ_1 and α is an embedded arc. Let γ_0 be a closed loop in Γ such that $[\sigma^n(\gamma_0)] = \gamma_1$. As α has bounded length, there is a constant K_1 such that $\mathrm{length}([\sigma'^n\tau(\alpha)]) \leqslant K_1\mu^n$. Hence we have

$$\begin{aligned}\mathrm{length}([\tau(\gamma_0)]) &= \mathrm{length}([\sigma'^n\tau(\gamma_1)])\\ &\geqslant \mathrm{length}([\sigma'^n\tau\sigma^n(\gamma)]) - \mathrm{length}([\sigma'^n\tau(\alpha)])\\ &\geqslant \mathrm{length}([\sigma'^n\tau\sigma^n(\gamma)]) - K_1\mu^n.\end{aligned}$$

There is a path β in Γ (unique up to homotopy) such that the concatenation of γ and β is homotopic to γ_0. Note that $[\sigma^n(\beta)] = \alpha$ as $[\sigma^n(\gamma)]$ is a subpath of γ_1. Now $\mathrm{length}([\tau(\gamma)]) \geqslant \mathrm{length}([\tau(\gamma_0)]) - \mathrm{length}([\tau(\beta)])$. We will show that β is the union of a bounded number of i-vanishing paths, where $i \leqslant n$ along with some segments of bounded length.

Let x and y be the endpoints of α, and let p and q be the endpoints of β with $\sigma^n(p) = x$ and $\sigma^n(q) = y$, respectively. Now decompose $\beta = \beta_0 \cdot \beta_1$ by subdividing β at the point in $(\sigma^n)^{-1}(y)$ that is closest (along β) to x. Then β_1 is an i-vanishing path for σ, with $i \leqslant n$, as α is embedded and $[\sigma^n(\beta)] = \alpha$. Thus $[\sigma^n(\beta_0)] = \alpha$. Similarly, decompose $\beta_0 = \beta_2 \cdot \beta_3$, where β_2 is an i'-vanishing path for σ with $i' \leqslant n$ and where $[\sigma^n(\beta_2)] = x$. Thus $[\sigma^n(\beta_3)] = \alpha$ and $(\sigma^n)^{-1}(x)$ and $(\sigma^n)^{-1}(y)$ only intersect β_3 in its endpoints. Now repeat at the vertices contained in α to decompose β_3 as a union of vanishing paths (the number of which is bounded by the number of vertices of α) connected by segments that are homeomorphically mapped to the edges of α. The length and number of such segments are bounded. Hence, by Proposition 5.5

and since τ induces a Lipschitz map between the universal covers of Γ and Γ', there is a constant K_2 such that length$([\tau(\beta)]) \leqslant K_2\mu^n$. Setting $K = K_1 + K_2$ completes the proof. □

To show the correct lower bound on the length of an i-vanishing path for $\sigma : \Gamma \to \Gamma$, we need the following lemma from [**3**].

LEMMA 5.7 [**3**, Lemma 2.9]. *Let $\sigma : \Gamma \to \Gamma$ be a train-track map representing a fully irreducible outer automorphism ϕ. Then, for every $C > 0$, there is a number $M > 0$ such that if γ is any path, then one of the following holds:*
(1) *$[\sigma^M(\gamma)]$ contains a legal segment of length greater than C;*
(2) *$[\sigma^M(\gamma)]$ has fewer illegal turns than γ;*
(3) *γ is a concatenation $x \cdot y \cdot z$ with $[\sigma^M(y)]$ (periodic) Nielsen and x and z have length at most $2C$ and at most one illegal turn.*

Using the above, we can show that when Γ does not contain an orbit of periodic Nielsen paths then short vanishing paths for σ vanish quickly. First, we need to understand paths in Γ that could satisfy conclusion (3) of Lemma 5.7.

LEMMA 5.8. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4 and suppose that Γ does not contain an indivisible orbit of periodic Nielsen paths. Then there exists an $N > 0$ such that, if γ is an i-vanishing path for $\sigma : \Gamma \to \Gamma$ with at most two illegal turns, then $i \leqslant N$.*

Proof. Suppose otherwise, therefore we have a sequence of i_j-vanishing paths γ_j with at most two illegal turns, where $i_j < i_{j+1}$. We will show that this implies that some power of σ has a Nielsen path, which contradicts our assumption.

First note that the lengths of γ_j are uniformly bounded, since legal subpaths of vanishing paths have length bounded by the critical constant. By passing to a subsequence, we can assume that the paths γ_j all have the same combinatorial type, that is, they cross the same turns of Γ in the same order. Further, by passing to a power of σ, possibly replacing the sequence γ_j by $\sigma^\ell(\gamma_j)$ for some ℓ, such that the vertex (or vertices) at the illegal turn is (are) fixed by σ, and again passing to a subsequence, we can assume that $[\sigma(\gamma_j)] \subseteq \gamma_j$.

We break the proof up into two cases depending on whether the paths γ_j contain a single illegal turn or two illegal turns. First, assume that the paths γ_j only contain a single illegal turn. Note that any i-vanishing path with a single illegal turn has the form $a \cdot b$, where a and b are legal segments of the same length such that $\sigma^i(a) = \sigma^i(\bar{b})$. Further any subpath of the form $a' \cdot b'$, where a' and b' are legal segments of the same length, is an i'-vanishing path where $i' \leqslant i$.

We have two claims about the vanishing paths γ_j.

CLAIM 1. *If ι is a subpath of γ_j and ι is an i-vanishing path for $\sigma : \Gamma \to \Gamma$, then $i \leqslant i_j$.*

Proof. If $[\sigma^{i_j}(\iota)]$ is not a point, then it is a path that does not contain any illegal turns (as there is a single illegal turn in γ_j) and hence ι is not a vanishing path. □

CLAIM 2. *If $j < k$, then γ_j is a subpath of γ_k.*

Proof. Since the two legal segments in each of γ_j and γ_k have equal lengths, we must have $\gamma_j \subset \gamma_k$ or $\gamma_k \subset \gamma_j$ because they have the same combinatorial type. The latter is not possible by Claim 1. □

Therefore, by the above claims, we have $\gamma_j \subset \gamma_{j+1}$ and hence there is a well-defined limit γ_∞. By construction, we have $\gamma_\infty = a \cdot b$, where a and b are legal segments of the same length as this holds for each of the γ_j. Further $[\sigma(a)]$ and $[\sigma(b)]$ are also legal segments of the same length and $[\sigma(\gamma_\infty)] \subseteq [\gamma_\infty]$. If $[\sigma(\gamma_\infty)] \neq \gamma_\infty$, then $[\sigma(\gamma_\infty)] \subset \gamma_{j'}$ for some j' and as $[\sigma(a)]$ and $[\sigma(b)]$ are also legal segments of the same length, $[\sigma(\gamma_\infty)]$ is a i'-vanishing path for some $i' \leqslant i$. However, since $\gamma_j \subset \gamma_\infty$ are all i_j-vanishing paths with $i_j \to \infty$, by Claim 1, the path γ_∞ is not a vanishing path. Therefore $[\sigma(\gamma_\infty)] = \gamma_\infty$, and hence γ_∞ is a Nielsen path.

It remains to consider the case when the γ_j have two illegal turns. We argue that this is not possible. Suppose that γ is a vanishing path with two illegal turns and $[\sigma(\gamma)] \subset \gamma$. If $[\sigma(\gamma)]$ has a single illegal turn, then we are in the case above. Otherwise, the middle legal segment b of $\gamma = a \cdot b \cdot c$ maps over itself and therefore has a unique fixed point, breaking it up as $b = b_1 b_2$. There are two subcases depending on whether b maps over itself, preserving or reversing the orientation; the cases are similar and we assume that the orientation is preserved. Therefore the iterates of $a \cdot b_1$ and of $b_2 \cdot c$ never cancel each other, and so both must be vanishing; but $[\sigma(a \cdot b_1)]$ has the form $a' \cdot b_1$ for some $a' \subset a$ (and both ab_1 and $a'b_1$ are vanishing). This contradicts the fact that vanishing paths with one illegal turn have legal segments of equal lengths.

This completes the proof of the lemma. □

Lemma 5.9. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention* 5.4 *and suppose that Γ does not contain an indivisible orbit of periodic Nielsen paths. For every $L > 0$ there exists an $N > 0$, such that if γ is an i-vanishing path for σ and length$(\gamma) \leqslant L$, then $i \leqslant N$.*

Proof. As there is a lower bound on the distance between any two illegal turns in Γ, for any vanishing path with length less than L, the number of illegal turns in a path of length at most L is at most some constant I.

Let C be larger than the critical constant for σ and let M be the constant from Lemma 5.7. If $[\sigma^M(\gamma)]$ had a legal segment of length greater than C, this would imply that $[\sigma^n(\gamma)]$ had positive length for all $n > M$, and thus that γ is not a vanishing path, contrary to hypothesis. Thus possibility (1) of Lemma 5.7 cannot occur for any vanishing path. On the other hand, if possibility (3) of Lemma 5.7 occurs, then y is trivial since by assumption Γ does not contain an indivisible orbit of periodic Nielsen paths; hence $\gamma = x \cdot z$ and has length less than $2C$ and at most two illegal turns.

Thus the possibility (2) of Lemma 5.7 must occur for every vanishing path with more than two illegal turns or length greater than $2C$, namely, $[\sigma^M(\gamma)]$ has strictly fewer illegal turns than γ. Hence $[\sigma^{IM}(\gamma)]$ has at most two illegal turns and length less than $2C$ for any vanishing path γ. Therefore, we see that $[\sigma^{N+IM}(\gamma)]$ is a point, where N is the constant from Lemma 5.8 and hence $i \leqslant N + IM$. □

To derive the more precise lower bound on the length of an i-vanishing path, we consider two types of *legality*. For a path γ with I illegal turns, we define the ratio

$$\mathrm{LEG}_1(\gamma) = \frac{\mathrm{length}(\gamma) - 2\lambda^{-1}BI}{\mathrm{length}(\gamma)}$$

where $B = \mathrm{BCC}(\sigma)$ is the bounded cancellation constant. Since $\mathrm{length}([\sigma(\gamma)]) \geqslant \lambda\,\mathrm{length}(\gamma) - 2BI$, we see that $\mathrm{length}([\sigma(\gamma)]) \geqslant \lambda\,\mathrm{LEG}_1(\gamma)\,\mathrm{length}(\gamma)$. The following lemma and corollary do not depend on the absence of Nielsen paths.

Lemma 5.10. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4. For any path $\gamma \subset \Gamma$, with $\mathrm{LEG}_1(\gamma) > \lambda^{-1}$, the ratio $\mathrm{LEG}_1([\sigma^n(\gamma)])$ converges to 1. Moreover, for any $0 < \epsilon < 1 - \lambda^{-1}$, there is a $\delta > 0$ such that if $\mathrm{LEG}_1(\gamma) \geqslant \lambda^{-1} + \epsilon$, then the infinite product $\prod_{n=0}^{\infty} \mathrm{LEG}_1([\sigma^n(\gamma)])$ converges to a positive number greater than or equal to δ.*

Proof. Since the number of illegal turns in $[\sigma(\gamma)]$ is at most the number of illegal turns in γ and $\mathrm{length}([\sigma(\gamma)]) \geqslant \lambda\,\mathrm{length}(\gamma) - 2BI$, we see that

$$\frac{1 - \mathrm{LEG}_1([\sigma(\gamma)])}{1 - \mathrm{LEG}_1(\gamma)} \leqslant \frac{2\lambda^{-1}BI}{\lambda\,\mathrm{length}(\gamma) - 2BI}\left(\frac{2\lambda^{-1}BI}{\mathrm{length}(\gamma)}\right)^{-1}$$
$$= \frac{\mathrm{length}(\gamma)}{\lambda\,\mathrm{length}(\gamma) - 2BI} = \frac{1}{\lambda\,\mathrm{LEG}_1(\gamma)} < 1.$$

Hence $\mathrm{LEG}_1([\sigma(\gamma)]) > \mathrm{LEG}_1(\gamma) > \lambda^{-1}$ and so repeating the above, we see that $\mathrm{LEG}_1([\sigma^n(\gamma)])$ converges to 1. By bounding $\mathrm{LEG}_1(\gamma)$ away from λ^{-1}, we can make the rate of convergence independent of the path γ. Further the above calculation shows that

$$\limsup_{n\to\infty} \frac{1 - \mathrm{LEG}_1([\sigma^{n+1}(\gamma)])}{1 - \mathrm{LEG}_1([\sigma^n(\gamma)])} \leqslant \lambda^{-1}.$$

Hence the product $\prod_{n=1}^{\infty} \mathrm{LEG}_1([\sigma^n(\gamma)])$ converges to a positive number, which can be bounded away from 0 independent of the path γ by bounding $\mathrm{LEG}_1(\gamma)$ away from λ^{-1}. □

The following corollary follows immediately.

Corollary 5.11. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4. For any $0 < \epsilon < 1 - \lambda^{-1}$, there is a constant $K > 0$ such that, for any $n \geqslant 0$ and any path $\gamma \subset \Gamma$ with $\mathrm{LEG}_1(\gamma) \geqslant \lambda^{-1} + \epsilon$, we have*

$$\mathrm{length}([\sigma^n(\gamma)]) \geqslant K\lambda^n\,\mathrm{length}(\gamma).$$

As a word of caution, we will be applying this corollary to paths in Γ' and the train-track map $\sigma' : \Gamma' \to \Gamma'$. Specifically, this corollary enables us to get a lower bound on the length of a path $\gamma \subset \Gamma$ when $[\tau(\sigma^n(\gamma))] \subset \Gamma'$ has large legality with respect to $\sigma' : \Gamma' \to \Gamma'$.

Lemma 5.12. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4. For any $0 < \epsilon < 1 - \mu^{-1}$, there exist constants $K > 0$ and $C, C' \geqslant 0$ such that, for any $n \geqslant 0$ and any path $\gamma \subset \Gamma$, if $\mathrm{LEG}_1([\tau\sigma^n(\gamma)]) \geqslant \mu^{-1} + \epsilon$ then*

$$\mathrm{length}(\gamma) \geqslant K\mu^n(\mathrm{length}([\tau\sigma^n(\gamma)]) - C) - C'.$$

Proof. Applying Corollary 5.11 to $\sigma' : \Gamma' \to \Gamma'$, we see that there is a constant $K_1 > 0$ such that, for any path $\gamma \subset \Gamma$ with $\mathrm{LEG}_1([\tau\sigma^n(\gamma)]) \geqslant \lambda^{-1} + \epsilon$,

$$\mathrm{length}([\sigma'^n\tau\sigma^n(\gamma)]) \geqslant K_1\mu^n\mathrm{length}([\tau\sigma^n(\gamma)]).$$

Combining this with Lemma 5.6, there is a constant $C \geqslant 0$ such that

$$\mathrm{length}([\tau(\gamma)]) \geqslant K_1\mu^n(\mathrm{length}([\tau\sigma^n(\gamma)]) - C).$$

Finally, since τ induces a quasi-isometry between the universal covers of Γ and Γ', there are constants $K > 0$ and $C' \geqslant 0$ such that

$$\text{length}(\gamma) \geqslant K\mu^n(\text{length}([\tau\sigma^n(\gamma)]) - C) - C'. \qquad \square$$

We now show that if γ is an i-vanishing path for σ, then $\text{LEG}_1([\tau(\gamma)])$ can be made close to 1 for large enough i, thus enabling us to use Lemma 5.12. To show this, we need to use the version of legality from [**3**]. Let C be larger than the critical constant for σ and the critical constant for σ' and define

$$\text{LEG}_2(\gamma) = \frac{\text{sum of the lengths of the legal segments of } \gamma \text{ of length } \geqslant C}{\text{length}(\gamma)}.$$

There is a constant η such that $\text{length}([\sigma^n(\gamma)]) \geqslant \eta\lambda^n \,\text{LEG}_2(\gamma)\,\text{length}(\gamma)$ for any path γ. We need the following lemma regarding this version of legality.

LEMMA 5.13 [**3**, Lemma 5.6]. *Assume that ϕ has no nontrivial periodic conjugacy classes (or equivalently that ϕ is not represented by a homeomorphism of a punctured surface). Then there are constants $L, N, \epsilon > 0$ such that, for any path γ with $\text{length}(\gamma) \geqslant L$ and every $n \geqslant N$, either $\text{LEG}_2([\sigma^n(\gamma)]) \geqslant \epsilon$ or $\text{LEG}_2([\sigma'^n\tau(\gamma)]) \geqslant \epsilon$.*

In [**3**], the above lemma is stated for conjugacy classes in F_k but the proof applies equally well to this setting.

LEMMA 5.14. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4 and suppose that Γ does not contain an indivisible orbit of periodic Nielsen paths. For all $0 < \delta < 1$, there exist constants $L > 0$ and $N > 0$ such that if γ is a vanishing path for σ with length $\geqslant L$, then $\text{LEG}_1([\sigma'^n\tau(\gamma)]) \geqslant \delta$ for all $n \geqslant N$.*

Proof. Let L_1, N_1 and ϵ be constants from Lemma 5.13. Since γ is a vanishing path for σ, we have $\text{LEG}_2([\sigma^n(\gamma)]) = 0$ for all $n \geqslant 0$ and hence we must have the second conclusion from this lemma, namely, $\text{LEG}_2([\sigma'^{N_1}\tau(\gamma)]) \geqslant \epsilon$ for vanishing paths γ with length at least L.

Therefore there is a constant $\eta > 0$ such that, for any vanishing path with length at least L, we have

$$\text{length}([\sigma'^n\tau(\gamma)]) \geqslant \eta\epsilon\mu^{n-N_1}\,\text{length}([\sigma'^{N_1}\tau(\gamma)])$$

for $n \geqslant N_1$. Further, there is a constant I such that the number of illegal turns in $\sigma'^n\tau(\gamma)$ is at most $I\,\text{length}([\sigma'^{N_1}\tau(\gamma)])$ for all $n \geqslant N_1$ (since legal turns go to legal turns and the distance between illegal turns is uniformly bounded from below). Hence, set $L = L_1$ and N large enough such that $1 - 2B\mu^{-1}I/\eta\epsilon\mu^{N-N_1} \geqslant \delta$. $\square$

Finally, we can prove that the i-vanishing paths have length approximately μ^i when there are no orbits of periodic Nielsen paths.

PROPOSITION 5.15. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4 and suppose that Γ does not contain a Nielsen path. There exist constants $K > 0$ and $C \geqslant 0$ such that, for every i-vanishing path γ for $\sigma : \Gamma \to \Gamma$, we have $\text{length}(\gamma) \geqslant K\mu^i - C$.*

Proof. Fix a small $0 < \epsilon < 1 - \mu^{-1}$, and let L_1 and N_1 be the constants from Lemma 5.14 using $\delta = \mu^{-1} + \epsilon$. Also let K_1 and C_1 and C_1' be the constants from Lemma 5.12 using $\epsilon/2$.

Let ℓ be large enough such that the τ-image of an ℓ-vanishing path has length at least $2C_1$ and such that an $(\ell - N_1)$-vanishing path has length at least L_1. Such an ℓ exists by Lemma 5.9.

Suppose that α is an ℓ-vanishing path for σ. Hence $\sigma^{N_1}(\alpha)$ has length at least L_1 and by Lemma 5.14 we have

$$\mathrm{LEG}_1([\sigma'^{N_1}\tau\sigma^{N_1}(\alpha)]) \geqslant \mu^{-1} + \epsilon.$$

Since τ is homotopic to $\sigma'^{N_1}\tau\sigma^{N_1}$, by further requiring that ℓ be sufficiently large, we can guarantee that the length of an ℓ-vanishing path is large enough such that

$$\mathrm{LEG}_1([\tau(\alpha)]) \geqslant \mathrm{LEG}_1([\sigma'^{N_1}\tau\sigma^{N_1}(\alpha)]) - \epsilon/2 \geqslant \mu^{-1} + \epsilon/2.$$

For $i > \ell$ we can apply the above to $\alpha = \sigma^{i-\ell}(\gamma)$, and by Lemma 5.12 we have

$$\mathrm{length}(\gamma) \geqslant K_1\mu^{i-\ell}(\mathrm{length}([\tau\sigma^{i-\ell}(\gamma)]) - C_1) - C_1'.$$

Since $\mathrm{length}([\tau\sigma^{i-\ell}(\gamma)]) \geqslant 2C_1$, we have $\mathrm{length}(\gamma) \geqslant K_1\mu^{i-\ell}C_1 - C_1'$. □

Recall that T_e^n is the span of the points in $(f^n)^{-1}(\widetilde{p_e})$, where $f : T \to T$ is a lift of $\sigma : \Gamma \to \Gamma$ and p_e is a lift of a point in the interior of e. The path joining any pair of points in $(f^n)^{-1}(p_e)$ is an i-vanishing path for some $i \leqslant n$. We define a set of equivalence relations on the set $(f^n)^{-1}(p_e)$ by $x \sim_i x'$ if the path connecting them is a j-vanishing path, where $j \leqslant i$. Equivalently, $x \sim_i x'$ if $f^i(x) = f^i(x')$ for $i \leqslant n$. Note that if $x \sim_j x'$ and $x' \sim_i x''$ for $j \leqslant i$, then $x \sim_i x''$. An $\sim_i$ equivalence class is called an *i-clump*. Note that the entire set $(f^n)^{-1}(p_e)$ is an n-clump. We shall also sometimes use the term i-clump to refer to the subtree of T_e^n spanned by the points in an individual i-clump.

Contained within an i-clump, there are several j-clumps, where $j < i$. Indeed, any two points in the i-clump connected by a j-vanishing path are part of the same j-clump. As the length of an ℓ-vanishing path is roughly μ^ℓ, if j is much smaller than ℓ, distinct j-clumps contained within an i-clump joined by an ℓ-vanishing path are far away from each other.

We fix *j-spanning trees* contained within the i-clumps by taking the subtree spanned by choosing one point from each j-clump within the i-clump. We shall estimate the volume of T_e^n by estimating the size and number of i-spanning trees contained in T_e^n.

LEMMA 5.16. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4 and let $f : T \to T$ be a lift of σ. With the notation above, for any $j < i \leqslant n$, the number of j-clumps contained in an i-clump of T_e^n is $\sim \lambda^{i-j}$.*

Proof. As $\sigma : \Gamma \to \Gamma$ is irreducible and linearly expands any edge by the factor λ, there exist constants $K_1 \geqslant 1$ and $C_1 \geqslant 0$ such that, for any point $x \in \Gamma$, we have $(1/K_1)\lambda^\ell - C_1 \leqslant$ the number of points in $\sigma^{-\ell}(x) \leqslant K_1\lambda^\ell + C_1$.

Now after applying f^j to T, any j-clump becomes a point and distinct j-clumps become distinct points. Further applying f^{i-j}, the images of the j-clumps which are contained within the i-clump all map to the same point x. Hence the j-clumps contained within the i-clump are parameterized by $\sigma^{-(i-j)}(x)$. □

In particular, the number of i-clumps contained in T_e^n is $\sim \lambda^{n-i}$.

LEMMA 5.17. *Let $\sigma : \Gamma \to \Gamma$ be as in Convention 5.4, let $f : T \to T$ be a lift of σ and suppose that Γ does not contain a Nielsen path. Fix an $\ell > 0$; then, for any $\ell < i \leqslant n$, the volume of a $(i - \ell)$-spanning tree contained in an i-clump is $\sim \mu^i$.*

Proof. As in Lemma 5.16, the number of $(i-\ell)$-clumps contained in an i-clump is bounded above and below by some constants independent of i. Thus any $(i-\ell)$-spanning tree is covered by a bounded number of j-vanishing paths, where $i-l<j\leqslant i$. By Propositions 5.5 and 5.15 the length of such a vanishing path is bounded above and below by linear functions of μ^i. □

Putting these together we can show the following proposition.

PROPOSITION 5.18. *Let $\sigma:\Gamma\to\Gamma$ be as in Convention 5.4, let $f:T\to T$ be a lift of σ and suppose that Γ does not contain a Nielsen path. Then, for any edge $e\subset T$ we have*

$$\mathrm{vol}(T_e^n)\sim\lambda^n+\lambda^{n-1}\mu+\ldots+\lambda\mu^{n-1}+\mu^n.$$

Proof. As T_e^n is covered by i-vanishing paths for $i\leqslant n$, we can also cover T_e^n using all of the fixed $(i-1)$-spanning trees contained in all of the i-clumps for $1\leqslant i\leqslant n$. The number of such $(i-1)$-spanning trees is equal to the number of i-clumps. Thus, putting together Lemmas 5.16 and 5.17, we get that $\mathrm{vol}(T_e^n)$ is bounded above by a linear function of $\lambda^n+\lambda^{n-1}\mu+\ldots+\lambda\mu^{n-1}+\mu^n$.

To get the lower bound, we need to look carefully at the overlap of these spanning trees. Let K be the constant appearing in Proposition 5.5 and choose ℓ such that μ^ℓ is much bigger than K. We proceed by induction on i and show that the volume of an i-clump is bounded below by a linear function of $\mu^i+\lambda^\ell\mu^{i-\ell}+\lambda^{2\ell}\mu^{i-2\ell}+\ldots$.

Fix an i-clump and let T_0 be an $(i-\ell)$-spanning tree contained within this i-clump. Let $j=i-\ell$ and denote the j-clumps contained within this i-clump by $T_1,\ldots,T_{M_j}$. By induction, the volume of any of the T_m is bounded below by a linear function of $\mu^j+\lambda^\ell\mu^{j-\ell}+\lambda^{2\ell}\mu^{j-2\ell}+\ldots$. Note that $M_j\approx\lambda^\ell$. For any T_m, the overlap of T_0 and T_m has volume bounded above by $K\mu^j$. This follows as the overlap is contained within a j-vanishing path. Therefore the volume of $\widehat{T_0}$, the subforest of T_0 obtained by removing any overlap with some T_m, has volume bounded below by a linear function of μ^i. As the i-clump equals the union of $\widehat{T_0}\cup T_1\cup\ldots\cup T_{M_j}$, applying induction this shows that the volume of an i-clump is bounded below by a linear function of $\mu^i+\lambda^\ell\mu^{i-\ell}+\lambda^{2\ell}\mu^{i-2\ell}+\ldots\sim\mu^i+\lambda\mu^{i-1}+\ldots+\lambda^{i-1}\mu+\lambda^i$. For $i=n$, we achieve the desired lower bound for $\mathrm{vol}(T_e^n)$. □

Putting together Propositions 4.3 and 5.18 with Theorem 5.1, we get the second case of Theorem 1 from the Introduction.

THEOREM 5.19. *Suppose that $\phi\in\mathrm{Out}(F_k)$ is a fully irreducible automorphism with a nongeometric stable tree. Let λ be the expansion factor of ϕ and let μ be the expansion factor of ϕ^{-1}. Then, for any $T,T'\in cv_k$, we have*

$$i(T,T'\phi^n)\sim\lambda^n+\lambda^{n-1}\mu+\ldots+\lambda\mu^{n-1}+\mu^n.$$

5.2. T^+ *geometric*

We now look at the case when T^+ is geometric. We need the following well-known result providing a special representative of an automorphism with a geometric stable tree.

PROPOSITION 5.20. *Suppose that T^+, the stable tree of a fully irreducible automorphism ϕ, is geometric. Then some power of ϕ has a train-track representative $\sigma:\Gamma\to\Gamma$ such that:*

(1) *σ has exactly one indivisible Nielsen path;*
(2) *σ has a unique illegal turn;*

(3) *σ is homotopic to the result of iteratively folding the indivisible Nielsen path along the illegal turn.*

Proof. Let $\sigma : \Gamma \to \Gamma$ be a stable train-track representative for ϕ. Then σ contains a unique indivisible orbit of periodic Nielsen paths, each path of which contains a unique illegal turn. This follows by [**4**, Lemmas 3.4 and 3.9]. Thus replacing σ by a power, we can achieve (1). Further, by [**4**, Lemma 3.9], after replacing σ by this power, $\sigma : \Gamma \to \Gamma$ has a unique illegal turn, and hence (2) holds.

Any homotopy equivalence between graphs factors into a sequence of folds [**28**]. Since edges in a legal turn do not get identified by σ, the first fold in the factored sequence is folding the indivisible Nielsen path. Folding a pair of edges of a stable train-track map results in a stable train-track map (remark on [**4**, p. 23]). Hence the factored sequence is iteratively folding the indivisible Nielsen path. □

We call a train-track representative satisfying the conclusions of the preceding proposition a *parageometric* train-track representative. For $t \geqslant 0$ let Γ_t be the graph resulting from Γ by iteratively folding the indivisible Nielsen path along the illegal turn at a constant rate, where, for $t \in \mathbb{Z}_{\geqslant 0}$, the induced map $\sigma_{0t} : \Gamma \to \Gamma_t$ is σ^t. For $0 \leqslant s \leqslant t$ let $\sigma_{st} : \Gamma_s \to \Gamma_t$ by the induced map.

The dichotomy in Theorem 1 stems from the following observation. Suppose that $\sigma : \Gamma \to \Gamma$ has an indivisible Nielsen path γ. For small $\epsilon > 0$ let γ_ϵ be the subpath of γ obtained by removing an ϵ neighborhood of its endpoints. Then γ_ϵ is an i-vanishing path for σ and, moreover, $i \to \infty$ as $\epsilon \to 0$. This observation follows as we can write γ as a concatenation of legal paths $\gamma = a_0 \cdot b_0 \cdot \overline{b_1} \cdot \overline{b_1}$, where $\sigma(a_i) = a_i \cdot b_i$ and $\sigma(b_0) = \sigma(b_1)$. Hence for large n (depending on ϵ), we have $[\sigma^n(\gamma_\epsilon)] \subset b_0 \cdot \overline{b_1}$ and therefore is an $(n+1)$-vanishing path. Therefore, in contrast with Lemma 5.9 and Proposition 5.15 from the nongeometric case, there are bounded length i-vanishing paths for large i. In fact, in the geometric case we claim that any i-vanishing path is a concatenation, tightened relative to its endpoints, of vanishing paths that are contained in Nielsen paths. A vanishing path contained in an indivisible Nielsen path is called a *special vanishing path*.

We begin with a simple criterion for finding special vanishing paths.

LEMMA 5.21. *Suppose that $\sigma : \Gamma \to \Gamma$ is a parageometric train-track representative. There is a constant $T > 0$ such that any vanishing path for $\sigma_{0t} : \Gamma \to \Gamma_t$, with $t \leqslant T$, is a special vanishing path.*

Proof. Choose T such that $\sigma_{0t} : \Gamma \to \Gamma_t$ does not identify a pair of vertices of Γ for $t \leqslant T$. Then any vanishing path of σ_{0t} is contained within a pair of edges that are partially folded together by σ_{0t}, and the turn between these edges is illegal. As σ has a unique illegal turn and this illegal turn is contained in the unique Nielsen path, any vanishing path for σ_{0t} is contained in the Nielsen path and hence is a special vanishing path. □

We can now show that vanishing paths can be covered by special vanishing paths.

PROPOSITION 5.22. *Suppose that $\sigma : \Gamma \to \Gamma$ is a parageometric train-track representative and γ is an i-vanishing path for σ. Then $\gamma = [\gamma_1 \ldots \gamma_\ell]$, where, for $j = 1, \ldots, \ell$, the path γ_j is a special vanishing path.*

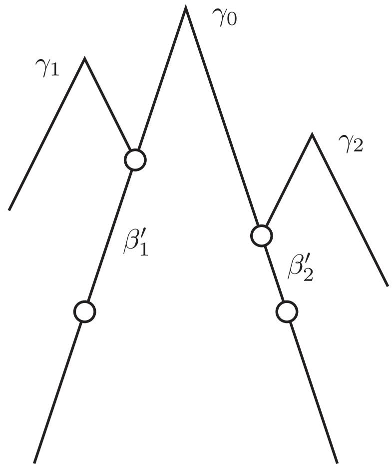

FIGURE 4. *Decomposing the vanishing path* γ *in the proof of Proposition* 5.22.

Proof. Since folding Γ does not collapse loops and since Γ has a unique illegal turn, for any path α in Γ, there is a lower bound on the distance between two illegal turns in $\sigma_{0t}(\alpha)$ independent of α and t. Therefore, there is an ϵ such that if α is a vanishing path with $\text{length}(\alpha) < \epsilon$, then α contains a single illegal turn. Hence, there is a $\delta > 0$ such that if γ is a vanishing path for $\sigma_{0t} : \Gamma \to \Gamma$ with t minimal, then, for $t' = t - \delta$, we find that $[\sigma_{0t'}(\gamma)]$ contains a single illegal turn.

We will prove the proposition by induction. The inductive claim is as follows: if γ is a vanishing path for $\sigma_{0t_\gamma} : \Gamma \to \Gamma_{t_\gamma}$ with t_γ minimal, then $\gamma = [a_1 \cdot a_0 \cdot a_2]$, where a_0 is a special vanishing path and a_1 and a_2 are vanishing paths for $\sigma_{0t_1} : \Gamma \to \Gamma_{t_1}$, where $t_1 = t_\gamma - \delta$. The base case, when $t_\gamma \leqslant T$, where T is the constant from Lemma 5.21, is proved by Lemma 5.21 as in this case $\gamma = a_0$ is a special vanishing path.

By construction a single illegal turn in γ is still present in $[\sigma_{0t_1}(\gamma)]$. In the universal cover of Γ, there are lifts of the vanishing path γ and the Nielsen path that share a lift of this illegal turn. Write the lift of γ as $\gamma_1 \cdot \gamma_0 \cdot \gamma_2$, where γ_0 is the common overlap between the lift of γ and the lift of the Nielsen path. This decomposes the lift of the Nielsen path into $\beta_1 \cdot \gamma_0 \cdot \beta_2$. Note that β_1 and β_2 are legal paths.

Since $[\sigma_{0t_1}(\gamma)]$ only contains a single illegal turn, as in Lemma 5.21 we see that $[\sigma_{0t_1}(\gamma)]$ is contained in the image of the Nielsen path in Γ_{t_1}. Therefore, we can find a subpath β'_1 of β_1 such that $\gamma_1 \cdot \overline{\beta'_1}$ is a vanishing path for σ_{0t_1}. Similarly we can find a subpath β'_2 of β_2 such that $\gamma_2 \cdot \overline{\beta'_2}$ is also a vanishing path for σ_{0t_1}. Then $\beta'_1 \cdot \gamma_0 \cdot \beta'_2$ is a special vanishing path; see Figure 4. By induction, we know that both $\gamma_1 \cdot \overline{\beta'_1}$ and $\gamma_2 \cdot \overline{\beta'_2}$ are the compositions of special vanishing paths pulled tight. Therefore, γ is the composition of special vanishing paths pulled tight. □

We can now prove the first case of Theorem 1 from the Introduction.

THEOREM 5.23. *Suppose that* $\phi \in \text{Out}(F_k)$ *is a fully irreducible automorphism with a geometric stable tree. Let* λ *be the expansion factor of* ϕ*. Then, for any* $T', T'' \in cv_k$*, we have*

$$i(T', T''\phi^n) \sim \lambda^n.$$

Proof. First suppose that ϕ is represented by a parageometric train-track map $\sigma : \Gamma \to \Gamma$. This lifts to a map $f : T \to T$. By Proposition 4.3, we have $i(T', T''\phi^n) \sim \text{vol}(T_e^n)$, where

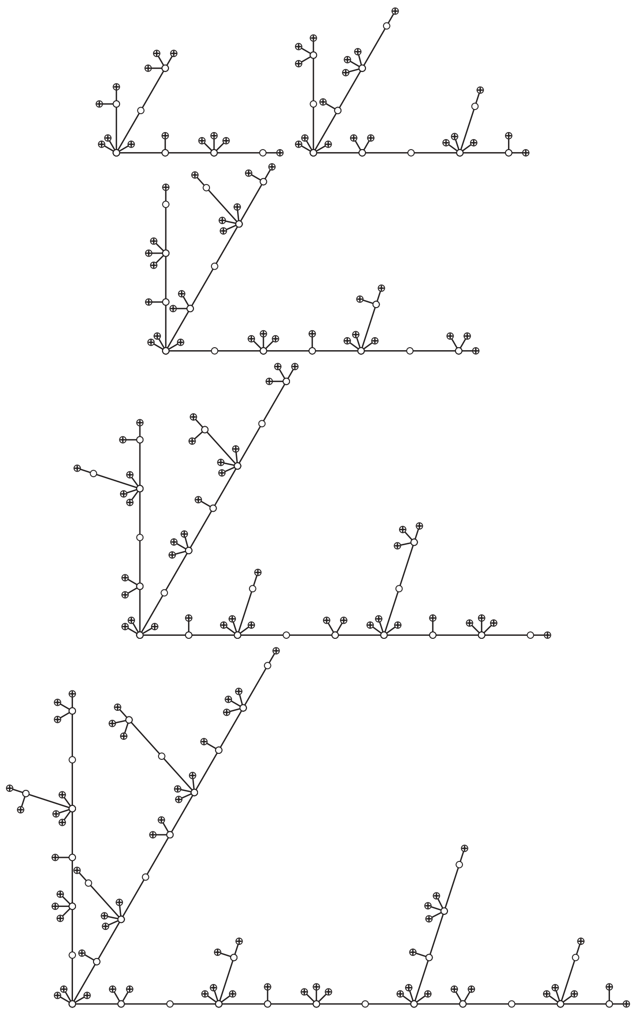

FIGURE 5. *The trees T_a^n in Example* 5.24 *for the automorphism with geometric stable tree.*

$T_e^n \subset T$ is the subtree spanned by the points in $(f^n)^{-1}(p_e)$ for any edge $e \subset T$. Then as T_e^n is a union of i-vanishing paths for $i \leqslant n$, by Proposition 5.22, T_e^m is also covered by special vanishing paths. As in Lemma 5.16 the number of special vanishing paths needed to cover T_e^n is $\sim \lambda^n$. As the length of a special vanishing path is bounded, $\mathrm{vol}(T_e^n) \sim \lambda^n$.

If ϕ is not represented by a parageometric train-track map, then by Proposition 5.20 some power ϕ^ℓ is. Then, for any i we have $i(T', T''\phi^{i+n\ell}) = i(T', (T''\phi^i)\phi^{n\ell}) \sim (\lambda^\ell)^n \sim \lambda^{i+n\ell}$, and hence $i(T', T''\phi^n) \sim \lambda^n$. □

Combining Theorems 5.19 and 5.23, we get Theorem 1 from the Introduction. We conclude with an example of a parageometric automorphism, illustrating the difference between the length of vanishing paths in the direction with a geometric stable tree and the direction with a nongeometric stable tree.

EXAMPLE 5.24. In this example we present subtrees T_a^n for the following fully irreducible automorphisms:

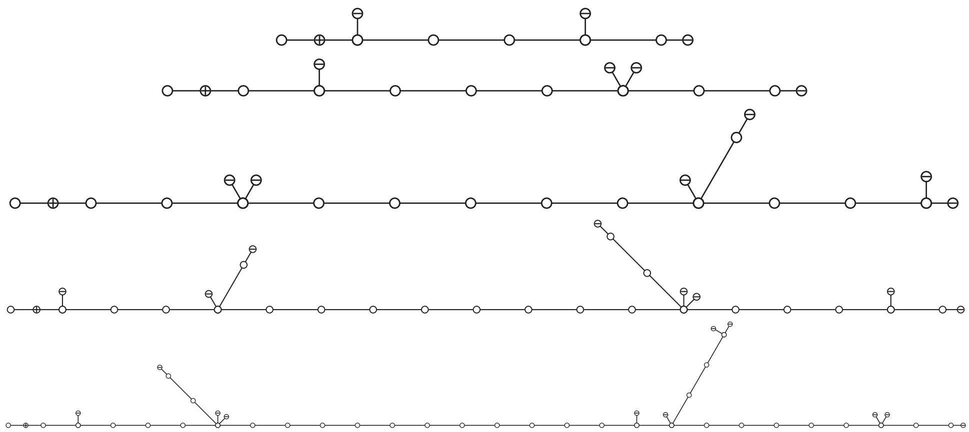

FIGURE 6. *The trees T_a^n in Example* 5.24 *for the automorphism with nongeometric stable tree.*

$$\begin{aligned} a &\longmapsto ac, & a &\longmapsto b,\\ \phi : b &\longmapsto a, & \psi : b &\longmapsto c,\\ c &\longmapsto b, & c &\longmapsto ab. \end{aligned}$$

As in Example 3.10, we let T be the universal cover of R_3, the 3-rose marked with petals labeled a, b and c and let $f_\phi : T \to T$ and $f_\psi : T \to T$ denote the lifts of the obvious homotopy equivalences of R_3 representing ϕ and ψ, respectively. Figure 5 shows $(f_\phi^n)^{-1}(p_a)$ and Figure 6 shows $(f_\psi^n)^{-1}(p_a)$ for $n = 6, 7, 8, 9, 10$. Note how the points in $(f^n)^{-1}(p_a)$ stay uniformly close together (since the stable tree for ϕ is geometric), where, as in $(f_\psi^n)^{-1}(p_a)$, they start to clump together (since the stable tree for ψ is nongeometric).

Acknowledgement. The authors would like to thank the referee for suggestions which improved this work.

References

1. J. Behrstock, B. Kleiner, Y. Minsky and L. Mosher, 'Geometry and rigidity of mapping class groups', Preprint, 2008, arXiv:0801.2006.

2. M. Bestvina and M. Feighn, 'Outer limts', Preprint, 1992.

3. M. Bestvina, M. Feighn and M. Handel, 'Laminations, trees, and irreducible automorphisms of free groups', *Geom. Funct. Anal.* 7 (1997) 215–244.

4. M. Bestvina and M. Handel, 'Train tracks and automorphisms of free groups', *Ann. of Math.* (2) 135 (1992) 1–51.

5. F. Bonahon, 'Bouts des variétés hyperboliques de dimension 3', *Ann. of Math.* (2) 124 (1986) 71–158.

6. B. H. Bowditch, 'Intersection numbers and the hyperbolicity of the curve complex', *J. Reine Angew. Math.* 598 (2006) 105–129.

7. M. R. Bridson and K. Vogtmann, 'Automorphism groups of free groups, surface groups and free abelian groups', *Problems on mapping class groups and related topics*, Proceedings of Symposia in Pure Mathematics (American Mathematical Society, Providence, RI, 2006), 301–316.

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

9. M. Culler and J. W. Morgan, 'Group actions on **R**-trees', *Proc. London Math. Soc.* (3) 55 (1987) 571–604.

10. M. Culler and K. Vogtmann, 'Moduli of graphs and automorphisms of free groups', *Invent. Math.* 84 (1986) 91–119.

11. K. Fujiwara and P. Papasoglu, 'JSJ-decompositions of finitely presented groups and complexes of groups', *Geom. Funct. Anal.* 16 (2006) 70–125.

12. V. Guirardel, 'Cœur et nombre d'intersection pour les actions de groupes sur les arbres', *Ann. Sci. École Norm. Sup.* (4) 38 (2005) 847–888.
13. M. Handel and L. Mosher, 'The expansion factors of an outer automorphism and its inverse', *Trans. Amer. Math. Soc.* 359 (2007) 3185–3208 (electronic).
14. M. Handel and L. Mosher, 'Parageometric outer automorphisms of free groups', *Trans. Amer. Math. Soc.* 359 (2007) 3153–3183 (electronic).
15. J. L. Harer, 'Stability of the homology of the mapping class groups of orientable surfaces', *Ann. of Math.* (2) 121 (1985) 215–249.
16. J. L. Harer, 'The virtual cohomological dimension of the mapping class group of an orientable surface', *Invent. Math.* 84 (1986) 157–176.
17. W. J. Harvey, 'Boundary structure of the modular group', *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference* (eds I. Kra and B. Maskit), Annals of Mathematocs Studies 97 (Princeton University Press, Princeton, 1981).
18. A. Hatcher and K. Vogtmann, 'The complex of free factors of a free group', *Q. J. Math.* (2) 49 (1998) 459–468.
19. N. V. Ivanov, 'Automorphisms of complexes of curves and of Teichmüller spaces', *Int. Math. Res. Notices 1997* (1997) 651–666.
20. I. Kapovich, 'Currents on free groups', *Topological and asymptotic aspects of group theory*, Contemporary Mathematics 394 (American Mathematical Society, Providence, RI, 2006) 149–176.
21. I. Kapovich and M. Lustig, 'Geometric intersection number and analogues of the curve complex for free groups', *Geom. Topol.* 13 (2009) 1805–1833 (electronic).
22. M. Korkmaz, 'Automorphisms of complexes of curves on punctured spheres and on punctured tori', *Topology Appl.* 95 (1999) 85–111.
23. G. Levitt and F. Paulin, 'Geometric group actions on trees', *Amer. J. Math.* 119 (1997) 83–102.
24. F. Luo, 'Automorphisms of the complex of curves', *Topology* 39 (2000) 283–298.
25. H. A. Masur and Y. Minsky, 'Geometry of the complex of curves I: Hyperbolicity', *Invent. Math.* 138 (1999) 103–149.
26. P. Scott, 'The symmetry of intersection numbers in group theory', *Geom. Topol.* 2 (1998) 11–29 (electronic).
27. P. Scott and G. A. Swarup, 'Splittings of groups and intersection numbers', *Geom. Topol.* 4 (2000) 179–218 (electronic).
28. J. R. Stallings, 'Topology of finite graphs', *Invent. Math.* 71 (1983) 551–565.
29. W. P. Thurston, 'On the geometry and dynamics of diffeomorphisms of surfaces', *Bull. Amer. Math. Soc. (N.S.)* 19 (1988) 417–431.

Jason Behrstock
Department of Mathematics
Lehman College
CUNY Bronx, NY 10468
USA

jason.behrstock@lehman.cuny.edu

Mladen Bestvina
Department of Mathematics
University of Utah
155 South 1400 East, JWB 233
Salt Lake City, UT 84102
USA

bestvina@math.utah.edu

Matt Clay
Department of Mathematics
Allegheny College
520 N. Main Street
Meadville, PA 16335
USA

mclay@allegheny.edu