

THE EXPANSION FACTORS OF AN OUTER AUTOMORPHISM AND ITS INVERSE

MICHAEL HANDEL AND LEE MOSHER

ABSTRACT. A fully irreducible outer automorphism ϕ of the free group F_n of rank n has an expansion factor which often differs from the expansion factor of ϕ^{-1} . Nevertheless, we prove that the ratio between the logarithms of the expansion factors of ϕ and ϕ^{-1} is bounded above by a constant depending only on the rank n . We also prove a more general theorem applying to an arbitrary outer automorphism of F_n and its inverse and their two spectrums of expansion factors.

1. INTRODUCTION

We denote the free group on n letters by F_n , the outer automorphism group of F_n by $\text{Out}(F_n)$ and the mapping class group of a closed surface by \mathcal{MCG} . There has been a great deal of success in developing the analogy between $\text{Out}(F_n)$ and \mathcal{MCG} in general and between Outer Space X_n and Teichmüller space \mathcal{T} in particular. The survey paper [Vog02] is an excellent reference. One notable exception is the lack in $\text{Out}(F_n)$ of an analogue for Teichmüller theory. This is the first of several papers on this topic; see also the companion paper [HM06]. Our focus is not on finding an analogue of the Teichmüller metric but rather on finding an analogue of Teichmüller geodesics.

If γ is a Teichmüller geodesic that is ψ -invariant for some $\psi \in \mathcal{MCG}$, then ψ is pseudo-Anosov and γ serves as an axis for the action of ψ on \mathcal{T} : the minimum translation distance of ψ is realized precisely on γ . This translation distance is $\log \lambda$, where λ is the expansion factor of ψ . Symmetry of the Teichmüller metric implies that ψ and ψ^{-1} have the same axis and the same expansion factor.

A strictly analogous result does not hold in $\text{Out}(F_n)$ no matter what metric is placed on X_n . An outer automorphism ϕ is *fully irreducible*, written $\phi \in \mathcal{FI}$, if there are no free factors of F_n that are invariant under an iterate of ϕ . Fully irreducible ϕ correspond to pseudo-Anosov mapping classes. They are the most well behaved and best understood elements of $\text{Out}(F_n)$. The action induced by $\phi \in \mathcal{FI}$ on the boundary of X_n has exactly two fixed points, one a source and the other a sink with all other orbits limiting on these points [LL03]. Each $\phi \in \mathcal{FI}$ has a well defined expansion factor that can be thought of as the exponential growth rate of the action of ϕ on conjugacy classes in F_n . In contrast to \mathcal{MCG} , the expansion factor λ of ϕ need not equal the expansion factor μ of ϕ^{-1} . This implies that the action of ϕ on

Received by the editors December 9, 2004 and, in revised form, April 22, 2005.

2000 *Mathematics Subject Classification*. Primary 20E05; Secondary 20E36, 20F65.

The first author was supported in part by NSF grant DMS0103435.

The second author was supported in part by NSF grant DMS0103208.

X_n cannot have an axis in the traditional sense: if there were a metric on X_n and an axis for ϕ that uniquely minimizes translation length, symmetry of the metric would imply that $\lambda = \mu$ as in Teichmüller space.

Any attempt to develop a Teichmüller theory for $\text{Out}(F_n)$ must address this phenomenon. Given that the expansion factor of ϕ^n is the n^{th} power of the expansion factor of ϕ , the correct way to measure the variation between λ and μ is by the ratio of $\log(\lambda)$ to $\log(\mu)$. The main result of this paper is that this ratio is bounded independently of ϕ .

It is generally believed that most fully irreducible outer automorphisms have different expansion factors than their inverses. In the companion paper [HM06] we exhibit this property for a large class of examples, the ‘parageometric’ outer automorphisms.

Throughout this paper we say that a bound or a constant is uniform if it depends only on n and not on the choice of $\phi \in \text{Out}(F_n)$. Suppose that $A, B > 0$. If A/B is uniformly bounded above, then we write $A \prec B$, and if A/B is uniformly bounded below, then we write $A \succ B$. If $A \prec B$ and $A \succ B$, then we write $A \sim B$.

Theorem 1.1. *Suppose that $\phi \in \text{Out}(F_n)$ is fully irreducible, that λ is the expansion factor of ϕ and that μ is the expansion factor of ϕ^{-1} . Then $\log(\lambda) \sim \log(\mu)$.*

Associated to an arbitrary outer automorphism $\phi \in \text{Out}(F_n)$ there is a finite, indexed set of expansion factors with values in $(1, \infty)$, in terms of which we can formulate a generalization of Theorem 1.1. We recall the definition of these expansion factors from sections 3.1–3.3 of [BFH00], expressed using the language of relative train track maps and attracting laminations.

To each $\phi \in \text{Out}(F_n)$ there is associated (Definition 3.1.5 of [BFH00]) a finite set of attracting laminations $\mathcal{L}(\phi)$. If $g: G \rightarrow G$ is a relative train track representative of ϕ , then $\mathcal{L}(\phi)$ contains one element Λ^+ for each subgraph $H \subset G$ which is an exponentially growing aperiodic stratum for some power of g ; generic leaves of Λ^+ are obtained by iterating edges of H and passing to a limit. The expansion factor for the action of ϕ on Λ^+ (Definition 3.3.2 of [BFH00]) measures the asymptotic expansion of the number of edges of H in a generic leaf segment of Λ^+ under iteration. By Proposition 3.3.3 of [BFH00], this expansion factor equals the Perron-Frobenius eigenvalue of the submatrix associated to H of the transition matrix for g . The set of attracting laminations $\mathcal{L}(\phi)$ and the expansion factor associated to each element of $\mathcal{L}(\phi)$ are independent of the choice of relative train track representative.

There is a natural bijective pairing between the expansion factors of ϕ and those of ϕ^{-1} , defined as follows. Associated to each $\Lambda^+ \in \mathcal{L}(\phi)$ is a free factor of F_n , called the *supporting free factor* of Λ^+ ; this is the smallest free factor on which each leaf of Λ^+ is carried in the appropriate sense. Distinct elements of $\mathcal{L}(\phi)$ have distinct supporting free factors. As proved in [BFH00, Lemma 3.2.4], there is a natural bijection between $\mathcal{L}(\phi)$ and $\mathcal{L}(\phi^{-1})$, where $\Lambda^+ \in \mathcal{L}(\phi)$ corresponds to $\Lambda^- \in \mathcal{L}(\phi^{-1})$ if and only if Λ^+, Λ^- have the same supporting free factor. Thus, the expansion factors of ϕ and of ϕ^{-1} can be indexed by the same set, namely the set of supporting free factors of their attracting laminations; this common indexing defines the desired bijective pairing.

Throughout this paper, we shall denote the paired expansion factors of ϕ, ϕ^{-1} as a pair of sequences $(\lambda_i), (\mu_i)$, respectively, where i varies over the same finite index set.

For completeness sake, in section 2 we shall give a self-contained definition of expansion factors and the bijective pairing between inverse expansion factors, formulated entirely in terms of filtered topological representatives, without recourse to the overhead needed to define relative train track maps and attracting laminations.

The natural generalization of Theorem 1.1 holds in the general case.

Theorem 1.2. *Given $\phi \in \text{Out}(F_n)$, if $(\lambda_i), (\mu_i)$ are the paired expansion factors of ϕ, ϕ^{-1} , respectively, then $\log(\lambda_i) \sim \log(\mu_i)$ for all i .*

Methods of proof. Denote the rose with n petals by R_n and identify $\pi_1(R_n)$ with F_n . A *marked graph* is a graph G , all of whose vertices have valence at least two, equipped with a homotopy equivalence $\rho_G : R_n \rightarrow G$ called the *marking*. A homotopy equivalence $f : G \rightarrow G'$ between marked graphs *respects the markings* if $f \circ \rho_G$ is homotopic to $\rho_{G'}$.

Consider a homotopy equivalence $f : G \rightarrow G'$ such that f takes vertices to vertices, f respects the markings, and for each edge e of G , either $f(e)$ is a single vertex or $f|_e$ is an immersion. If $f(e)$ is a vertex, define $|f(e)| = 0$. If $f|_e$ is an immersion, define $|f(e)|$ to be the edge length of the immersed path $f(e)$; equivalently, if $f|_e$ is an immersion, then $f|_e$ is one plus the number of points in the interior of e that map to vertices. Define $L(f) = \log(\sum_e |f(e)|)$ where the sum is taken over all edges e of G . Thus $L(f)$ is the log of the ‘total edge length’ of f .

Define $d(G, G') = \min_f \{L(f)\}$. Theorem 1.1 reduces (see section 2) to the following proposition.

Proposition 1.3. *$d(G, G') \sim d(G', G)$ for all marked graphs G and G' that have no valence 2 vertices.*

Recall [Vog02] that the spine K_n of Outer Space has the structure of a locally finite simplicial complex whose vertices are equivalence classes of marked graphs without valence 2 vertices, where G and \hat{G} are equivalent if there is a simplicial homeomorphism $h : G \rightarrow \hat{G}$ that respects the markings. Thus $d(G, G')$ descends to a well defined function on ordered pairs of vertices of K_n . It is easy to check that there exists $C > 0$ so that $d(G, G) < C$ for all G and so that $d(G, G'') \leq d(G, G') + d(G', G'') + C$ for all G, G' and G'' . Proposition 1.3 can therefore be interpreted as the statement that $d(G, G')$ defines a ‘quasi-metric’ on the vertex set of K_n and hence a ‘quasi-metric’ on K_n .

One question which arises is whether the ‘quasi-metric’ d is quasi-isometric to the standard simplicial metric d_s on K_n , which itself is quasi-isometric to any word metric on $\text{Out}(F_n)$. In other words, do there exist constants $L \geq 1, C \geq 0$ such that $\frac{1}{L} d_s(G, G') - C \leq d(G, G') \leq L d_s(G, G') + C$, for any two marked graphs G, G' ? This question is easily answered in the negative. Suppose that G_0 is a marked rose with edges e_1, \dots, e_n and that $f : G_0 \rightarrow G_0$ is defined by $f(e_2) = e_2 e_1$ and $f(e_i) = e_i$ for $i \neq 2$. Let $G_m, m > 0$, be the marked graph obtained from G_0 by changing its marking so that $f^m : G_0 \rightarrow G_m$ respects the markings. Then $d(G_0, G_m) \leq \log(m)$, as opposed to the linear growth of $d_s(G_0, G_m)$ proved in [Ali02]. This argument implies, moreover, that d is not quasi-isometric to any $\text{Out}(F_n)$ -equivariant proper, geodesic metric on K_n , because all such metrics are quasi-isometric to d_s . Whether a precise relation between d and d_s can be described remains unclear.

We conclude this section with a discussion of the proof of Proposition 1.3 and Theorems 1.1 and 1.2.

A *filtration* \mathcal{F} of a marked graph G with length N is an increasing sequence of distinct subgraphs $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ such that no G_i has a valence one vertex. For $1 \leq i \leq N$, the closure of $G_i \setminus G_{i-1}$ is called the i th *stratum* of the filtration and is denoted by either H_i or $G(i, i)$. For any $f : G \rightarrow G'$ as above, denote $f(G_i)$ by $G'_i \subset G'$. If G'_i has the same rank as G_i for all i , we say that f *respects the filtration*. In this case $\emptyset = G'_0 \subset G'_1 \subset \cdots \subset G'_N = G'$ is a filtration for G' and f restricts to a homotopy equivalence between G_i and G'_i for each i .

Every $f : G \rightarrow G'$ respects the unique filtration $\emptyset = G_0 \subset G_1 = G$ of length one. We may therefore assume that $f : G \rightarrow G'$ is equipped with a filtration \mathcal{F} that it respects. In the context of Theorem 1.2, one expects to work with filtrations that have length greater than one. It turns out that even in the context of Theorem 1.1, it is helpful to consider such filtrations. In section 4 we state our main technical result, Proposition 4.1. It is a generalization of Proposition 1.3 that incorporates filtrations into its statement. The statement of Proposition 4.1 also includes properties that are required for our induction argument and so is somewhat lengthy. Using results from the theory of relative train track maps, we show in section 4 that Proposition 4.1 implies Proposition 1.3 and Theorems 1.1 and 1.2.

For the sake of simplicity we focus our remaining remarks on Proposition 1.3.

A homotopy equivalence $f : G \rightarrow G'$ as above can be factored via [Sta83] as

$$G = K^0 \xrightarrow{p_1} K^1 \xrightarrow{p_2} \cdots \xrightarrow{p_k} K^k \xrightarrow{\theta} G'$$

where each p_i is a fold and θ is a homeomorphism. Each p_k and θ in this *folding sequence* has a natural homotopy inverse. Composing these in reverse order defines a homotopy inverse $g : G' \rightarrow G$ for $f : G \rightarrow G'$. If $L(g) \prec L(f)$, then we say that the folding sequence has a controlled inverse.

We prove that every $f : G \rightarrow G'$ as above has a folding sequence with a controlled inverse. The proof is by downward induction on the length of a filtration \mathcal{F} respected by f . Thus we first consider those f that respect a filtration of maximal length $2n - 1$ and work our way down to those f that respect only the unique filtration of length one.

Assume that \mathcal{F} has been specified. A fold p identifies points in a pair of edges. If both edges belong to G_j and if p restricts to a simplicial homeomorphism of G_{j-1} , then we say that p is supported on H_j . In Lemma 6.3, we show that it is sufficient to consider only those f for which there exists $1 \leq j \leq N$ and a folding sequence in which every fold is supported on H_j . We say that such an f is supported on a single stratum (with respect to \mathcal{F}).

If \mathcal{F} has maximal length and f is supported on a single stratum, then every folding sequence for f has a controlled inverse. The main steps in proving this are Lemma 6.4 and Lemma 7.3.

At the other extreme, suppose that \mathcal{F} has length one. There being only one stratum there is nothing gained by assuming that f is supported on a single stratum. In this case we begin with a randomly chosen folding sequence for f with notation as above. For $1 \leq a \leq b \leq k$, define $f_{a,b} = p_b \circ \cdots \circ p_a : K^{a-1} \rightarrow K^b$ and say that $f_{a,b}$ is reducible if it respects a filtration of length greater than 1. Let a_1 be the largest value for which f_{1,a_1} is reducible. If $a_1 < k$, consider the factorization $f = f_{a_1+2,k} p_{a_1+1} f_{1,a_1}$. Repeating this operation with $f_{a_1+2,k}$ replacing f leads to a decomposition of f as an alternating concatenation of maximal length reducible

maps and single folds. As our proof is by downward induction on the length of the filtration, we may assume that there are folding sequences with controlled inverses for each of the maximal length reducible maps. Together with the single folds, they define a folding sequence for f which we prove has a controlled inverse.

For filtration lengths between the minimum and the maximum, a more general notion of reducibility is required. This is addressed in section 7.

Remark. Although the techniques of this paper share some concepts with the techniques of relative train track maps, particularly the use of filtrations, this paper is almost entirely independent of the theory of relative train track maps. The sole exceptions are as follows. While the statement of Lemma 3.1 is independent of relative train track maps, the proof is an extended exercise in the methods of relative train track maps. Also, we will need to quote results on relative train track maps to deduce Theorem 1.1 from Proposition 1.3 and to justify our definitions of expansion factors in section 2.

2. PRELIMINARIES

All maps $h: G \rightarrow G'$ are assumed to be homotopy equivalences of marked graphs of a fixed rank n . We say that $h: G \rightarrow G'$ is *simplicial* if for each edge E of G , either $h(E)$ is a vertex or $h(E)$ is a single edge of G' and $h|_{\text{Int}(E)}$ is injective. We assume throughout that $h: G \rightarrow G'$ takes vertices to vertices and is simplicial with respect to some subdivision of G .

Filtrations. By a *weak filtration* (G_i) of a marked graph G with length N we mean a properly nested sequence of subgraphs

$$\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$$

some of which may have valence one vertices. A *filtration* is a weak filtration in which no G_i has a valence one vertex, and thus all components of G_i are non-contractible and G_k is not homotopy equivalent to G_{i-1} . For sets $A \subset X$ we use the notation $X \setminus A$ for the complement of A in X ; in general $X \setminus A$ is not closed in X . We denote $\text{Cl}(G_b \setminus G_{a-1})$ by $G(b, a)$. In the special case that $b = a$ we say that $G(b, b) = \text{Cl}(G_b \setminus G_{b-1})$ is the *b-stratum* of G . A marked graph with a (weak) filtration is called a *(weak) filtered marked graph*. Our standing notation will be that the filtration on a graph G has filtration elements G_i .

Given a filtration (G_i) of length N of a rank n marked graph G , the length N satisfies the bound $N \leq 2n - 1$, because for each pair of properly nested filtration elements $G_i \subset G_{i+1}$, either $\text{rk}(H_1(G_{i+1})) > \text{rk}(H_1(G_i))$ or $\text{rk}(H_0(G_{i+1})) < \text{rk}(H_0(G_i))$.

Suppose that G and G' are marked graphs with weak filtrations of length N . A homotopy equivalence $h: G \rightarrow G'$ respects the weak filtrations if:

- $h|_{G_i}: G_i \rightarrow G'_i$ is a homotopy equivalence for each $1 \leq i \leq N$.

If in addition the weak filtrations are filtrations, so that no G_i or G'_i has any valence one vertices, then $G'_i = h(G_i)$.

Suppose that G and G' are marked graphs with weak filtrations of length N and that $h: G \rightarrow G'$ respects the weak filtrations. Suppose also that $\{E_1, \dots, E_r\}$ and $\{E'_1, \dots, E'_s\}$ are the edges of G and G' numbered so that the edges in G_{i+1} have larger indices than the edges in G_i and similarly for G' . A path $\sigma \subset G$ with its orientation reversed is denoted by $\bar{\sigma}$. Since h is simplicial with respect to a subdivision

of G , we can identify $h(E_k)$ with a (possibly trivial) word in $\{E'_j, \bar{E}'_j : 1 \leq j \leq s\}$ called *the edge path associated to $h(E_k)$* . We allow the possibility that $h|_{E_k}$ is not an immersion or equivalently that the edge path associated to $h(E_k)$ is not a reduced word. *For the remainder of the paper we identify the image of an edge with its edge path.* The *transition matrix associated to $h: G \rightarrow G'$* is the matrix $M(h)$ whose jk th entry is the number of times that E'_j or \bar{E}'_j occurs in $h(E_k)$. For $1 \leq r \leq s \leq N$ let $M_{r,s}(h)$ be the submatrix of $M(h)$ corresponding to the edges of $G(s, r)$ and the edges of $G'(s, r)$.

It is important that we bound the size of the transition matrices. Since n is fixed, we can bound the number of edges in G and G' , and hence the size of the transition matrix, by bounding the number of valence two vertices in G and G' . There are several steps in the proof in which we factor a map. The following definition and lemma are used to control the number of valence two vertices of the intermediate graphs.

The link $lk(v)$ of a vertex v in G is the set of oriented edges with v as initial endpoint. For $h: G \rightarrow G'$ let $lk_h(v) \subset lk(v)$ be those oriented edges whose image under h is non-trivial. There is an induced map $Dh: lk_h(v) \rightarrow lk(h(v))$ that sends E to the first edge of $h(E)$. We denote by $T(h, v)$ the *cardinality of the image $Dh(lk_h(v))$* . We record the following obvious properties of $T(h, v)$.

Lemma 2.1. *The valence of v in G is greater than or equal to $T(h, v)$. If $h_1: G \rightarrow K$ and $h_2: K \rightarrow G'$, then $T(h_1, v), T(h_2, h_1(v)) \geq T(h_2 \circ h_1, v)$.*

Matrix lemmas. Our definitions of $A \prec B$ and $A \succ B$ require that $A, B > 0$. Many of our comparisons involve $\log(C)$ for some $C \geq 1$. To remove $C = 1$ as a special case we define

$$\text{mlog}(C) = \max\{1, \log(C)\}.$$

For any matrix M denote the largest coefficient of M by $LC(M)$ and the sum of all entries in M by $L(M)$. The following lemma is obvious but useful.

Lemma 2.2. *Suppose that M_1 and M_2 are non-negative integer matrices with at most α rows and columns. Then $LC(M_1 M_2) \leq \alpha LC(M_1) LC(M_2)$ and $LC(M_1) \leq L(M_1) \leq \alpha^2 LC(M_1)$.*

We will often work with non-square matrices. For this and other reasons it is more convenient to think about largest coefficients than about Perron-Frobenius eigenvalues. We recall the following well known fact.

Lemma 2.3. *Suppose that M is a non-negative integral irreducible matrix and that λ is its Perron-Frobenius eigenvalue. If the size of M is uniformly bounded, then $\text{mlog}(\lambda) \sim \text{mlog}(LC(M)) \sim \text{mlog}(L(M))$.*

Proof. Suppose that all matrices being considered have at most α rows and columns. Since $\lambda \geq 1$ is an eigenvalue with positive eigenvector for M , $\lambda \leq \alpha \cdot LC(M)$. On the other hand, every non-zero coordinate of M^α is greater than or equal to $LC(M)$. This implies that $\lambda^\alpha \geq LC(M)$. This proves the first comparison. The second comparison follows from Lemma 2.2. \square

Reduction of Theorem 1.1 to Proposition 1.3. A homotopy equivalence $f: G \rightarrow G$ of a marked graph determines an outer automorphism of $\pi_1(G)$ and so an element $\phi \in \text{Out}(F_n)$. We say that $f: G \rightarrow G$ *represents ϕ* .

Proposition 1.3 implies Theorem 1.1. By Theorem 1.7 of [BH92] there exists $f: G \rightarrow G$ representing ϕ such that $M(f)$ is irreducible and has Perron-Frobenius eigenvalue λ . Lemma 2.3 and Proposition 1.3 imply that there is a homotopy inverse $g: G \rightarrow G$ for f such that $\text{mlog}(LC(g)) \prec \log(\lambda)$. Remark 1.8 of [BH92] implies that some submatrix of $M(g)$ is irreducible with Perron-Frobenius eigenvalue greater than or equal to μ . Another application of Lemma 2.3 shows that $\log(\mu) \prec \log(\lambda)$. Symmetry completes the proof. \square

3. THE EXPANSION FACTORS OF AN OUTER AUTOMORPHISM

In this section we define the expansion factors of an outer automorphism ϕ , and the bijective pairing between the expansion factors of ϕ and of ϕ^{-1} . We shall state these definitions in a self-contained manner, solely in terms of filtered marked graphs and free factor systems for F_n , although several results from [BH92] and [BFH00] regarding relative train track maps and attracting laminations will be quoted in order to prove that these definitions make sense and that they agree with the definitions given in the introduction.

We also state and prove Lemma 3.1, which collects in one statement several properties of relative train track maps that will be useful in what follows.

The reader who is interested only in Theorem 1.1 can skip this section entirely. The reader who is interested in Theorem 1.2 and who has a good understanding of the concepts of relative train track maps, attracting laminations, and expansion factors can just read the statement of Lemma 3.1 and skip the rest of this section.

Definition of expansion factors. We now formulate the definition of the expansion factors of an arbitrary outer automorphism $\phi \in \text{Out}(F_n)$.

Let $f: G \rightarrow G$ be a homotopy equivalence of a marked graph. Among all weak filtrations (G_i) that f respects there is a unique maximal one, and it has the property that the transition matrix $M_{ii}(h)$ for the i -stratum $G(i, i)$ is either an irreducible matrix or a zero matrix ([BH92] section 5). If $M_{ii}(h)$ is irreducible and has Perron-Frobenius eigenvalue $\lambda_i > 1$, then we say that $G(i, i)$ is an *exponentially growing* or *EG*-stratum. Let $\Gamma(f)$ denote the sequence of Perron-Frobenius eigenvalues of EG-strata written in decreasing order $\lambda_{i_1} \geq \dots \geq \lambda_{i_R}$, where $\{i_1, \dots, i_R\}$ is an enumeration of the indices of the EG-strata (see pages 32–33 of [BH92] where $\Gamma(f)$ is denoted $\Lambda(f)$). Although the sequence $\Gamma(f)$ may already have repeated entries, the multiplicity is not quite correct for defining the expansion factors of ϕ . The correct multiplicity is determined as follows. For each irreducible square matrix M there exists a unique integer $p > 1$, the *multiplicity* of M , such that some positive power M^j is a block diagonal matrix with p -blocks each having all positive entries. Let p_i be the multiplicity of $M_{ii}(f)$; if $p_i = 1$, then $G(i, i)$ is called an *EG-aperiodic* stratum. Let $\hat{\Gamma}(f)$ be the sequence obtained from $\Gamma(f)$ where λ_i is replaced by p_i equal entries of the same value as λ_i . To put it another way, letting p be the least common multiple of the multiplicities of all EG-strata of f , then each EG-stratum $G(i, i)$ of f of multiplicity p_i breaks into p_i EG-strata of f^p of multiplicity 1, and so $\hat{\Gamma}(f)$ is obtained from $\Gamma(f^p)$ by taking the p th root of each entry of $\Gamma(f^p)$. Note that the operation of replacing $\Gamma(f)$ by $\hat{\Gamma}(f)$ preserves lexicographic order.

Now define $\hat{\Gamma}(\phi)$ to be the lexicographic minimum of the set

$$\{\hat{\Gamma}(f) \mid f: G \rightarrow G \text{ represents } \phi\}.$$

To prove that the lexicographic minimum exists, the subset where G has no valence 2 vertices is cofinal with respect to the “greater than” relation, according to [BH92, Lemma 5.4]. But that subset is well-ordered, because if G has no valence 2 vertices, then it has at most $3n - 3$ edges, so f has at most $3n - 3$ strata, and so $\Gamma(f)$ has at most $3n - 3$ entries. The indexed set $\hat{\Gamma}(\phi)$ is defined to be the expansion factors of ϕ .

To check that this definition of the expansion factors agrees with the one given in the introduction, we quote Theorem 5.12 from [BH92] to conclude that $\Gamma(f)$ is minimized by some relative train track representative $f_0: G_0 \rightarrow G_0$, and so $\hat{\Gamma}(f)$ is minimized as well by f_0 . We then pass to a power of ϕ so that each EG-stratum of f_0 is aperiodic; the effect is to take the p th power of each element of $\hat{\Gamma}(f_0)$, and the p th power of the expansion factor of any each expanding lamination in $\mathcal{L}(\phi)$. Next, consider *any* relative train track representative $f: G \rightarrow G$ of ϕ . By Lemma 3.1.14 of [BFH00], each EG-stratum of f is aperiodic. Next we quote Definition 3.1.12 of [BFH00] to conclude that the attracting laminations of ϕ correspond bijectively with the EG-strata of f . Finally we quote Proposition 3.3.3 from [BFH00] to conclude that the expanding lamination associated to an aperiodic EG-stratum of f has expansion factor equal to the Perron-Frobenius eigenvalue of that stratum.

The results just quoted also show that for *every* relative train track representative $f: G \rightarrow G$ of ϕ we have $\hat{\Gamma}(f) = \hat{\Gamma}(\phi)$.

The bijective pairing between $\hat{\Gamma}(\phi)$ and $\hat{\Gamma}(\phi^{-1})$. A *free factor* of F_n is a non-trivial subgroup $A < F_n$ such that $F_n = A * B$ for some subgroup B . Let $[A]$ denote the conjugacy class of A . Define a partial ordering on free factor conjugacy classes, where $[A] \sqsubset [A']$ if A is conjugate to a free factor of A' . Note that if $F = A * A' * B$ where A, A' are non-trivial, then $[A] \neq [A']$. Given a free factorization $F_n = A_1 * \cdots * A_I * B$ where A_1, \dots, A_I are non-trivial, we say that $\{A_1, \dots, A_I\}$ is an *independent set of free factors*, and that $\mathcal{A} = \{[A_1], \dots, [A_I]\}$ is a *free factor system*. By convention the empty set \emptyset is a free factor system. We extend the partial ordering \sqsubset to free factor systems, where $\mathcal{A} \sqsubset \mathcal{A}'$ if for each $[A] \in \mathcal{A}$ there exists $[A'] \in \mathcal{A}'$ such that $[A] \sqsubset [A']$; the relation $\mathcal{A} \sqsubset \mathcal{A}'$ is *proper* if $\mathcal{A} \neq \mathcal{A}'$. A *free factor filtration* (\mathcal{A}^k) of length K is a finite, properly nested sequence of free factor systems $\emptyset = \mathcal{A}^0 \sqsubset \mathcal{A}^1 \sqsubset \cdots \sqsubset \mathcal{A}^K = \{[F_n]\}$. The length K satisfies the same bound $K \geq 2r - 1$ for the same reason as the length of a marked graph filtration: for each proper relation $\mathcal{A}^k \sqsubset \mathcal{A}^{k+1}$, either the sum of the ranks of \mathcal{A}^{k+1} is strictly greater than the sum of the ranks of \mathcal{A}^k , or the cardinality of \mathcal{A}^{k+1} is strictly less than the cardinality of \mathcal{A}^k .

The group $\text{Out}(F_n)$ acts on free factor conjugacy classes, systems, and filtrations. Given $\phi \in \text{Out}(F_n)$, a ϕ -invariant free factor filtration (\mathcal{A}^k) is *reduced* if the following holds: if a free factor system \mathcal{A}' is invariant under the action of an iterate of ϕ , and if $\mathcal{A}^k \sqsubset \mathcal{A}' \sqsubset \mathcal{A}^{k+1}$, then either $\mathcal{A}^k = \mathcal{A}'$ or $\mathcal{A}' = \mathcal{A}^{k+1}$. If this does not hold, so that \mathcal{A}' exists with $\mathcal{A}^k \sqsubset \mathcal{A}' \sqsubset \mathcal{A}^{k+1}$ and $\mathcal{A}' \neq \mathcal{A}^k, \mathcal{A}^{k+1}$, then by passing to a power of ϕ which preserves \mathcal{A}' , we can insert \mathcal{A}' and get a longer ϕ -invariant free factor filtration. Since the length of a free factor filtration is bounded above, it follows that some power of ϕ respects a reduced free factor filtration.

To each marked graph G and each subgraph $G' \subset G$ there corresponds a free factor system $\mathcal{F}(G')$ as follows: if G' is connected, then $\mathcal{F}(G')$ is a singleton consisting of the (well-defined) conjugacy class of $\pi_1(G')$ in $\pi_1(G)$; and in general if G' has components $G' = G'_1 \cup \cdots \cup G'_m$, then the (well-defined) conjugacy classes of

$\pi_1(G'_1), \dots, \pi_1(G'_m)$ in $\pi_1(G)$ collect together to form a free factor system $\mathcal{F}(G')$. To each filtered marked graph $\emptyset = G_0 \subset G_1 \subset \dots \subset G_K = G$ there corresponds a free factor filtration $\emptyset \sqsubset \mathcal{F}(G_1) \sqsubset \dots \sqsubset \mathcal{F}(G_K) = \{[F_n]\}$. If $\phi \in \text{Out}(F_n)$ has a topological representative $f: G \rightarrow G$ that respects the filtration of G , then ϕ preserves the corresponding free factor filtration. Conversely, the proof of Lemma 2.6.7 of [BFH00] shows that for any ϕ -invariant free factor filtration (\mathcal{A}^k) there exists a filtered marked graph (G_k) and a filtration-respecting representative $f: G \rightarrow G$ of ϕ such that $\mathcal{A}^k = \mathcal{F}(G_k)$. Moreover, we may choose f to be a relative train track representative of ϕ .

Consider now $\phi \in \text{Out}(F_n)$. We shall define the bijective pairing between $\hat{\Gamma}(\phi)$ and $\hat{\Gamma}(\phi^{-1})$ after passing to a certain power of ϕ , and one easily checks that the definition is independent of which power is chosen. Replace ϕ by a power so that there exists a ϕ -invariant reduced free factor filtration (\mathcal{A}^k) . Choose a filtered marked graph $\emptyset = G_0 \subset G_1 \subset \dots \subset G_K = G$ and a filtration-respecting representative $f: G \rightarrow G$ of ϕ such that $\mathcal{A}^k = \mathcal{F}(G_k)$; we may assume that f is a relative train track representative of ϕ , and so in particular $\hat{\Gamma}(f)$ realizes $\hat{\Gamma}(\phi)$. Since the free factor filtration is reduced, it follows that each stratum $G(k, k)$ is aperiodic.

The key result of [BFH00], which establishes the pairing between $\hat{\Gamma}(\phi)$ and $\hat{\Gamma}(\phi^{-1})$, is Lemma 3.2.4. The statement of this lemma requires full knowledge of the attracting laminations of ϕ and of ϕ^{-1} . But the proof of Lemma 3.2.4 yields the following statement, which does not require knowledge of attracting laminations, and which can still be used to support the definition of the pairing, because it characterizes EG-strata purely in terms of free factor filtrations:

- If $f: G \rightarrow G$ is a relative train track map whose strata are aperiodic and whose associated free factor filtration $(\mathcal{A}^k = \mathcal{F}(G_k))$ is reduced, then for each k , the stratum $G(k, k)$ is *not* EG if and only if the following holds:
 - (*): *either* there exist two elements $[A], [A'] \in \mathcal{A}^{k-1}$ represented by independent free factors A, A' so that

$$\mathcal{A}^k = (\mathcal{A}^{k-1} \setminus \{[A], [A']\}) \cup \{[A * A']\}$$

or there exist $[A] \in \mathcal{A}^{k-1}$ and $[A'] \in \mathcal{A}^k$ represented by free factors $A \subset A'$ so that $\text{rank}(A') = \text{rank}(A) + 1$, and

$$\mathcal{A}^k = (\mathcal{A}^{k-1} \setminus \{[A]\}) \cup \{[A']\}.$$

Using this statement we can define the pairing between $\hat{\Gamma}(\phi)$ and $\hat{\Gamma}(\phi^{-1})$ as follows. We are assuming that (\mathcal{A}^k) is a reduced free factor filtration invariant under ϕ , and so also invariant under ϕ^{-1} . Choose relative train track representatives $f: G \rightarrow G$ for ϕ and $f': G' \rightarrow G'$ for ϕ^{-1} so that $\mathcal{F}(G_k) = \mathcal{F}(G'_k) = \mathcal{A}^k$, for all $k = 0, \dots, K$. Applying the above proposition, it follows that for each $k = 0, \dots, K$, $G(k, k)$ is an EG-stratum of f if and only if $G'(k, k)$ is an EG-stratum of f' , because both of these are equivalent to the failure of property (*). We therefore define the pairing by requiring that the Perron-Frobenius eigenvalue of f on $G(k, k)$ be paired with the Perron-Frobenius eigenvalue of f^{-1} on $G'(k, k)$.

The fact that this is well defined independent of the choice of relative train track representatives f and f' , and that this definition agrees with the one given in the introduction, is again a direct consequence of the proof of Lemma 3.2.4, which shows that the attracting lamination of ϕ associated to $G(k, k)$ is paired with the attracting lamination of ϕ' associated with $G'(k, k)$, and so the expansion factor λ_k

of ϕ associated with $G(k, k)$ is paired with the expansion factor μ_k of ϕ^{-1} associated with $G'(k, k)$.

A relative train track lemma. The proof of Theorem 1.1 followed from Proposition 1.3 plus a modest amount of train track theory quoted from [BH92], as we saw at the end of section 2. For Theorem 1.2 we will need more. The proof of the following lemma is essentially an extended exercise in relative train track theory.

Lemma 3.1. *Suppose that $\phi \in \text{Out}(F_n)$ and that λ and μ are paired expansion factors of ϕ . Then there is a filtered marked graph G with filtration (G_k) , a homotopy equivalence $f: G \rightarrow G$ that respects the filtration and that represents ϕ^p for some $p \geq 1$, and there is a stratum $G(j, j)$ such that:*

- (1): f restricts to an immersion on each edge of G ;
- (2): $T(f|_{G_i, v}) \geq 2$ for all vertices $v \in G_i$ and for all i ;
- (3): $\text{mlog}(LC(M_{jj}(f))) \sim p \log(\lambda)$;
- (4): if $g: G \rightarrow G$ is a homotopy inverse for $f: G \rightarrow G$ that respects the filtrations, then $\text{mlog}(LC(M_{jj}(g))) \succ p \log(\mu)$.

Moreover, G has at most V'_n vertices for some uniform constant V'_n .

Remark 3.2. The uniformity of V'_n expresses the fact that relative train track graphs can be chosen with a uniformly bounded number of valence two vertices. One can choose G so that no embedded arc in G contains three consecutive valence two vertices, thereby giving an explicit value for V'_n of $6n - 6$.

Notation 3.3. Choose a uniform constant $V_n \geq 2V'_n$ such that a marked graph of rank n has at most $V_n/2$ vertices of valence at least three. Thus a marked graph of rank n with at most $V_n/2$ vertices of valence two has at most V_n vertices. Choose a uniform constant $Edge_n$ such that a graph of rank n with at most V_n vertices has at most $Edge_n$ edges.

The proof of Lemma 3.1. Fix $\phi \in \text{Out}(F_n)$ and paired expansion factors λ of ϕ and μ of ϕ^{-1} .

A relative train track map $f: G \rightarrow G$ is a homotopy equivalence of a marked graph G which has useful properties for computation — explicit details can be found on page 38 of [BH92] and in Theorem 5.1.5 of [BFH00]. We will not use the most technical properties listed in this theorem and will be explicit about what is necessary.

Notation: the maximal weak filtration preserved by a relative train track map $f: G \rightarrow G$ will be denoted by (\hat{G}_i) , with strata denoted by \hat{H}_i , as is conventional when considering relative train track maps, rather than $\hat{G}(i, i)$. Also let \hat{M}_{ii} denote the transition matrix of the restriction of f to the stratum \hat{H}_i . The “hat” notation \hat{G}_i is used to distinguish this weak filtration from the (non-weak) filtration (G_k) with strata $H_k = G(k, k)$ whose description is the goal of the proof of Lemma 3.1. Also, we must describe a particular index j of an EG-stratum H_j associated to λ .

Choose $p > 0$ so that Theorem 5.1.5 of [BFH00] applies to both ϕ^p and ϕ^{-p} and let $f: G \rightarrow G$ be a relative train track map $f: G \rightarrow G$ representing ϕ^p and satisfying the conclusions of that theorem. Item (1) of Lemma 3.1 holds for every relative train track map (and can be arranged simply by tightening the image of the edges). In particular, Df is defined on every oriented edge in the link of a vertex.

The strata \hat{H}_i of the weak filtration (\hat{G}_i) are divided into four mutually exclusive types.

Exponentially growing: \hat{H}_i is an EG-stratum. As we have seen in section 2, there is a bijection between the set of exponentially growing strata \hat{H}_i and the set of paired expansion factors λ_i, μ_i of ϕ, ϕ^{-1} , respectively. If \hat{H}_i corresponds to the pair λ_i, μ_i , then $PF(\hat{M}_{ii}) = \lambda_i^p$.

Zero stratum: If \hat{G}_i has contractible components, then \hat{H}_i is the union of these components and is called a *zero stratum*. In this case $f(\hat{H}_i) \subset \hat{G}_{i-1}$.

All the remaining strata \hat{H}_i are single oriented edges \hat{E}_i .

Fixed edge: If $f(\hat{E}_i) = \hat{E}_i$, then both endpoints of \hat{E}_i have valence at least two in \hat{G}_{i-1} and \hat{H}_i is a *fixed edge*. In some contexts fixed non-loop edges are collapsed, but the resulting filtration might not be reduced, so we do not do that here.

Non-exponentially growing: Otherwise (see (ne-i) and (ne-ii) in the statement of Theorem 5.1.5 of [BFH00]) $f(\hat{E}_i) = \hat{E}_i \hat{u}_i$ for some non-trivial closed path $\hat{u}_i \subset \hat{G}_{i-1}$ that is immersed both as a path and as a loop. In this case the terminal endpoint of \hat{E}_i has valence at least two in \hat{G}_{i-1} . Strata of this fourth kind are said to be *NEG*.

We define the filtration (G_k) from the bottom up with each stratum of (G_k) being a union of strata of (\hat{G}_i) . The stratum H_j associated to λ will be the one that contains the weak EG-stratum \hat{H}_i associated to λ . Since $\hat{G}_0 = \emptyset$, \hat{H}_1 is either a fixed loop or is EG and we define $G_1 = \hat{G}_1$. In the former case, (2) is clear and $j \neq 1$. Suppose then that \hat{H}_1 is EG and that $v \in \hat{H}_1$. Inductively define $Df^m(e) = Df^{m-1}(Df(e))$. We say that a pair of edges (e, e') in the link of v determines a *legal turn* if $Df^m(e)$ and $Df^m(e')$ are distinct for all $m > 0$. Clearly (2) is satisfied at v if there is a legal turn at v . Since \hat{H}_1 is EG, there exists $m > 0$ and a point x in the interior of an edge such that $f^m(x) = v$. The ‘link’ at x contains two directions and condition (RTT-3) on page 38 of [BH92] implies that their images in $lk(v)$ define a legal turn. This verifies (2) for G_1 . If λ corresponds to \hat{H}_1 , then Condition (3) for G_1 follows from Lemma 2.3 and the fact that $PF(\hat{M}_{11}) = \lambda^p$. This completes our analysis of G_1 .

Let $l > 1$ be the first parameter value for which the non-contractible components of \hat{G}_l do not deformation retract to \hat{G}_1 . Since a zero stratum is a union of contractible components, \hat{H}_l is not a zero stratum. We will refer to the strata above G_1 and below \hat{H}_l as *intermediate strata*. Since every bi-infinite path in G_{l-1} is contained in G_1 , no intermediate stratum is EG. Property (z-i) on page 562 of [BFH00] implies that at most one intermediate stratum is a zero stratum and if there is one, then it is \hat{H}_{l-1} and \hat{H}_l is EG.

If \hat{H}_l is a fixed edge, then push all intermediate edges of $\hat{G}_l \setminus G_1$ up the weak filtration by reordering the strata. Thus $l = 2$ and we define $G_2 := \hat{G}_2$. As in the G_1 case, (2) is satisfied and $j \neq 2$. Suppose next that \hat{H}_l is NEG, and hence each intermediate stratum, is NEG. Each of these strata is a single edge with terminal endpoint in G_1 . The initial endpoint of E_l is either in G_1 or is shared with one other of the intermediate edges. Push all other intermediate edges of $\hat{G}_l \setminus G_1$ up the weak filtration by reordering the strata. Thus l is two or three and $G_2 := \hat{G}_l$ is obtained from G_1 by topologically adding an arc which may be subdivided at one point. As in the previous case (2) is satisfied and $j \neq 2$.

It remains to consider the case that \hat{H}_l is EG. If \hat{E}_i is the edge of an intermediate NEG stratum, then its initial endpoint w_i is fixed and so is not a vertex in a zero stratum. If w_i is not a vertex in \hat{H}_l , then we push \hat{E}_i up the weak filtration as before. Once this is done define $G_2 = \hat{G}_l$, so each remaining intermediate stratum is either a zero stratum comprising \hat{H}_{l-1} or an NEG stratum with initial point on \hat{H}_l and terminal point on G_1 .

Condition (RTT-1) on page 38 of [BH92] implies that if $e, e' \in lk(v)$ where e is an edge in \hat{H}_l and e' is an edge in \hat{G}_{l-1} , then (e, e') is legal. If the link of v in G_2 is entirely contained in \hat{H}_l , then we produce legal turns at v as we did for the G_1 case. If v is in a zero stratum, then (2) follows from conditions z-(ii) and (z-iii) on page 562 of [BFH00]. This verifies (2). It also shows that G_2 is f -invariant since no valence one vertex w_i in the original \hat{G}_l could be the image of either a vertex in G_2 or the interior of an edge in G_2 .

We now consider (3) for G_2 . It is clear that the only columns of $M_{22}(f)$ that have entries greater than one are those that correspond to the edges of \hat{H}_l . These columns record the number of times that the f -image α of an edge in \hat{H}_l crosses an edge of H_2 . Denote the union of the intermediate strata by X and write α as an alternating concatenation of subpaths $\alpha = \sigma_1\tau_1\sigma_2\dots$ where τ_i is contained in X and σ_i is disjoint from X . Since X is a forest, the number of edges in τ_i is uniformly bounded and either σ_{i-1} or σ_{i+1} is contained in \hat{H}_l . It follows that the largest coefficient of $M_{22}(f)$ determined by α is comparable to the largest coefficient of $\hat{M}_l(f)$ determined by α . Condition (3) now follows as in the previous cases.

After finitely many such steps we have defined (G_k) satisfying (1)–(3). We say that a vertex of G is mixed if its link intersects more than one stratum of (G_k) . By the fourth item in the statement of Theorem 5.1.5 of [BFH00], $f(v) \in \text{Fix}(f)$ for each vertex v . Redefine the simplicial structure on G by eliminating from the zero skeleton any valence two vertex of G that is not mixed, is not the image of a mixed vertex, and is not the image of a vertex of valence greater than two. In this new structure, G has a uniformly bounded number of vertices and so a uniformly bounded number of edges. Properties (1)–(3) are not effected by this reverse subdivision, so it remains to verify (4).

There is no loss in replacing G by the component of G_j that contains H_j . We may therefore assume that $j = N$. We first verify (4) in the special case that $g|G_{N-1}: G_{N-1} \rightarrow G_{N-1}$ is a relative train track map.

By Theorem 5.1.5 of [BFH00] there is a relative train track map $g': G' \rightarrow G'$ representing ϕ^{-p} , a stratum G'_m of the associated filtration and a bijection between the components $\{C'_i\}$ of G'_m and the components $\{C_i\}$ of G_{N-1} such that $\mathcal{F}(C'_i) = \mathcal{F}(C_i)$. Let A_i be a free factor representing $\mathcal{F}(C'_i) = \mathcal{F}(C_i)$, so $[\phi^p(A_i)] = [A_i]$, and by conjugating A_i back to $\phi^p(A_i)$ we obtain a well-defined outer automorphism of A_i that we denote by $\phi^p | [A_i]$. In our special case, both $g | C_i$ and $g' | C'_i$ are relative train track maps for the outer automorphism $\psi_i := \phi^{-p}|[A_i]$. Thus $\Gamma(g | C_i) = \Gamma(g' | C'_i)$. Lemma 3.2.4 of [BFH00] implies that if Λ_j^\pm is a lamination pair for ϕ^p and if Λ_j^+ is carried by A_i , then Λ_j^- is carried by A_i . The expansion factor for the action of ψ_i on any such Λ_j^- is the same as the expansion factor for the action of ϕ^{-p} on Λ_j^- . Thus the elements of $\Gamma(g | C_i)$ are the expansion factors for the action of ϕ^{-p} on those laminations Λ_j^- whose paired Λ_j^+ is carried by A_i . It follows that the elements of $\Gamma(\phi^{-p})$ and the elements of $\Gamma(g)$ differ only in that μ^p

is replaced by the finite collection of eigenvalues corresponding to irreducible weak strata in H_N . At least one of these eigenvalues must be greater than or equal to μ^p . Lemma 2.3 completes the proof in the special case.

We reduce to the special case as follows. For each component C_i of G_{N-1} , $g|_{C_i}$ determines an outer automorphism of some lower rank free group; let $h_i: X^i \rightarrow X^i$ be a relative train track map representing this outer automorphism. Define $X = \bigcup X_i$ and $h: X \rightarrow X$ by $h|_{X_i} = h_i$. There is a sequence of Whitehead moves that deform G_{N-1} to X . These extend to Whitehead moves on G in the obvious way to produce a deformation of G to a marked graph G^* with X as a subgraph and with a natural bijection between the edges of $G^* \setminus X$ and the edges of H_N . Let $\alpha: G \rightarrow G^*$ be the induced marking-preserving homotopy equivalence and let $\beta: G^* \rightarrow G$ be its homotopy inverse. Let $\hat{g} = \alpha \circ g \circ \beta: G^* \rightarrow G^*$. Then

- \hat{g} represents ϕ^{-p} ;
- $M_{NN}(\hat{g}) = M_{NN}(g)$;
- $\hat{g}|_X$ is homotopic to a relative train track map.

The homotopy of the third item can be performed while maintaining the second, and this completes the proof. □

4. THE MAIN PROPOSITION

We will prove Proposition 1.3 by working in a more general context, one that is suited to our induction technique. We have already started laying the groundwork for this by introducing filtrations. Our main technical result, Proposition 4.1, is stated below. In this section we show that it implies Proposition 1.3 and that it and Lemma 3.1 imply Theorem 1.2. The proof of Proposition 4.1 occupies the rest of the paper.

Suppose that G and G' are marked graphs with filtrations of length N and that $f: G \rightarrow G'$ respects the filtrations. We say that $f: G \rightarrow G'$ is *supported on $G(b, a)$* if:

- (s1): $f|_{G_{a-1}}: G_{a-1} \rightarrow G'_{a-1}$ is a simplicial homeomorphism.
- (s2): $f|_{Cl(G \setminus G_b)}: Cl(G \setminus G_b) \rightarrow Cl(G' \setminus G'_b)$ fails to be a simplicial homeomorphism only in that it may identify distinct vertices in G_b . Moreover, f maps vertices in $G \setminus G_b$ to vertices in $G' \setminus G'_b$.

In the special case that $b = a$ we say that f is *supported on a single stratum*.

The *frontier of G_b* is the set of vertices of G_b whose link contains an edge in $G(N, b + 1)$ or equivalently in some stratum higher than b . For $a \leq b$, define $Fr(G, b, a)$ to be the set of vertices in the frontier of G_b that are not contained in G_{a-1} . If f is supported on $G(b, a)$, then $Fr(G', b, a) \subset f(Fr(G, b, a))$ and hence $|Fr(G', b, a)| \leq |Fr(G, b, a)|$. There are two ways that this inequality can be strict. One is that f identifies a pair of vertices in $Fr(G, b, a)$, and the other is that f identifies a vertex in $Fr(G, b, a)$ with a vertex in G_{a-1} .

Proposition 4.1. *Suppose that G and G' are marked graphs of rank n with filtrations of length N and that $f: G \rightarrow G'$ respects the filtrations and is supported on $G(b, a)$. Suppose further that :*

- f restricts to an immersion on each edge;
- $T(f|_{G_i}, v) \geq 2$ for all vertices $v \in G_i$ and all i ;
- G has at most $V_n/2$ vertices with $T(f, v) = 2$.

Then there is a homotopy inverse $g: G' \rightarrow G$ that respects the filtrations such that

- (1): $\text{mlog}(LC(M_{rs}(g))) \prec \text{mlog}(LC(M_{rs}(f)))$ for all $1 \leq r \leq s \leq N$;
- (2): $g|_{G'_{a-1}}: G'_{a-1} \rightarrow G_{a-1}$ is a simplicial homeomorphism;
- (3): if $|\text{Fr}(G, b, a)| = |\text{Fr}(G', b, a)|$, then g is supported on $G'(b, a)$.

Proposition 4.1 implies Proposition 1.3. Suppose that $d(G, G') = L(f)$ for $f: G \rightarrow G'$. We want to apply Proposition 4.1 with $N = a = b = r = s = 1$. The first and third item in the hypothesis are satisfied, so consider the second. If $T(f, v) = 1$, then there is a vertex v in G and an edge e' in G' so that the edge path $f(e)$ begins with e' for every edge e in G incident to v . There is an obvious homotopy of f that slides $f(v)$ across e' . Denote the resulting map by $f': G \rightarrow G'$. Then $L(f') < L(f)$ in contradiction to our choice of f . This verifies the second item. Let $g: G' \rightarrow G$ be the homotopy inverse produced by Proposition 4.1, and let $g': G' \rightarrow G$ be obtained from g by a homotopy relative to the vertices of G' that tightens the g -image of each edge into an immersed path. The entries in $M(g')$ are all less than or equal to the corresponding entries of $M(g)$. Conclusion (1) of Proposition 4.1 therefore implies that $d(G', G) \prec d(G, G')$. Symmetry completes the proof. \square

Lemma 3.1 and Proposition 4.1 imply Theorem 1.2. Given $\phi \in \text{Out}(F_n)$ with paired expansion factors λ and μ , let $f: G \rightarrow G$, p and $G(j, j)$ be as in Lemma 3.1. Apply Proposition 4.1 to produce a homotopy inverse $g: G \rightarrow G$ such that

$$\text{mlog}(LC(M_{jj}(g))) \prec \text{mlog}(LC(M_{jj}(f))).$$

By Lemma 3.1,

$$p \log \mu \prec \text{mlog}(LC(M_{jj}(g)))$$

and

$$\text{mlog}(LC(M_{jj}(f))) \sim p \log \lambda.$$

Thus $\log \mu \prec \log \lambda$. The symmetric argument shows that $\log \lambda \prec \log \mu$ and hence that $\log \mu \sim \log \lambda$ as desired. \square

5. FOLDING

We ultimately construct $g: G' \rightarrow G$ by factoring $f: G \rightarrow G'$ as a sequence of elementary maps (folds and homeomorphisms) and inverting each one. Most of this is done recursively. In this section we construct g explicitly in two basic cases.

We begin with the easiest case.

Lemma 5.1. *Proposition 4.1 holds for a homeomorphism $f: G \rightarrow G'$.*

Proof. If f is simplicial, then $g = f^{-1}$. In general, f is a finite subdivision of $G(b, a)$ followed by a simplicial homeomorphism. The inverse h of a finite subdivision is a simplicial map that collapses finitely many edges, and we define $g = hf^{-1}$. Thus g is supported on $G'(b, a)$ and $LC(M(g)) = 1$. \square

We set notation for the folding method of Stallings [Sta83] as follows.

Notation 5.2. Suppose that $f: G \rightarrow G'$ is a homotopy equivalence of marked graphs and that f restricts to an immersion on each edge. Suppose further that

- (F1): \hat{e}_1 is a non-trivial initial segment of an edge e_1 in G with initial vertex v_0 and terminal vertex v_1 .
- (F2): \hat{e}_2 is a non-trivial initial segment of an edge e_2 in G with initial vertex v_0 and terminal vertex v_2 .

(F3): $f(\hat{e}_1) = f(\hat{e}_2)$ is the maximal common initial subpath of $f(e_1)$ and $f(e_2)$.

Then $f = f_1 \circ p$ where $p: G \rightarrow G^*$ is the quotient map that identifies \hat{e}_1 with \hat{e}_2 and $f_1: G^* \rightarrow G'$ is the induced homotopy equivalence. The map p is called a fold.

Assume now that f is supported on $G(b, a)$. After interchanging e_1 and e_2 if necessary we may assume that

(F4): $e_1 \subset G(b, a)$ and $e_2 \subset G_b$.

(F5): If e_2 is contained in G_{a-1} , then $\hat{e}_2 = e_2$.

(F6): If $\hat{e}_1 = e_1$, then $\hat{e}_2 = e_2$ and $v_1 \notin G_{a-1}$; moreover if $v_1 \in \text{Fr}(G, b, a)$, then either $v_2 \in \text{Fr}(G, b, a)$ or $v_2 \in G_{a-1}$.

There are several places in the proof where we factor $f: G \rightarrow G'$ as a finite composition. The following lemma is used to equip the intermediate graphs with filtrations.

Lemma 5.3. *Suppose that G and G' are filtered marked graphs, that $f: G \rightarrow G'$ respects the filtrations and that $G \xrightarrow{p} G^* \xrightarrow{f_1} G'$ are as in Notation 5.2. Suppose further that*

- (1) f restricts to an immersion on each edge;
- (2) $T(f|_{G_i}, v) \geq 2$ for all vertices $v \in G_i$ and for all i ;
- (3) G has at most $V_n/2$ vertices with $T(f, v) = 2$.

Then $G_i^* = p(G_i)$ defines a filtration \mathcal{F}^* on G^* such that p and f_1 respect the filtrations. Moreover, p and f_1 satisfy (1)–(3).

Proof. This is immediate from the definitions except perhaps for showing that G^* has at most $V_n/2$ vertices with $T(f_1, v) < 3$. If every vertex of G^* is the image of a vertex of G , then this follows from Lemma 2.1. If there is a new vertex w^* of G^* , then $\hat{e}_1 \neq e_1$ and $\hat{e}_2 \neq e_2$. The vertex w^* is the image of the terminal endpoint of \hat{e}_1 and also the image of the terminal endpoint of \hat{e}_2 . There are three edges of G^* incident to w^* . We label them \bar{e}^*, e_1^* and e_2^* where $p(e_i) = e^*e_i^*$ for $i = 1, 2$. The maximality condition on \hat{e}_1 and \hat{e}_2 guarantees that $Df_1(e_1^*), Df_1(e_2^*)$ and $Df_1(\bar{e}^*)$ are all distinct. Thus $T(f_1, w^*) = 3$. \square

The filtration \mathcal{F}^* on G^* defined in Lemma 5.3 is called the push-forward filtration.

We now explicitly invert a single fold $p: G \rightarrow G^*$ that is supported on $G(b, a)$. The construction depends not only on p but also on a and b .

Definition 5.4. There are three cases to consider.

(Case 1: $\hat{e}_1 \neq e_1$ and $\hat{e}_2 = e_2$). In this case G^* is obtained from G by replacing e_1 with an edge e_1^* that has initial endpoint v_2 and terminal endpoint v_1 . Identify $G \setminus \{e_1\}$ with $G^* \setminus \{e_1^*\}$. Then $p(e_1) = e_2e_1^*$ and p is the identity on all other edges. The homotopy inverse $q: G^* \rightarrow G$ is the identity on edges other than e_1^* and satisfies $q(e_1^*) = \bar{e}_2e_1$. We may view q as the fold of an initial segment of e_1^* with all of \bar{e}_2 . In this case q is supported on $G^*(b, a)$.

(Case 2: $\hat{e}_1 \neq e_1$ and $\hat{e}_2 \neq e_2$). By (F5), $e_1, e_2 \subset G(b, a)$. In this case G^* is obtained from G by ‘blowing up’ v_0 . More precisely, we add a vertex w^* and an edge e^* connecting v_0 to w^* ; the edge $e_i, i = 1, 2$ is replaced by an edge e_i^* that initiates at w^* and terminates at v_i . Identify $G \setminus \{e_1, e_2\}$ with $G^* \setminus \{e^*, e_1^*, e_2^*\}$. Then $p(e_i) = e^*e_i^*$ for $i = 1, 2$ and p is the identity on all other edges. The

link $lk(w^*) = \{\bar{e}^*, e_1^*, e_2^*\}$ is contained entirely in $G^*(b, a)$. The homotopy inverse $q: G^* \rightarrow G$ collapses e^* to v , maps e_i^* to e_i for $i = 1, 2$ and is the identity on all other edges. As in the previous case, q is supported on $G^*(b, a)$.

(Case 3: $\hat{e}_1 = e_1$). In this case (F6) applies. The graph G^* is obtained from $G \setminus \{e_1\}$ by identifying v_1 and v_2 ; we label the resulting vertex v^* . The quotient map p satisfies $p(e_1) = e_2$ and is the identity on all other edges except that it identifies the vertices v_1 and v_2 .

By (F6), $|\text{Fr}(G^*, b, a)| < |\text{Fr}(G, b, a)|$ if and only if $v_1 \in \text{Fr}(G, b, a)$. The map p induces an injection (also called p) of $lk(v_1)$ into $lk(v^*)$. If $e_k^* = p(e_k)$ for $e_k \in lk(v_1) \setminus \bar{e}_1$, then $q(e_k^*) = \bar{e}_2 e_1 e_k$. All other edges are mapped by the identity. Thus q is supported on $G^*(b, a)$ if and only if $v_1 \notin \text{Fr}(G, b, a)$.

Lemma 5.5. *Proposition 4.1 holds for single folds $p: G \rightarrow G^*$.*

Proof. We have constructed a homotopy inverse $q: G^* \rightarrow G$ that satisfies (2) and (3) and has $LC(M(q)) = 1$. □

6. BOUNDED PRODUCTS

The following lemma states that the set of maps satisfying Proposition 4.1 is closed under composition of a uniformly bounded number of factors.

Lemma 6.1. *To prove Proposition 4.1 it suffices to prove that there exists $k > 0$ such that every f satisfying the hypotheses of Proposition 4.1 factors as $G = K^0 \xrightarrow{f_1} K^1 \xrightarrow{f_2} \dots \xrightarrow{f_k} K^k = G'$ where the filtration on K^j is defined, inductively, as the push-forward by f_j of the filtration on K^{j-1} and where each $f_j: K^{j-1} \rightarrow K^j$ satisfies the conclusions of Proposition 4.1.*

Proof. Let $g_j: K^j \rightarrow K^{j-1}$ be a homotopy inverse for $f_j: K^{j-1} \rightarrow K^j$ that satisfies (1)–(3) of Proposition 4.1 and define $g = g_1 \circ \dots \circ g_k$. Properties (2) and (3) are preserved under composition and so are satisfied by g .

By Lemma 2.2,

$$LC(M_{rs}(g)) \leq (\text{Edge}_n)^{k-1} \prod_{j=1}^k LC(M_{rs}(g_j))$$

and so by (1) of Proposition 4.1,

$$\text{mlog}(LC(M_{rs}(g))) \prec \sum_{j=1}^k \text{mlog}(LC(M_{rs}(g_j))) \prec \sum_{j=1}^k \text{mlog}(LC(M_{rs}(f_j))).$$

For each $j = 1, \dots, k$ there is an edge $e_{j-1} \subset K^{j-1}(s, r)$ and an oriented edge $e_j^* \subset K^j(s, r)$ that occurs at least $LC(M_{rs}(f_j))/2$ times in $f_j(e_{j-1})$. We may divide e_{j-1} into $(LC(M_{rs}(f_j))/2) - 1$ subintervals each of which has image as an immersed loop that contains e_j^* . The image of each of these immersed loops under $f_k \circ \dots \circ f_{j+1}$ contains an edge of $G(s, r)$. This proves that

$$\text{mlog}(LC(M_{rs}(f_j))) \prec \text{mlog}(LC(M_{rs}(f)))$$

and hence that

$$\text{mlog}(LC(M_{rs}(g))) \prec k \text{mlog}(LC(M_{rs}(f))) \prec \text{mlog}(LC(M_{rs}(f)))$$

as desired. □

The first application of Lemma 6.1 is to prove a special case of Proposition 4.1. In the context of Notation 5.2 F4–F6, where the fold $p: G \rightarrow G^*$ is supported on $G(b, a)$, if $e_2 \subset G_{a-1}$, then we say that $p: G \rightarrow G^*$ folds into lower strata.

Lemma 6.2. *Proposition 4.1 holds under the additional hypothesis that f factors as $G = K^0 \xrightarrow{p_1} K^1 \xrightarrow{p_2} \dots \xrightarrow{p_l} K^l = G'$ where each p_i folds into lower strata (with respect to the push-forward filtrations).*

Proof. The folds p_i are necessarily in the first or third cases of Definition 5.4. The number of edges in $K^i(b, a)$ is less than or equal to the number of edges in $K^{i-1}(b, a)$ with equality if and only if p_i is a case one fold. It follows that there is a uniform bound to the number of case three folds. By Lemma 6.1 and Lemma 5.5 we are reduced to the case that each p_i is a case one fold.

There is a natural bijection between the edges of K^{i-1} and the edges of K^i . With respect to this bijection, $M(p_i)$ has ones on the diagonal and has exactly one non-zero off-diagonal entry in a column corresponding to an edge in $K^{i-1}(b, a)$ and a row corresponding to an edge of K^i_{a-1} . Recall that in the indexing convention for transition matrices, the indices for edges in $K^{i-1}(b, a)$ are all larger than the indices for edges in K^i_{a-1} . The $M(p_i)$'s therefore commute.

Let $q_i: K^i \rightarrow K^{i-1}$ be the inverse of p_i as specified in Definition 5.4, and note that $M(p_i) = M(q_i)$. Define $g = q_l \circ \dots \circ q_1$. Then $M(f) = M(p_l) \cdot M(p_{l-1}) \cdots M(p_1) = M(q_l) \cdot M(q_{l-1}) \cdots M(q_1) = M(g)$. This implies (1) of Proposition 4.1. Conditions (2) and (3) follow from Lemma 5.5 and the observation that these conditions are preserved under composition. \square

The second application of Lemma 6.1 is a reduction to the case that f is supported on $G(j, j)$ for some $1 \leq j \leq N$.

Lemma 6.3. *It suffices to prove Proposition 4.1 with the additional hypothesis that f is supported on a single stratum.*

Proof. We claim that f factors as

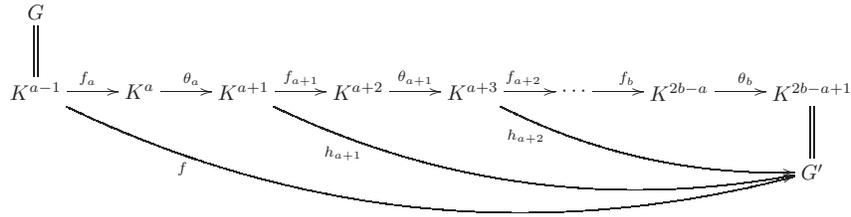
$$G = K^{a-1} \xrightarrow{f_a} K^a \xrightarrow{\theta_a} K^{a+1} \xrightarrow{f_{a+1}} K^{a+2} \xrightarrow{\theta_{a+1}} \dots \xrightarrow{f_b} K^{2b-a} \xrightarrow{\theta_b} K^{2b-a+1} = G'$$

where for $a \leq j \leq b$:

- $f_j: K^{2j-a-1} \rightarrow K^{2j-a}$ is a homotopy equivalence that respects the filtrations and is supported on $K^{2j-a-1}(j, j)$.
- $\theta_j: K^{2j-a} \rightarrow K^{2j-a+1}$ is a homeomorphism that respects the filtrations and is supported on $K^{2j-a}(j, j)$.

Recall the assumption of Proposition 4.1 that f is supported on $G(b, a)$. By [Sta83], $f|_{G_a}: G_a \rightarrow G'_a$ factors as a sequence of folds $\hat{\alpha}_i$ supported on $G(a, a)$ followed by a homeomorphism $\hat{\theta}_a$ that is simplicial except perhaps on the a -stratum. A fold of edges in G_a can be viewed as a fold of edges in G . Thus each $\hat{\alpha}_i$ extends to a fold α_i with support in $G(a, a)$. Define $f_a: G \rightarrow K^a$ to be the composition of the α_i 's. If $\hat{\theta}_a$ is simplicial, let $K^{a+1} = K^a$ and let θ_a be the identity; otherwise K^{a+1} is obtained from K^a by subdividing edges in $K^a(a, a)$, and θ_a is the extension of $\hat{\theta}_a$ over K^a . There is an induced homotopy equivalence $h_{a+1}: K^{a+1} \rightarrow G'$ such that $f = h_{a+1}\theta_a f_a$ and such that h_{a+1} is supported on $K^{a+1}(b, a+1)$. Lemma 5.3 implies that when K^a and K^{a+1} are equipped with the push-forward filtrations, then f_a ,

θ_a and h_{a+1} respect the filtrations. This construction can be repeated, applying it next to h_{a+1} and continuing by induction, to produce the desired factorization of f as shown in the following commutative diagram:



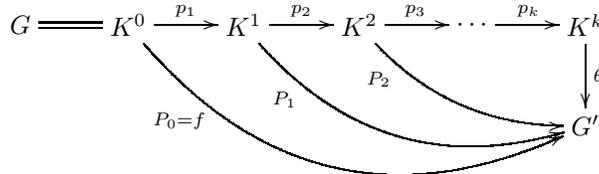
Lemma 5.1 and Lemma 6.1 complete the proof. □

We conclude this section with a third application of Lemma 6.1.

Lemma 6.4. *Proposition 4.1 holds under the additional hypotheses that f is supported on a single stratum H_j and $LC(M_{jj}(f)) < 6$.*

Remark 6.5. At the end of the proof of Lemma 7.3, it is required that the bound $C = 6$ on $LC(M_{jj}(f))$ given above satisfies $C/2 - 1 \geq 2$. This explains our choice of $C = 6$.

Proof. A factorization $G = K^0 \xrightarrow{p_1} K^1 \xrightarrow{p_2} \dots \xrightarrow{p_k} K^k \xrightarrow{\theta} G'$ of f into folds p_i followed by a homeomorphism θ is not unique. For this proof we factor with a preference for those p_i that fold into lower strata. We make this precise as follows. Let $P_0 := f: K^0 \rightarrow G'$. Recall that for any vertex $v \in K^0$ and any oriented edge $e \in lk(v)$, $DP_0(e) \in lk(P_0(v))$ is the first oriented edge in $P_0(e)$. Edges $e_1, e_2 \in lk(v)$ can be folded if and only if $DP_0(e_1) = DP_0(e_2)$. If there is a vertex $v \in K_{j-1}^0$ and an edge $e_1 \in lk(v) \cap K^0(j, j)$ such that $DP_0(e_1) \in G'_{j-1}$, then choose the first fold $p_1: K^0 \rightarrow K^1$ to use e_1 and an edge $e_2 \in K_{j-1}^0$; this is possible because $P^0|_{K_{j-1}^0}: K_{j-1}^0 \rightarrow G'_{j-1}$ is a simplicial homeomorphism. Then $P_0 = P_1 p_1$ where $P_1: K^1 \rightarrow G'$. Iterate this process replacing P_0 with P_1 and so on to produce the indicated factorization, where we denote $\theta p_k \dots p_{i+1}$ by $P_i: K^i \rightarrow G'$ as shown in the following commutative diagram:



Let C_i be the sum of all entries of $M_{jj}(P_i)$. By assumption C_i is uniformly bounded. It is helpful to think of C_i as follows. There is a subdivision of K^i with respect to which P_i is simplicial; we refer to the edges of this subdivision as *edgelets* of K^i . Edgelets that map into $G'(j, j)$ are colored red while all other edgelets are colored white. All the red edgelets are contained in $K^i(j, j)$ and C_i is the number of red edgelets. Two edgelets identified by p_{i+1} must have the same color. Thus $C_{i+1} \leq C_i$ for all i and $C_i < C_{i-1}$ if p_{i+1} identifies any red edgelets. In the latter case we say that p_{i+1} *decreases j -length*. There is a uniform bound to the number of p_{i+1} 's that decrease j -length.

If e is an edge of $K^{i-1}(j, j)$, then $p_i(e)$ is an edge path of length one or two. If the length is two and if both edges in $p_i(e)$ contain red edgelets, then part of e was folded with all of some other edge that contained red edgelets; in particular, some red edgelets of e were identified by p_i . In all other cases the edge paths $p_i(e)$ and e have the same number of edges (either zero or one) that contain red edgelets. We conclude, by induction on m , that if e is an edge of $K^{i-1}(j, j)$ and if the edge path $p_m p_{m-1} \cdots p_i(e)$ contains two edges in $K^m(j, j)$ that contain red edgelets, then p_l decreases j -length for some $i \leq l \leq m$.

If p_i does not fold into lower strata, then $p_i(e) \subset K^i(j, j)$ for each edge $e \subset K^{i-1}(j, j)$. There exists $A \geq 0$ with the following property. If p_l does not fold into lower strata for all $i \leq l \leq i + A$, then there exists an edge e of $K^{i-1}(j, j)$ such that the edge path $p_{i+A} p_{i+A-1} \cdots p_i(e)$ contains two loops in $K^{i+A}(j, j)$ by an argument similar to the one used in the proof of Lemma 6.1. Since every loop in $K^{i+A}(j, j)$ contains red edgelets, p_l decreases j -length for some $i \leq l \leq i + A$.

By Lemma 6.1, Lemma 5.1, Lemma 5.5 and Lemma 6.2 it suffices to show that there is a uniform bound to the number of p_i 's that do not fold into lower strata. It therefore suffices to show that if p_i does not fold into lower strata, then p_l decreases j -length for some $i \leq l \leq i + A$.

We say that an edge e in $K^{i-1}(j, j)$ with initial vertex $v \in K_{j-1}^{i-1}$ is a *frontier edge* and that K^{i-1} has the *frontier red edgelet property* if the initial edgelet of every frontier edge is red. Assume that p_i does not fold into lower strata. Given our preference for folding into lower strata, K^{i-1} must have the frontier red edgelet property. If an initial segment of a frontier edge is folded by p_i , then p_i decreases j -length. If not, then K^i has the frontier red edgelet property. Thus either p_l decreases j -length for some $i \leq l \leq i + A$ or p_l does not fold into lower strata for all $i \leq l \leq i + A$. By our choice of A , p_l decreases j -length for some $i \leq l \leq i + A$. \square

7. IRREDUCIBILITY

A filtration $\hat{\mathcal{F}}$ of length $N + 1$ refines the j -stratum of \mathcal{F} if $\hat{\mathcal{F}}$ is obtained from \mathcal{F} by inserting a new filtration element between G_{j-1} and G_j . More precisely, $\hat{G}_i = G_i$ for $i < j$ and $\hat{G}_i = G_{i-1}$ for $i > j$.

Suppose that $f: G \rightarrow G'$ respects the filtrations \mathcal{F} and \mathcal{F}' and restricts to an immersion on each edge of G . We say that f is (\mathcal{F}, j) -reducible if:

- there is a filtration $\hat{\mathcal{F}}$ that refines the j -stratum of \mathcal{F} ;
- $T(f|\hat{G}_j, v) \geq 2$ for all vertices $v \in \hat{G}_j$;
- $f|\hat{G}_j: \hat{G}_j \rightarrow f(\hat{G}_j)$ is a homotopy equivalence.

If f is not (\mathcal{F}, j) -reducible, then it is (\mathcal{F}, j) -irreducible.

We will apply (\mathcal{F}, j) -irreducibility via the following lemma and corollary.

Lemma 7.1. *Suppose that f is (\mathcal{F}, j) -irreducible. Suppose also that X is a subgraph of G such that $G_{j-1} \subset X$ and $X \subset G_j$ are proper inclusions and such that $T(f|X, v) \geq 2$ for all vertices $v \in X$. Then either $\text{rk}(H_1(f(X))) > \text{rk}(H_1(X))$ or $\text{rk}(H_0(f(X))) < \text{rk}(H_0(X))$.*

Proof. The conditions on X imply that $f(X) \subset G'_j$ properly contains G_{j-1} and has no valence one vertices. If there is a rank-preserving bijection between the components of X and the components of $f(X)$, then $f|X: X \rightarrow f(X)$ is a homotopy equivalence in contradiction to the assumption that f is (\mathcal{F}, j) -irreducible. \square

Corollary 7.2. *Suppose that K^i , $0 \leq i \leq 2n$, are marked graphs of rank n with filtrations \mathcal{F}^i of length N , that $1 \leq j \leq N$, that each $f_i: K^{i-1} \rightarrow K^i$ respects the filtrations and that each f_i is (\mathcal{F}^{i-1}, j) -irreducible. Suppose further that $\gamma \subset K_j^0$ is an immersed loop that is not contained in K_{j-1}^0 and that $f_{2n} \circ \dots \circ f_1|_\gamma$ is an immersion. Then $f_{2n} \circ \dots \circ f_1(\gamma)$ contains $K^{2n}(j, j)$.*

Proof. Let X_i be the smallest subgraph of K^i that contains both K_{j-1}^i and the immersed loop $f_i \circ \dots \circ f_1(\gamma)$. Then $f_{i+1}(X_i) = X_{i+1}$ and if $X_i \neq K_j^i$, then X_i satisfies the hypotheses of Lemma 7.1 with respect to $f_{i+1}: K^i \rightarrow K^{i+1}$. The corollary therefore follows from Lemma 7.1, the fact that $\text{rk}(H_1(X_{2n})) \leq n$ and the fact that $\text{rk}(H_0(X_0)) \leq \text{rk}(H_1(X_0))$. \square

The next lemma is inspired by the following observation. If M is a product of positive integer matrices M_i , then $\text{mlog}(LC(M)) \sim \sum \text{mlog}(LC(M_i))$. This fails without the assumption that M is positive; for example, if the M_i are upper triangular with ones on the diagonal, then (cf. Lemma 6.2) $LC(M) \sim \sum LC(M_i)$.

Lemma 7.3. *Suppose that $f: G \rightarrow G'$ restricts to an immersion on edges, f respects filtrations \mathcal{F} and \mathcal{F}' , f is supported on $G(j, j)$, and f factors as $G = K^0 \xrightarrow{f_1} K^1 \xrightarrow{f_2} \dots \xrightarrow{f_k} K^k \xrightarrow{\theta} G'$ where:*

- (1) θ is a homeomorphism;
- (2) f_i respects filtrations \mathcal{F}^{i-1} , \mathcal{F}^i where $\mathcal{F}^0 = \mathcal{F}$ and θ pushes \mathcal{F}^k forward to \mathcal{F}' ;
- (3) each $LC(M_{jj}(f_i)) \geq 6$;
- (4) each $f_i: K^{i-1} \rightarrow K^i$ has a homotopy inverse $g_i: K^i \rightarrow K^{i-1}$ satisfying properties (2) and (3) of Proposition 4.1 and such that $\log(LC(M_{jj}(f_i))) \succ \text{mlog}(LC(M_{jj}(g_i)))$;
- (5) each f_i is (\mathcal{F}^{i-1}, j) -irreducible.

Let $\theta': G' \rightarrow K^k$ be the homotopy inverse to θ given by Lemma 5.1. Then the homotopy inverse $g = g_1 \circ g_2 \circ \dots \circ g_k \circ \theta': G' \rightarrow G$ satisfies properties (2) and (3) of Proposition 4.1 and

$$\log(LC(M_{jj}(f))) \succ \sum_{i=1}^k \log(LC(M_{jj}(f_i))) \succ \text{mlog}(LC(M_{jj}(g))).$$

Proof. There is no loss in assuming that θ is the identity and that $k > 1$. Items (2) and (3) of Proposition 4.1 are closed under composition and so are satisfied by g .

Choose $1 \leq m_1 \leq k$ so that $LC(M_{jj}(f_{m_1}))$ is maximal. Choose m_2 such that $LC(M_{jj}(f_{m_2}))$ is maximal among all $LC(M_{jj}(f_i))$ with $|i - m_1| > 2n$. Choose m_3 such that $LC(M_{jj}(f_{m_3}))$ is maximal among all $LC(M_{jj}(f_i))$ with $|i - m_1| > 2n$ and $|i - m_2| > 2n$. Continue in this way as long as possible to choose $\{m_1, \dots, m_p\}$. It is clear that

$$\sum_{l=1}^p \log(LC(M_{jj}(f_{m_l}))) < \sum_{i=1}^k \log(LC(M_{jj}(f_i))) < (4n + 1) \sum_{l=1}^p \log(LC(M_{jj}(f_{m_l}))),$$

so

$$\sum_{l=1}^p \log(LC(M_{jj}(f_{m_l}))) \sim \sum_{i=1}^k \log(LC(M_{jj}(f_i))).$$

Lemma 2.2 and item 4 imply that

$$\sum_{i=1}^k \log(LC(M_{jj}(f_i))) \succ \sum_{i=1}^k \text{mlog}(LC(M_{jj}(g_i))) \succ \log(LC(M_{jj}(g))).$$

Reorder the m_l 's so that they are increasing. Define

$$\begin{aligned} \alpha &:= f_{m_1-1} \circ \cdots \circ f_1, \\ \omega &:= f_k \circ \cdots \circ f_{m_p+1}, \end{aligned}$$

and for $1 \leq l \leq p - 1$,

$$h_l := f_{m_{l+1}-1} \circ \cdots \circ f_{m_l+1}.$$

Thus f factors as

$$K^0 \xrightarrow{\alpha} K^{m_1-1} \xrightarrow{h_1 \circ f_{m_1}} K^{m_2-1} \xrightarrow{h_2 \circ f_{m_2}} \cdots \xrightarrow{h_{p-1} \circ f_{m_{p-1}}} K^{m_p-1} \xrightarrow{f_{m_p}} K^{m_p} \xrightarrow{\omega} K^k.$$

For each $1 \leq l \leq p - 1$, there is an edge e_l in $K^{m_l-1}(j, j)$ and an oriented edge e_l^* in $K^{m_l}(j, j)$ that occurs at least $LC(M_{jj}(f_{m_l}))/2$ times in $f_{m_l}(e_l)$. There are $(LC(M_{jj}(f_{m_l}))/2) - 1$ subintervals of e_l with disjoint interior whose f_{m_l} -images are immersed loops $\{\gamma_u\}$ that contain e_l^* . Since $(h_l \circ f_{m_l})$ restricts to an immersion on edges, $h_l|_{\gamma_u}$ is an immersion and Corollary 7.2 implies that each $h_l(\gamma_u)$ contains $K^{m_l}(j, j)$. Thus $(h \circ f_{m_l})(e_{l-1})$ covers every edge of $K^{m_l}(j, j)$ at least $(LC(M_{jj}(f_{m_l}))/2) - 1$ times. Note that by item 3, $(LC(M_{jj}(f_{m_l}))/2) - 1 \geq 2$.

It follows that

$$\log(LC(M_{jj}(f))) \succ \sum_{l=1}^p \log(LC(M_{jj}(f_{m_l}))),$$

which completes the proof. □

8. PROOF OF PROPOSITION 4.1

We now present the proof of Proposition 4.1. Recall that G and G' are marked graphs of rank n with filtrations \mathcal{F} and \mathcal{F}' of length N , $f : G \rightarrow G'$ respects the filtrations, restricts to an immersion on each edge and is supported on $G(a, b)$. Moreover, G has at most $V_n/2$ vertices with $T(f, v) = 2$ and $T(f|_{G_i}, v) \geq 2$ for all vertices $v \in G_i$ and all i . In particular, G has at most V_n vertices and at most $Edge_n$ edges.

The proof of Proposition 4.1 is by downward induction on N . Recall that the maximal length of a filtered marked graph is $2n - 1$. Proposition 4.1 is therefore vacuously satisfied for $N = 2n$, providing the basis step for our downward induction. We assume that $N < 2n$ and that the proposition holds for filtrations with length greater than N . By Lemma 6.3, we may assume that f is supported on a single stratum $G(j, j)$.

As a special case, suppose that f is (\mathcal{F}, j) -reducible. Choose a refinement $\hat{\mathcal{F}}$ of \mathcal{F} as in the definition of (\mathcal{F}, j) -reducible. Define $\hat{\mathcal{F}}'$ by inserting $f(\hat{G}_j)$ between G'_{j-1} and G'_j . Then $f : G \rightarrow G'$ respects $\hat{\mathcal{F}}$ and $\hat{\mathcal{F}}'$ and satisfies the hypotheses of Proposition 4.1; note that f is supported on $G(j, j)$ with respect to the original filtrations and is supported on $G(j + 1, j)$ with respect to the new filtrations. By the induction hypothesis, the conclusions of Proposition 4.1 hold for the transition matrices defined with respect to $\hat{\mathcal{F}}$ and $\hat{\mathcal{F}}'$. But for each r, s , there exists r', s' such

that $M_{rs}(f)$ defined with respect to \mathcal{F} and \mathcal{F}' equals $M_{r's'}(f)$ defined with respect to $\hat{\mathcal{F}}$ and $\hat{\mathcal{F}}'$. This verifies the proposition in the special case.

Sublemma 8.1. *f factors as*

$$G = K^0 \xrightarrow{f_1} K^1 \xrightarrow{f_2} \dots \xrightarrow{f_m} K^m \xrightarrow{\theta} G'$$

where:

- (1) θ is a homeomorphism;
- (2) Proposition 4.1 holds for each $f_i: K^{i-1} \rightarrow K^i$;
- (3) f_i is (\mathcal{F}^i, j) -irreducible and $LC(M_{jj}(f_i)) \geq 6$ for $i \leq m - 1$.

Proof. Begin with a factorization $f = \theta \circ p_w \circ \dots \circ p_1$ where each p_l is a fold and θ is a homeomorphism. Define a finite sequence of integers $0 = k_0 < k_1 < \dots < k_m = w$ as follows. If f is (\mathcal{F}, j) -reducible or if $LC(M_{jj}(f)) < 6$, let $k_1 = w$. Otherwise k_1 is the smallest positive value for which $p_{k_1} \circ \dots \circ p_1$ is (\mathcal{F}, j) -irreducible and for which $LC(M_{jj}(p_{k_1} \circ \dots \circ p_1)) \geq 6$. Define $k_2 > k_1$ similarly, replacing f with $\theta \circ p_w \circ \dots \circ p_{k_1+1}$. This process stops in finite time to produce $0 = k_0 < k_1 < \dots < k_m = w$. Denote G by K^0 and $p_{k_i} \circ \dots \circ p_{k_{i-1}+1}$ by $f_i: K^{i-1} \rightarrow K^i$. Thus $f = \theta \circ f_m \circ \dots \circ f_1$ and it remains to verify item 2.

Suppose at first that $i < m$. Then $f_i = p_{k_i} \circ \hat{f}_i$ where either \hat{f}_i is (\mathcal{F}^{i-1}, j) -reducible or $LC(M_{jj}(\hat{f}_i)) < 6$. Lemma 6.4 and the special case considered above imply that Proposition 4.1 holds for \hat{f}_i . Lemma 5.5 and Lemma 6.1 then imply that Proposition 4.1 holds for f_i . If this argument does not also apply to f_m , then either the f_m is (\mathcal{F}^{m-1}, j) -reducible or $LC(M_{jj}(\hat{f}_m)) < 6$. In either case Proposition 4.1 holds for f_m . □

Let $g_i: K^k \rightarrow K^{k-1}$ be the homotopy inverse of f_i determined by item (2) of Sublemma 8.1. In particular,

$$\text{mlog}(LC(M_{rs}(g_i))) \prec \text{mlog}(LC(M_{rs}(f_i))).$$

For all but uniformly many i , $|\text{Fr}(K^{i-1}, j, j)| = |\text{Fr}(K^i, j, j)|$. We are therefore reduced by Lemma 6.1 to the case that θ is the identity, that each g_i is supported on $K^k(j, j)$ and that item (3) of Sublemma 8.1 is satisfied for $i \leq m$. We may assume without loss that $j = s = N$.

Consider the homotopy commutative diagram

$$K^0 \begin{array}{c} \xrightarrow{f_1} \\ \xleftarrow{g_1} \end{array} K^1 \begin{array}{c} \xrightarrow{f_2} \\ \xleftarrow{g_2} \end{array} \dots \begin{array}{c} \xrightarrow{f_{m-2}} \\ \xleftarrow{g_{m-2}} \end{array} K^{m-2} \begin{array}{c} \xrightarrow{f_{m-1}} \\ \xleftarrow{g_{m-1}} \end{array} K^{m-1} \begin{array}{c} \xrightarrow{f_m} \\ \xleftarrow{g_m} \end{array} K^m$$

and for $l = 1, \dots, m$, denote the homotopy inverses by

$$\begin{aligned} \hat{g}_l &= g_l \circ \dots \circ g_m: K^m \rightarrow K^{l-1}, \\ \hat{f}_l &= f_m \circ \dots \circ f_l: K^{l-1} \rightarrow K^m. \end{aligned}$$

We set notation as follows. For $h: X \rightarrow Y$ supported on $X(j, j)$, let $A(h) = M_{j,j}(h)$ and let $B(h)$ be the submatrix of $M(h)$ whose columns correspond to $X(j, j)$ and whose rows correspond to $Y(j - 1, r)$. Note that

$$LC(M_{rs}(h)) = \max\{LC(A(h)), LC(B(h))\}.$$

Let $g = g_1 \circ g_2 \circ \dots \circ g_m: G' \rightarrow G$. Lemma 7.3 implies that $\log(LC(A(g))) \prec \log(LC(A(f)))$, and we are reduced to showing that $\log(LC(B(g))) \prec \log(LC(M_{rs}(f)))$.

Our first estimate follows from Lemma 2.2 and Lemma 7.3:

$$\begin{aligned} \text{mlog}(LC(A(\hat{g}_{l-1}))) &< \sum_{i=l-1}^m \text{mlog}(LC(A(g_i))) \\ &< \sum_{i=l-1}^m \log(LC(A(f_i))) \\ &< \sum_{i=1}^m \log(LC(A(f_i))) \\ &< \log(LC(A(f))). \end{aligned}$$

For our next estimate, it is convenient to work without log. Keep in mind that $\log(x) \prec \log(y)$ if and only if $x < y^c$ for some uniform constant $c \geq 1$. By definition,

$$B(g) = B(g_m) + \sum_{l=1}^{m-1} B(g_l)A(\hat{g}_{l+1}).$$

Thus for some uniform constant $c \geq 1$,

$$\begin{aligned} LC(B(g)) &\leq LC(B(g_m)) + Edge_n \cdot \left(\sum_{l=1}^{m-1} LC(A(\hat{g}_{l+1})) \cdot (LC(B(g_l))) \right) \\ &\leq Edge_n \cdot LC(A(f))^c \cdot \left(\sum_{l=1}^m LC(B(g_l)) \right) \\ &\leq Edge_n \cdot LC(A(f))^c \cdot \left(\sum_{l=1}^m LC(M_{rs}(f_l))^c \right) \\ &\leq Edge_n \cdot LC(A(f))^c \cdot \left(\sum_{l=1}^m LC(M_{rs}(f_l)) \right)^c. \end{aligned}$$

We restate this back in terms of log and continue, using both Lemma 7.3 and the fact that $B(f) = B(f_1) + \sum_{l=2}^m B(f_l)A(\hat{f}_{l-1})$ to verify the next to the last inequality:

$$\begin{aligned} \text{mlog}(LC(B(g))) &< \log(LC(A(f))) + \log \left(\sum_{l=1}^m LC(M_{rs}(f_l)) \right) \\ &< \log(LC(A(f))) + \log \left(\sum_{l=1}^m LC(A(f_l)) + \sum_{l=1}^m LC(B(f_l)) \right) \\ &< \log(LC(A(f))) + \sum_{l=1}^m \log(LC(A(f_l))) + \log \left(\sum_{l=1}^m LC(B(f_l)) \right) \\ &< \log(LC(A(f))) + \log(LC(A(f))) + \log(LC(B(f))) \\ &< \log(LC(M_{rs}(f))), \end{aligned}$$

which completes the proof.

REFERENCES

- [Ali02] Emina Alibegović, *Translation lengths in $\text{Out}(F_n)$* , *Geom. Dedicata* **92** (2002), 87–93. MR1934012 (2003m:20045)
- [BFH00] M. Bestvina, M. Feighn, and M. Handel, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, *Ann. of Math.* **151** (2000), no. 2, 517–623. MR1765705 (2002a:20034)
- [BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, *Ann. of Math.* **135** (1992), 1–51. MR1147956 (92m:20017)
- [HM06] M. Handel and L. Mosher, *Parageometric outer automorphisms of free groups*, *Trans. Amer. Math. Soc.*, this issue.
- [LL03] G. Levitt and M. Lustig, *Irreducible automorphisms of F_n have north-south dynamics on compactified outer space*, *J. Inst. Math. Jussieu* **2** (2003), no. 1, 59–72. MR1955207 (2004a:20046)
- [Sta83] J. Stallings, *Topology of finite graphs*, *Inv. Math.* **71** (1983), 551–565. MR0695906 (85m:05037a)
- [Vog02] K. Vogtmann, *Automorphisms of free groups and outer space*, *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, vol. 94, 2002, pp. 1–31. MR1950871 (2004b:20060)

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, LEHMAN COLLEGE - CUNY, 250 BEDFORD PARK BOULEVARD W, BRONX, NEW YORK 10468

E-mail address: michael.handel@lehman.cuny.edu

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, RUTGERS UNIVERSITY AT NEWARK, NEWARK, NEW JERSEY 07102

E-mail address: mosher@andromeda.rutgers.edu