QUASI-GEODESICS IN $\text{Out}(\mathbb{F}_n)$ AND THEIR SHADOWS IN SUB-FACTORS

YULAN QING AND KASRA RAFI

Abstract. We study the behaviour of quasi-geodesics in $\text{Out}(\mathbb{F}_n)$. Given an element ϕ in $\text{Out}(\mathbb{F}_n)$ there are several natural paths connecting the origin to ϕ in $\text{Out}(\mathbb{F}_n)$; for example, paths associated to sequences of Stallings folds and paths induced by the shadow of greedy folding paths in Outer Space. We show that none of these paths is, in general, a quasi-geodesic in $\text{Out}(\mathbb{F}_n)$. In fact, in contrast with the mapping class group setting, we construct examples where any quasi-geodesic in $\text{Out}(\mathbb{F}_n)$ connecting ϕ to the origin will have to back-track in some free factor of $\mathbb{F}_n$.

1. Introduction

Let $\mathbb{F}_n = \langle s_1, s_2, ..., s_n \rangle$ denote the free group of rank n and let $\text{Out}(\mathbb{F}_n)$ denote the group of outer automorphisms of $\mathbb{F}_n$,

$$\text{Out}(\mathbb{F}_n) := \text{Aut}(F_n)/\,\text{Inn}(F_n).$$

This group is finitely presented [Nie24]. For example, it can be generated by the set of right transvections and left transvections:

$$\begin{cases} s_i \to s_i s_k \\ s_j \to s_j \qquad \text{for all } j \neq i \end{cases} \qquad \begin{cases} s_i \to s_k s_i \\ s_j \to s_j \qquad \text{for all } j \neq i \end{cases}$$

together with elements permuting the basis and elements sending some s_i to its inverse $\overline{s}_i$. For an explicit presentation see [McC89]. Equip $\text{Out}(\mathbb{F}_n)$ with the word metric associated to this generating set. We aim to understand the geometry $\text{Out}(\mathbb{F}_n)$ as a metric space. Specifically, we want to understand the quasi-geodesics in $\text{Out}(\mathbb{F}_n)$.

A common way to generate a path connecting a point in $\text{Out}(\mathbb{F}_n)$ to the identity is to use the Stallings' Folding Algorithm [Sta83]. One can create a model for $\text{Out}(\mathbb{F}_n)$ by considering the space of graphs x of rank n where the oriented edges are labeled by elements of $\mathbb{F}_n$ inducing an isomorphism $\mu\colon \pi_1(x) \to \mathbb{F}_n$ (defined up to conjugation and graph automorphism). We refer to μ as the *marking map*. A Stallings fold is a quotient map from x to another labeled graph x',

$$\text{fold}\colon x \to x',$$

where edges are identified according to their labels (see Section 2.4 for details). Stallings Folding theorem provides an algorithm for finding a sequence of folds connecting any two such labeled graphs.

Let R_0 be a rose with labels $s_1, \dots, s_n$ inducing an isomorphism

$$\rho_0\colon \pi_1(R_0) \to \mathbb{F}_n,$$

Date: June 25, 2018.

and, for $\phi \in \mathrm{Out}(\mathbb{F}_n)$, let $R = \phi(R_0)$ be the rose where the marking map

$$\mu\colon \pi_1(R) \to \mathbb{F}_n \qquad \text{is given by} \qquad \mu = \phi\, \mu_0.$$

Then, a sequence

$$R = x_m \to x_{m-1} \ldots \to x_1 \to x_0 = R_0$$

of Stallings folds produces a path in $\mathrm{Out}(\mathbb{F}_n)$ connecting the identity to ϕ as follows: for $0 < i \le m$, consider a quotient map $x_i \to R_i$ by collapsing a spanning sub-tree, resulting in a rose R_i. Also, let $q_i\colon x_i \to x_0 = R_0$ be the composition of the folding maps. Then, we have the following diagram of homotopy equivalences:

$$R_0 \to R_i \leftarrow x_i \to R_0$$

where the first arrow is any graph automorphism between the two roses. We associate x_i to the induced map $\phi_i\colon \pi_1(R_0) \to \pi_1(R_0)$. The map $\phi_i \in \mathrm{Out}(\mathbb{F}_n)$ is coarsely well defined, depending on the chosen quotient maps $x_i \to R_i$ and the graph automorphism $R_0 \to R_i$. But the set of all possible resulting maps has a uniformly bounded diameter in $\mathrm{Out}(\mathbb{F}_n)$ and the distance between ϕ_i and ϕ_{i+1} in $\mathrm{Out}(\mathbb{F}_n)$ is uniformly bounded. That is, we have a coarse path

$$\phi = \phi_m, \ldots, \phi_0 = \mathrm{id}$$

in $\mathrm{Out}(\mathbb{F}_n)$ connecting ϕ to the identity which we refer to as a *Stallings folding paths* in $\mathrm{Out}(\mathbb{F}_n)$.

However, Stallings folding paths do not in general give efficient paths in $\mathrm{Out}(\mathbb{F}_n)$.

Theorem A. *For any $K, C > 0$, there exists an element $\phi \in \mathrm{Out}(\mathbb{F}_n)$ such that any Stallings folding path connecting $R = \phi(R_0)$ to R_0 does not yield a (K, C)–quasi-geodesic path in $\mathrm{Out}(\mathbb{F}_n)$.*

Proof. Consider the following automorphism

$$\phi\colon \begin{cases} a \to & a \\ b \to & b \\ c \to & c\,(ab^s)^t \end{cases}$$

That is, if R_0 is the rose with edge labels a, b and c, then $R = \phi(R_0)$ has labels a, b, and $(ab^s)^t$. Following Stallings folding algorithm, to go from R to R_0 we need to fold the third edge around the edges labeled a and b. At each step, there is only one fold possible. The first and second edges remain unchanged since they are also present in R_0. This takes $t\ (s+1)$ steps. That is, the associated path in $\mathrm{Out}(\mathbb{F}_n)$ has a length comparable to $t\,s$.

However, one can see that there is another path connecting R to R_0 that takes $2s + t$ steps, namely:

$$\big\langle a, b, c(ab^s)^t \big\rangle \longrightarrow \big\langle ab^s, b, c(ab^s)^t \big\rangle \longrightarrow \big\langle ab^s, b, c \big\rangle \longrightarrow \langle a, b, c \rangle$$

Choosing s, t sufficiently large compared with the given K and C, we have shown that the path given by Stallings folds was not a (K, C)–quasi-geodesic. □

One can also represent an element of $\mathrm{Out}(\mathbb{F}_n)$ using *train-track maps* (see [BH92]) and consider a folding sequence according to the train-track structure. Or similarly, consider a geodesic in Culler-Vogtmann Outer Space and take the shadow of it to $\mathrm{Out}(\mathbb{F}_n)$[BF14]. The same example above shows that none of these paths would, in general, produce a quasi-geodesic in $\mathrm{Out}(\mathbb{F}_n)$.

On the other hand, we observe in the above example, that even along the shorter paths, it takes at least s steps to form or to eliminate an s-th power of b and t steps to eliminate the t-th power of (ab^s). We examine this phenomenon through the language of *relative twisting number* [CP12].

The *relative twisting number* [CP12] of two labeled graphs x and x' around a loop α measures the difference between x and x' from the point of view of α and is denoted by $\mathrm{tw}_\alpha(x, x')$ (See Definition 3.1). We show that, the length of any path in $\mathrm{Out}(\mathbb{F}_n)$ connecting ϕ to the identity is bounded below by the relative twisting number $\mathrm{tw}_\alpha(R, R_0)$ where $R = \phi(R_0)$.

Again, instead of considering a path in the Cayley graph $\mathrm{Out}(\mathbb{F}_n)$, we consider a general folding sequence $R_m, \ldots, R_0$. This is a sequence of labeled graphs where R_i is obtained from R_{i+1} by a fold that it is not necessary coming from Stallings' algorithm or any train-track structure (see Section 2.4). We also show that if the length of the loop α remains long along the path, it takes longer to twist around α.

Theorem B. *For any general folding sequence $R_m, \ldots, R_0$ and any loop α, we have*

$$m \geq \mathrm{tw}_\alpha(R_0, R_n).$$

Further, if $\ell_{R_i}(\alpha) \geq L > 50$ for every i, then

$$m \geq \mathrm{tw}_\alpha(R_0, R_n)\left(\log_5 \frac{L}{50}\right).$$

One might suspect that, in the above theorem, $\log_5 L$ can be replaced with L. However, we will show that the above inequality is sharp with an example (see Example 3.7).

Theorem B can be viewed in the context of an attempt to have a distance formula for the word length of an element in $\mathrm{Out}(\mathbb{F}_n)$ in analogy with the work of Masur-Minsky in the setting of the mapping class group [MM00]:

Let $S = S_{g,s}$ be a surface of genus g with s punctures and $\mathrm{Mod}(S)$ denote the mapping class group of S, that is, the group of orientation preserving self-homeomorphisms of S up to isotopy. One can try to understand an element $f \in \mathrm{Mod}(S)$ inductively by measuring the contribution of every subsurface to the complexity of f. This is done explicitly as follows: a marking μ_0 on a surface S is a set of simple closed curves that fill the surface, that is to say, every other curve on S intersects some curve in μ_0. Masur and Minsky introduced a measure of complexity $d_Y(\mu_0, f(\mu_0))$ between μ_0 and $f(\mu_0)$ called the subsurface projection distance. Namely, they defined a projection map

$$\pi_Y \colon C(S) \to C(Y)$$

from the curve graph of S to the curve graph of a sub-surface Y and defined $d_Y(\mu_0, f(\mu_0))$ to be the distance in $C(Y)$ between the projection $\pi_Y(\mu_0)$ and $\pi_Y(f(\mu_0))$ to $C(Y)$. They showed the sum of these subsurface projections is a good estimate for the word length of f (see [MM00] for more details).

To produce the upper-bound for the distance formula, Masur and Minsky constructed a class of quasi-geodesics called *hierarchy paths*, whose lengths is the coarse sum of all subsurface projection distances. An important characteristic of these quasi-geodesics is that they do not back-track in any subsurface Y. That is, there is a quasi-geodesic,

$$[0, m] \to \mathrm{Mod}(S), \qquad i \to f_i,$$

in $\mathrm{Mod}(S)$ so that the projection to the curve graph of Y

$$[0, m] \to \mathrm{Mod}(S), \qquad i \to \pi_Y(f_i(\mu_0))$$

is a quasi-geodesic for every subsurface Y of S.

There are several analogues for the curve graph in the setting $\mathrm{Out}(\mathbb{F}_n)$, most importantly, the free splitting graph $\mathcal{S}(\mathbb{F}_n)$ and the free factor graph $\mathcal{F}(\mathbb{F}_n)$ both have been shown to be Gromov hyperbolic spaces [HM13, BF14]. For every sub-factor $\mathcal{A}$ of $\mathbb{F}_n$, we have projection maps [BF14]

$$\mathrm{Out}(\mathbb{F}_n) \to \mathcal{S}(\mathcal{A}) \to \mathcal{F}(\mathcal{A})$$

and it is known that every quasi-geodesic in $\mathcal{S}(\mathbb{F}_n)$ projects to a quasi-geodesic in $\mathcal{F}(\mathbb{F}_n)$ [KR14]. One may hope to construct quasi-geodesic in $\mathrm{Out}(\mathbb{F}_n)$ where the projections the free splitting or the free factor graph of a sub-factor is always a quasi-geodesic. However, we use Theorem B to prove:

Theorem C. *There is $\phi \in \mathrm{Out}(\mathbb{F}_n)$ and a free factor $\mathcal{A}$, such that every quasi-geodesic connecting ϕ to the identity backtracks in $\mathcal{F}(\mathcal{A})$. In other words, there does not exists a quasi-geodesics between the identity and ϕ that projects to a quasi-geodesics in $\mathcal{F}(\mathcal{A})$. Same is true for $\mathcal{S}(\mathcal{A})$.*

Another application of Theorem B is in the understanding of relationships between Outer Space geodesics and $\mathrm{Out}(\mathbb{F}_n)$ geodesics. The Outer Space, denoted CV_n, is the set of metric graphs whose fundamental group is identified with $\mathbb{F}_n$. It is a CW–complex with $\mathrm{Out}(\mathbb{F}_n)$ action and it was defined by Culler and Vogtmann to study $\mathrm{Out}(\mathbb{F}_n)$ as an analogue of Teichmüller space which has $\mathrm{Mod}(S)$ action [CV86]. One can project a path in CV_n to a path in $\mathrm{Out}(\mathbb{F}_n)$, by considering the associated difference of markings maps along a path in CV_n. Bestvina and Feighn [BF14] showed that greedy folding paths in CV_n projects to quasi-geodesics in free factor graphs for all sub-factors. This was used to produce a weak version of a distance formula which give a lower bound for the word length in terms of projection distance to $\mathcal{S}(\mathcal{A})$ [BBF15].

However, it follows from Theorem C that shadows of greedy folding paths are not quasi-geodesics in $\mathrm{Out}(\mathbb{F}_n)$.

Theorem D. *The shadow of geodesics in CV_n do not in general behave well in $\mathrm{Out}(\mathbb{F}_n)$. More specifically,*

- *(i) There are points $x, y \in \mathrm{CV}_n$ so that there does not exist a geodesic in CV_n from x to y whose shadow in $\mathrm{Out}(\mathbb{F}_n)$ a quasi-geodesic.*
- *(ii) There are points x and y in the thick part of CV_n that are connected by a greedy folding path where the shadow of this greedy folding path in $\mathrm{Out}(\mathbb{F}_n)$ is not a quasi-geodesic.*

Our methods do not say anything about the projection of quasi-geodesics in $\mathrm{Out}(\mathbb{F}_n)$ to $\mathcal{S}(\mathbb{F}_n)$ or $\mathcal{F}(\mathbb{F}_n)$. It is interesting to know if a geodesic in $\mathcal{S}(\mathbb{F}_n)$ can be used as a guide to construct efficient paths in $\mathrm{Out}(\mathbb{F}_n)$.

Question E. *For a given $\phi \in \mathrm{Out}(\mathbb{F}_n)$, does there always exists a quasi-geodesic in $\mathrm{Out}(\mathbb{F}_n)$ connecting ϕ to the identity whose projection to $\mathcal{S}(\mathbb{F}_n)$ is also a quasi-geodesic?*

Acknowledgements. The second author was partially supported by Discovery Grants from the Natural Sciences and Engineering Research Council of Canada (RGPIN 06486).

2. Background

2.1. **Labeled graphs.** Recall from the introduction that a labeled graph x induces an isomorphism $\mu\colon\ \pi_1(x) \to \mathbb{F}_n$ called a *marking*. Two labeled graphs x and x' are equivalent if there is a graph automorphism $f\colon\ x \to x'$ such that the following diagram commutes up to conjugation

$$\begin{array}{ccc} \pi_1(x') & \xrightarrow{\ \mu'\ } & \mathbb{F}_n \\ {\scriptstyle f^\star}\uparrow & \nearrow{\scriptstyle \mu} & \\ \pi_1(x) & & \end{array}$$

Let w denote an element of $\mathbb{F}_n$. We refer to a conjugacy class $[w]$ of w as a *loop*. For any labeled graph x, and any loop α, there is an immersion of a circle in x representing α which (abusing the notation) we also denote by α. We always assume this immersion to be the shortest in its free homotopy class in terms of the number of edges. The number of edges of a loop α in the given marked graph x is denoted $\ell_x(\alpha)$, and is called the *combinatorial length* of α in x.

Sometimes it is more convenient to work with the universal cover of a marked graph. The universal cover of x is an $\mathbb{F}_n$-tree (A simplicial tree with free $\mathbb{F}_n$ action). Given such a tree T, an element $w \in \mathbb{F}_n$ in the conjugacy class α acts hyperbolically on T, and we use $\mathrm{axis}_T(w)$ to denote the axis of its action. Consistent with the definition of combinatorial lengths in the graphs, we use $\ell_T(w)$ to denote the number of edges in a fundamental domain of the action of w. Note that this is independent of the choice of $w \in \alpha$. Hence, we can also use the notation $\ell_T(\alpha)$ which is equal to $\ell_x(\alpha)$.

2.2. **Free Factor and Free Splitting Graphs.** There are several analogues of the curve graph in the setting of Out($\mathbb{F}_n$). The two important ones are the *free factor graph* $\mathcal{F}(\mathbb{F}_n)$ and the free splitting graph $\mathcal{S}(\mathbb{F}_n)$.

A *free factor* $\mathcal{A}$ of $\mathbb{F}_n$ is a subgroup such that there exists another subgroup $\mathcal{B}$ where

$$\mathbb{F}_n = \mathcal{A} * \mathcal{B}.$$

A free factor $\mathcal{A}$ is *proper* in $\mathbb{F}_n$ if the rank of $\mathcal{A}$ is strictly less than n. Then $\mathcal{F}(\mathbb{F}_n)$ is a graph whose vertices are proper free factors of $\mathbb{F}_n$ and edges are pairs of free where one is contained in the other. Similarly for each free factor $\mathcal{A}$, one defines the free factor graph $\mathcal{F}(\mathcal{A})$ whose vertices are proper free factors of $\mathcal{A}$.

A *free splitting* over $\mathbb{F}_n$ is a minimal, simplicial (but possibly not free) action of the group $\mathbb{F}_n$ on a simplicial tree T with trivial edge stabilizer. Then $\mathcal{S}(\mathbb{F}_n)$ is a graph whose vertices are free splittings of $\mathbb{F}_n$ and two splittings are connected by an edge if one can be obtained from the other by a collapse map (see [HM13] for more details). As above, for any free factor $\mathcal{A}$, $\mathcal{S}(\mathcal{A})$ denotes the free splitting graph of $\mathcal{A}$.

There is a projection map

$$\pi\colon\ \mathcal{S}(\mathbb{F}_n) \to \mathcal{F}(\mathbb{F}_n)$$

defined as follows. Let α be a *primitive* loop, that is, $\alpha = [w]$ and $\langle w \rangle$ is a free factor. Then, for any free splitting T, the translation length of α in T, $\ell_T(\alpha)$, can be defined as before but may be zero. We define $\pi(T) = \alpha$ where α is the primitive loop with shortest translation length. Similarly, there is a projection map $\mathcal{S}(\mathcal{A}) \to \mathcal{F}(\mathcal{A})$ which we also denote by π. These maps are coarsely well defined (see [BF14]) and further we have

Theorem 2.1 ([KR14]). *The projection under π of a quasi-geodesic in $\mathcal{S}(\mathcal{A})$ is a re-parameterized quasi-geodesic in $\mathcal{F}(\mathcal{A})$.*

Remark 2.2. The above theorem implies that, given a path p in $\mathrm{Out}(\mathbb{F}_n)$, if its projection to the free factor graph is not a quasi-geodesic, then its projection to the free splitting graph was also not a quasi-geodesic.

Let x be a labeled graph and T be the universal cover. For a proper free factor $\mathcal{A}$ let $T|\mathcal{A}$ be the minimal $\mathcal{A}$–invariant subtree of T. Note that $T|\mathcal{A} \in \mathcal{S}(\mathcal{A})$. Letting T_0 be the universal cover R_0, the labeled rose fixed in the introduction, we define a shadow map as follows:

$$\Theta^{\mathcal{S}}_{\mathcal{A}} \colon\ \mathrm{Out}(\mathbb{F}_n) \to \mathcal{S}(\mathcal{A}) \qquad \phi \to \phi(T_0)|\mathcal{A}.$$

That is, we change the action on T_0 according to ϕ and take the minimal $\mathcal{A}$–invariant subtree. Composing with $\pi \colon \mathcal{S}(\mathcal{A}) \to \mathcal{F}(\mathcal{A})$ we can define a shadow map to $\mathcal{F}(\mathcal{A})$,

$$\Theta^{\mathcal{F}}_{\mathcal{A}} \colon\ \mathrm{Out}(\mathbb{F}_n) \to \mathcal{F}(\mathcal{A}) \qquad \phi \to \pi\big(\Theta^{\mathcal{S}}_{\mathcal{A}}(\phi)\big).$$

The map $\Theta^{\mathcal{F}}_{\mathcal{A}}$ is used more often and hence we shorten the notation to $\Theta_{\mathcal{A}}$. In fact, the shadow to $\mathcal{F}(\mathcal{A})$ makes sense for all marked graphs. For a marked graph x with the universal cover T, we define

$$\Theta_{\mathcal{A}}(x) = \pi(T|\mathcal{A}).$$

We recall the following lemma in [BF14] on the upper bound for the distance in the free factor graph. It states that if a loop $\alpha \in \mathcal{A}$ is short in x then the shadow of x in $\mathcal{F}(\mathcal{A})$ is near α.

Lemma 2.3 (Lemma 3.3 in [BF14]). *Let $\mathcal{A}$ be a proper free factor and α be a primitive element in $\mathcal{A}$. Let x be a marked graph so that $\ell_x(\alpha) \leq L$. Then*

$$d_{\mathcal{F}(\mathcal{A})}\big(\Theta_{\mathcal{A}}(x), \alpha\big) \leq 6L + 13.$$

2.3. **Intersection Core.** To define the relative twisting, we first need to introduce the Guirardel Core associated to a pair of $\mathbb{F}_n$–trees. We give a characterization of the 2–skeleton of the Guirardel Core that is different from (but equivalent to) the one given in [Gui05]. Given an $\mathbb{F}_n$–tree T, let $\mathfrak{o}$ be a fixed vertex of T that we call the base-point. We refer to the vertex $w(\mathfrak{o})$ of T, $w \in \mathbb{F}_n$, simply by w_T and we refer to an edge by a pair of words $(w, ws)_T$ where s is the *label* of the oriented edge. We say the edge $(w, ws)_T$ is an *s-edge*. We say an s–edge $(w, ws)_T$ is *preceded* by a t–edge $(u, ut)_T$ if $w = u$.

There is a one-to-one correspondence between the set of infinite geodesic rays originating from $\mathfrak{o}$ and the set of infinite freely reduced words in $\mathbb{F}_n$. Hence, the Gromov boundary of the group $\partial\mathbb{F}_n$, i.e. the equivalence class of quasi-geodesics rays in a Cayley graph $\mathbb{F}_n$, can be identified with the set of all geodesic rays starting from $\mathfrak{o}$ which in turn can be identified with the Gromov boundary ∂T of T. An (oriented) edge $e = (w, ws)_T$ in T defines a decomposition of $\partial\mathbb{F}_n$ into two sets in the following way: a vertex v_T in T is in *front* of e if the geodesic connecting v_T and $(ws)_T$ does not contain the edge e; likewise a vertex v_T is *behind* e if the geodesic connecting v_T and w_T does not contain e. Let $\partial^+(e)$ be the set of all the geodesic rays originating from $\mathfrak{o}$ that eventually lie in front of e and $\partial^-(e)$ be set of all geodesic rays originating from $\mathfrak{o}$ that eventually lie behind e.

Note that the sets $\partial^{\pm}(e)$ are independent of the choice of $\mathfrak{o}$. Also,

$$\partial^+(e) \cup \partial^-(e) = \partial\mathbb{F}_n \qquad \text{and} \qquad \partial^+(e) \cap \partial^-(e) = \emptyset$$

That is, e induces a partition of $\partial\mathbb{F}_n$.

Definition 2.4. Consider an edge e_1 in T_1 and an edge e_2 in T_2. We say e_1 and e_2 are *boundary equivalent* if they induce the same partition of $\partial\mathbb{F}_n$. In contrast, we say $e_1 \times e_2$ is an *intersection square* if all of the following four intersections, as subsets of $\partial\mathbb{F}_n$, are nonempty:

$$\partial^+(e_1) \cap \partial^+(e_2) \neq \emptyset \qquad \partial^+(e_1) \cap \partial^-(e_2) \neq \emptyset$$
$$\partial^-(e_1) \cap \partial^+(e_2) \neq \emptyset \qquad \partial^-(e_1) \cap \partial^-(e_2) \neq \emptyset$$

Let $\overline{\mathrm{Core}(T_1, T_2)}$ be the sub-complex of $T_1 \times T_2$ that is the union of all intersection squares:

$$\overline{\mathrm{Core}(T_1, T_2)} = \Big\{ e_1 \times e_2 \,\Big|\, e_i \in T_i, \quad e_1 \times e_2 \text{ is an intersection square} \Big\}$$

We define the *intersection core*, or simply the *core* of T_1, T_2, to be:

$$\mathrm{Core}(T_1, T_2) = \overline{\mathrm{Core}(T_1, T_2)}/\mathbb{F}_n$$

It follows from Lemma 3.4 in [CQR17] that the above definition is the same as the definition given by Guirardel.

It is clear from the definition that the core is symmetric: $\mathrm{Core}(T_1, T_2)$ is isomorphic to $\mathrm{Core}(T_2, T_1)$. For an edge $e_2 \in T_2$, the e_2*–slice* of the intersection core is the subtree in T_1 that is a collection of edges that form intersection squares with e_2:

$$C_{e_2}(T_1) = \Big\{ e \in T_1 \,\Big|\, e \times e_2 \text{ is a square in } \overline{\mathrm{Core}(T_1, T_2)} \Big\}.$$

Guirardel showed [Gui05] that the e_2–slice is always convex and finite. That is, if two edges $e, e' \in T/\mathbb{F}_n$ are in a given slice, then all the edges on the geodesic path in $T_1/\mathbb{F}_n$ connecting e and e' are also in the slice.

2.4. **Stallings folding path.** We describe Stallings folding path here as needed in this paper; for full generality, see [Sta83]. Let x be a labeled directed graph. Then the edges of x can be subdivided into *edgelets* where each edgelet is labeled with an element in the fixed basis $\{s_1, \dots, s_n\}$ of $\mathbb{F}_n$ and the concatenation of edgelets into an edge yields the original labeling of the edge. A *Stallings fold* from x to x' is a map from x to x' that identifies two edgelets e_1, e_2 for which both of the following are satisfied:

(i) e_1 and e_2 share the same origin vertex.
(ii) e_1 and e_2 shares the same label

The resulting quotient map from x to x' is a homotopy equivalence respecting the markings. We recall Stallings' Folding Theorem per the context of this paper([Sta83]).

Theorem 2.5 (Stallings' Folding Theorem [Sta83])**.** *For any labeled graph x, there exists a finite sequence of Stallings foldings $x = x_m \to \cdots \to x_1 \to x_0 = R_0$ connecting x to R_0.*

For the labeled rose R in Theorem A, there is a unique Stallings fold path connecting R to R_0 because at every step, there are only two edgelets with the same label and the same original vertex.

2.5. **General Folding for roses.** Here, describe a general folding map and its association to partitions of $\partial\mathbb{F}_n$ which is similar to Stallings fold but one does not need to match the labels. For simplicity, we restrict our attention to cases where the labeled graph is a rose which is sufficient for all our examples.

Definition 2.6. Consider a labeled rose R and choose two edges of R labeled s and t. Let R' be a labeled rose whose edge labels are the same as R except the edge-label t has changed to $\overline{s}t$. We say R' is obtained from R by a *fold*, and write

$$R' = \mathrm{fold}(R, t, s) \quad \text{or simply} \quad R' = \mathrm{fold}(R).$$

We also write $T' = \mathrm{fold}(T, t, s)$ for the equivariant map in the universal covers.

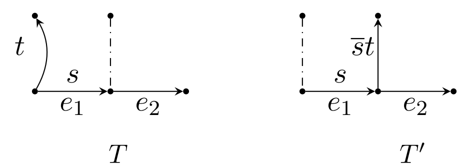

Remark 2.7. In this paper, a fold can occur between any two edges. In contrast, Stallings' folds [Sta83] follow the labeling of the graphs.

Given $\mathbb{F}_n$–trees T, T', and a fixed base-points $\mathfrak{o} \in T$ and $\mathfrak{o}' \in T'$, there is a natural morphism $f\colon T \to T'$ constructed as follows. Send $\mathfrak{o}$ to $\mathfrak{o}'$ and, for $w \in \mathbb{F}_n$, send the vertex w_T in T to the vertex $w_{T'}$ in T'. Also, send an edge $(w, ws)_T$ to the unique embedded edge path connecting $w_{T'}$ to $(ws)_{T'}$ in T'. The morphism f also induces an $\mathbb{F}_n$–equivariant homeomorphism

$$f_\infty\colon \partial T \to \partial T'.$$

Let T and T' be the universal covers of R and R' respectively. A fold from R to R' induces a morphism from T to T' where an edge of the form $(w, wt)_T$ is mapped to the edge path

$$\big[(w, ws)_{T'}, (ws, ws(\overline{s}t))_{T'}\big]$$

and every other edge is mapped to a single edge. Similarly, the edge $(w, ws)_{T'} \in T'$ has two pre-images, $(w, ws)_T$ and $(w, wt)_T$. All other edges have exactly one pre-images. We now describe how partitions given by edges in T differ from that of edges in T'.

Proposition 2.8. *Let $R' = \mathrm{fold}(R, t, s)$ and $f\colon T \to T'$ be the above morphism. Then:*

(i) If an edge e in T is not an s–edge or a t–edge then $f(e) \in T'$ is boundary equivalent to e.

(ii) If $e = (w, wt)_T$ is a t–edge, then $e' = (ws, ws(\overline{s}t))_{T'}$, which is contained in $f(e)$, is boundary equivalent to e.

(iii) If e_2 is an s–edge and e_1 is the t–edge preceding e_2 then $e' = f(e_2)$ partitions $\partial\mathbb{F}_n$ in the following way

$$\partial^-(e') = \partial^-(e_2) \smallsetminus \partial^+(e_1) \qquad \text{and} \qquad \partial^+(e') = \partial^+(e_2) \cup \partial^+(e_1).$$

(iv) In general, for adjacent edges $e_1 = (u, v)_T$, $e_2 = (v, w)_T$ and $e' \subset f(e_2)$, we have

$$\partial^-(e_1) \subset \partial^-(e').$$

Proof. Let $\mathfrak{o}$ and $\mathfrak{o}'$ be the base points in T and T'; $F(\mathfrak{o}) = \mathfrak{o}'$. Note that changing a base-point $\mathfrak{o}$ to a point w in T moves both the partition given by e in T and given by e' in T' by the action of w. Hence, it is sufficient to prove the statement for any desired base-point.

Consider an edge $e' \in T'$ that has only one pre-image under the morphism $f\colon T \to T'$. That is, e is the only edge where $e' \subseteq f(e)$. Then, e and e' are boundary equivalent. To see this, assume $\mathfrak{o} = w_T$. A ray r starting from $\mathfrak{o}$ crosses e if an only if $f(r)$ crosses e'. Hence, a ray is eventually in front of e if and only if it is eventually in front of e' which, by definition, means e and e' are boundary equivalent. This finished the proof of the first two parts.

Let $e_1 = (u, ut)_T$ and $e_2 = (u, us)_T$. We choose $\mathfrak{o} = u$. Then e_1 and e_2 are the only two edges that are mapped over $f(e_1)$. But a ray r in T starting from $\mathfrak{o}$ crosses at most one of e_1 or e_2. In fact, $f(r)$ crosses either edge if and only if it is eventually in front of $f(e_1)$. That is

$$\partial^+ f(e) = \partial^+(e) \cup \partial^+(\hat{e}).$$

The other equality in part (iii) holds because $\partial^- f(e_1)$ is the complement of $\partial^+ f(e_1)$.

To see part (iv), let $\mathfrak{o} = u_T$. If e_1 is not an s–edge, this follows from part (i) or part (ii). Otherwise, $e_1 = (u, us)_T$. The only ray r in $\partial^-(e_1)$ where $f(r)$ crosses $f(e_1)$ is a ray starting with $(u, ut)_T$. Then $f(r)$ starts with $\big[(u, us)_{T'}, (us, us(\bar{s}t))_{T'}\big]$. But e' is different from $(us, us(\bar{s}t))_{T'}$ (which has only one pre-image). Hence, $f(r)$ does not cross e' and lands in $\partial^-(e')$. □

3. Twisting estimate

Let $|G|$ denote the number of edges in the given graph G. We now define *relative twisting number*, which is an analogue of the Masur-Minsky twisting number [MM00]. Recall that, for a given pair of trees T and T_0, a slice $C_{e_0}(T)$ over an edge $e_0 \in T_0$ is the subtree of T consisting of edges that form intersection squares with e_0.

Definition 3.1. Given a loop α and two $\mathbb{F}_n$–trees T, T_0, the *relative twisting number* of $\mathbb{F}_n$–trees T, T_0 around α is

$$\operatorname{twist}_\alpha(T, T_0) = \max_{e_0 \in T_0,\, w \in \alpha} \frac{|\operatorname{axis}_T(w) \cap C_{e_0}(T)|}{\ell_T(\alpha)}.$$

If T and T_0 are universal covers of R and R_0, we define $\operatorname{twist}_\alpha(R, R_0) = \operatorname{twist}_\alpha(T, T_0)$.

Remark 3.2. The twisting number defined above is a rational number. The integer part of $\operatorname{twist}_\alpha(T, T_0)$, which we denote by $\operatorname{tw}_\alpha(T, T_0)$, is equal to the Clay-Pettet definition of relative twisting [CP12].

We show that the relative twisting number changes slowly along loops with large lengths. For a real number $r > 0$, let $[r]$ be the integer part of r and $\{r\}$ be the fractional part of r.

Theorem 3.3. *Let R' be obtained from R by a single fold:*

$$R' = \operatorname{fold}(R, t, s)$$

Then, for any loop α,

$$\operatorname{tw}_\alpha(R', R_0) \geq \operatorname{tw}_\alpha(R, R_0) - 1 \tag{1}$$

where R_0 is as defined before. Furthermore, if $L = \ell_R(\alpha)$, then,

$$\operatorname{twist}_\alpha(R', R_0) \geq \left\lfloor \operatorname{twist}_\alpha(R, R_0) - \frac{2}{L} \right\rfloor + \frac{\left\{\operatorname{twist}_\alpha(R, R_0) - \frac{2}{L}\right\}}{4}. \tag{2}$$

We need to prepare for the proof by establishing a few lemmas. Note that R can also be obtained from R' by a single fold:

$$R' = \operatorname{fold}(R, t, s), \qquad R = \operatorname{fold}(R', \overline{s}t, \overline{s})$$

For the rest of this section, we assume T and T' are universal covers of R and R' respectively and that

$$f\colon\ T \to T' \qquad \text{and} \qquad g\colon\ T' \to T$$

are the morphism associated to these folds. For an embedded edge path $E = [e_1 \dots e_k]$ in T, let $f(E)$ denote the image of E under the morphism f and let $f(E)_w$ denote the embedded edge path that is the intersection of $f(E)$ and $\operatorname{axis}_{T'}(w)$. We call e_1 and e_k the *end edges* of E.

Lemma 3.4. *For any loop α, we have*

$$\ell_T(\alpha) \geq \frac{\ell_{T'}(\alpha)}{2}.$$

Also, if $E = [e_1 \dots e_k]$ is an edge path on the $\operatorname{axis}_T(w)$, *for some $w \in \mathbb{F}_n$, so that both $f(e_1)$ and $f(e_k)$ contain an edge on* $\operatorname{axis}_{T'}(w)$, *then*

$$\big|f(E)_w\big| \geq \frac{|E|}{2}.$$

Proof. Recall that f maps an edge $(u, ut)_T$ to the edge path

$$\big[(u, us)_{T'}, (us, us(\overline{s}t))_{T'}\big]$$

and maps every other edge to one edge. Therefore, for any embedded edge path E in T,

$$\big|f(E)\big| \leq 2|E|. \tag{3}$$

If E is an edge path that realizes $\ell_T(\alpha)$, then $\big|f(E)\big| \geq \ell_{T'}(\alpha)$, thus

$$\ell_T(\alpha) = |E| \geq \frac{\big|f(E)\big|}{2} \geq \frac{\ell_{T'}(\alpha)}{2}.$$

Now assume $E = [e_1 \dots e_k]$ is an edge path on the $\operatorname{axis}_T(w)$ and both $f(e_1)$ and $f(e_k)$ contain an edge on $\operatorname{axis}_{T'}(w)$. Applying, Equation 3 to the morphism g and the edge path $f(E)_w$, we have

$$\big|g(f(E)_w)\big| \leq 2\big|f(E)_w\big|. \tag{4}$$

We need to show $E \subseteq g(f(E)_w)$. In fact, it suffices to show that any end vertex of E is contained in $g(f(E)_w)$.

Let u_T be the first vertex in e_1 (the argument for the last vertex of e_k is similar). If $u_{T'} \in \operatorname{axis}_{T'}(w)$, then $u_{T'} \in f(E)_w$, which means $u_T \in g(f(E)_w)$. Otherwise, $u_{T'} \notin \operatorname{axis}_{T'}(w)$. By assumption, $f(e_1)$ contains an edge on $\operatorname{axis}_{T'}(w)$. It follows that e_1 is mapped to two edges $[e', e'']$ under f, which means e_1 is necessarily a t–edge. Furthermore, the edge e' has two pre-images. This is because, $u_{T'} \notin \operatorname{axis}_{T'}(w)$ and the edge preceding e_1 along the

$\text{axis}_T(w)$ must be mapped over e' as well. But the edges in T' with label $\bar{s}t$ have only one pre-image. Hence,

$$e_1 = (u, ut)_T, \qquad e' = (u, us)_{T'}, \qquad \text{and} \qquad e'' = (us, ut)_{T'}.$$

We then have

$$(us)_{T'}, (ut)_{T'} \in \text{axis}_{T'}(w) \qquad \Longrightarrow \qquad (us)_T, (ut)_T \in g\big(f(E)_w\big).$$

But T is a tree and the vertex u_T necessarily lies on the path connecting $(us)_T$ and $(ut)_T$. Thus $(us)_T, (ut)_T \in g(f(E)_w)$ implies $u_T \in g(f(E)_w)$.

We have shown that the end vertices of E are both in $g(f(E)_w)$, which implies $E \subseteq g(f(E)_w)$. The lemma follows from Equation (4). □

Let $\alpha \in \mathbb{F}_n$ be a loop. Let $e_0 \in T_0$ and $w \in \alpha$ be so that the edge path

$$E := [e_1 e_2 \dots e_k] = \text{axis}_T(w) \cap C_{e_0}(T)$$

is the one realizing the maximum in the definition of $\text{twist}_\alpha(T, T_0)$.

Lemma 3.5. *Assume $|E| \geq 2$. Then $f([e_1 e_2])$ contains an edge in $C_{e_0}(T') \cap \text{axis}_{T'}(w)$. The same holds for $f([e_{k-1} e_k])$.*

Proof. The edge path $f(E)$ is an immersed path in T' that contains $f(E)_w$. Specifically, a point in $f(E)$ is on $f(E)_w$ if and only if it has a unique pre-image in E.

If, for $i = 1$ or 2, $e_i = (u, v)$ and $u\bar{v}$ is not s, $\bar{s}$, t or $\bar{t}$, then $f(e_i)$ is a single edge, has one pre-image and it is boundary equivalent to e_i. Hence, it lies on $C_{e_0}(T') \cap \text{axis}_{T'}(w)$.

If $e_i = (u, ut)$, by the definition of $\text{fold}(T, t, s)$, e_i is mapped to the edge path

$$[(u, us)(us, us(\bar{s}t))].$$

Since the edge $(us, us(\bar{s}t))$ has only one pre-image, it lies on $f(E)_w$ and by Proposition 2.8, it is boundary equivalent to e_i and thus $(us, us(\bar{s}t)) \in C_{e_0}(T')$. Similar argument works when $e_i = (ut, u)$.

There are two remaining cases. Assume $e_1 = (u, us)$ and $e_2 = (us, us^2)$. Then the t–edge preceding e_2, (us, ust), does not lie on E. Hence, $f(e_2)$ has one pre-image in E and thus is on $\text{axis}_{T'}(w)$. Also, by part (iii) of Proposition 2.8,

$$\partial^+(e_2) \subset \partial^+(f(e_2))$$

and by part (iv) of Proposition 2.8,

$$\partial^-(e_1) \subset \partial^-(f(e_2)).$$

Since $e_1, e_2 \in C_T(e_0)$, each of $\partial^+(e_2)$ and $\partial^-(e_1)$ intersects each of $\partial^+(e_0)$ and $\partial^-(e_0)$. Therefore, each of $\partial^+(f(e_2))$ and $\partial^-(f(e_2))$ intersects each of $\partial^+(e_0)$ and $\partial^-(e_0)$. That is, $f(e_2) \in C_{T'}(e_0)$.

The remaining case when $e_1 = (us^2, us)$ and $e_2 = (us, u)$ is identical, except in this case, $f(e_1)$ is in $C_{e_0}(T') \cap \text{axis}_{T'}(w)$. □

Proof of Theorem 3.3. Recall $\alpha \in \mathbb{F}_n$ is a loop and the edge path

$$E := [e_1 e_2 \dots e_k] = \text{axis}_T(w) \cap C_{e_0}(T)$$

is the edge path that realizes the maximum in the definition of $\text{twist}_\alpha(T, T_0)$. By Lemma 3.5, either $f(e_1)$ or $f(e_2)$ contains an edge that is in $C_{e_0}(T') \cap \text{axis}_{T'}(w)$. Call the associated edge in T (either e_1 or e_2), e_{first}. Likewise, one of $f(e_{k-1})$ or $f(e_k)$ has this property. Call the associated edge in T (either e_{k-1} or e_k) e_{last}. The combinatorial length of the edge path $[e_{\text{first}} \dots e_{\text{last}}]$ is at least $k-2$. Let

$$p := \left[\operatorname{twist}_\alpha(T, T_0) - \frac{2}{L}\right] \tag{5}$$

Note that $p \leq \operatorname{tw}_\alpha(T, T_0)$ hence, $\alpha^p(e_{\text{first}})$ lies on the path $[e_{\text{first}} \dots e_{\text{last}}]$ Let E_{remain} be the edge path $[\alpha^p(e_{\text{first}}) \dots e_{\text{last}}]$. Then

$$\operatorname{twist}_\alpha(T', T_0) \geq \frac{f(E)_w}{\ell_{T'}(\alpha)} = \frac{p\ell_{T'}(\alpha) + f(E_{\text{remain}})_w}{\ell_{T'}(\alpha)} = p + \frac{|f(E_{\text{remain}})_w|_{T'}}{\ell_{T'}(\alpha)}. \tag{6}$$

Also, by Lemma 3.4

$$\ell_{T'}(\alpha) \leq 2\ell_T(\alpha) \qquad \text{and} \qquad |f(E_{\text{remain}})_w|_{T'} \geq \frac{|E_{\text{remain}}|_T}{2},$$

and

$$\frac{|E_{\text{remain}}|_T}{\ell_T(\alpha)} \geq \left\{\operatorname{twist}_\alpha(T, T_0) - \frac{2}{L}\right\}.$$

Hence

$$\frac{|f(E_{\text{remain}})_w|_{T'}}{\ell_{T'}(\alpha)} \geq \frac{|E_{\text{remain}}|_T/2}{2\ell_{T'}(\alpha)} \geq \frac{\left\{\operatorname{twist}_\alpha(T, T_0) - \frac{2}{L}\right\}}{4}. \tag{7}$$

Equation 2 follows from Equations (5), (6)and (7).

For $L \geq 2$, Equation 2 implies Equation 1. If $L = 1$, following the proof of Lemma 3.5 we see that if α is not s or $\overline{s}$ then $e_1 = e_{\text{first}}$ and $e_k = e_{\text{last}}$. If α is either s or $\overline{s}$, we still have either $e_1 = e_{\text{first}}$ or $e_k = e_{\text{last}}$. Hence, $f(E)_w$ contains at least $(k-1)$ fundamental domains of action of w and $\operatorname{tw}_\alpha(T', T_0) \geq k-1$. But $\operatorname{tw}_\alpha(T, T_0) = k$. Therefore, Equation 1 still holds. □

Theorem 3.6. *For any general folding sequence $T_m, \dots, T_0$ and any loop α, we have*

$$m \geq \operatorname{tw}_\alpha(T_0, T_m).$$

Further, if $\ell_{T_i}(\alpha) \geq L > 50$ for every i, then

$$m > \operatorname{tw}_\alpha(T_0, T_m)\left(\log_5 \frac{L}{50}\right).$$

Proof. The first assertion of the theorem follows directly from the first assertion of Theorem 3.3. We prove the second assertion.

For a real number $\frac{10}{L} < r$, we have

$$r - \frac{2}{L} > \frac{4r}{5} \qquad \text{and} \qquad \frac{r - \frac{2}{L}}{4} > \frac{r}{5}. \tag{8}$$

For any $0 \leq N < \operatorname{tw}_\alpha(T_0, T_m)$, consider the first time $i = i(N)$ when $\operatorname{tw}_\alpha(T_i, T_m) = N$. By Theorem 3.3, the fractional part of relative twisting is not very small:

$$\left\{\operatorname{twist}(T_i, T_m)\right\} \geq \frac{1 - \frac{2}{L}}{4} \geq \frac{1}{5}.$$

Let $(i + k + 1)$ be first time when $\operatorname{tw}_\alpha(T_{i+k+1}, T_m) = N$. Then

$$\left\{\operatorname{twist}_\alpha(T_{i+k}, T_m)\right\} \leq \frac{10}{L}.$$

That is, in k-steps, the fractional twist has been reduced from above $\frac{1}{5}$ to below $\frac{10}{L}$. Thus, from Proposition 3.3 and Equation (8), we have

$$\frac{\frac{1}{5}}{5^k} < \frac{10}{L} \qquad \Longrightarrow \qquad k > \log_5 \frac{L}{50}.$$

Since this is true for every $0 \leq N < \mathrm{tw}_\alpha(T_0, T_m)$, we have

$$\mathrm{tw}_\alpha(T_0, T_m) < \frac{m}{\log_5 \frac{L}{50}} \qquad \Longrightarrow \qquad m > \mathrm{tw}_\alpha(T_0, T_m) \left(\log_5 \frac{L}{50} \right),$$

which is as desired. □

3.1. **The lower bound is sharp.** At first glance, one might think, that the factor $\log_5 \frac{L}{50}$ in Theorem 3.3 could be replaced with a linear function of L. However, we show that the above estimate is sharp up to a uniform multiplicative error. We now construct, for an arbitrarily large L a folding path $R = R_m, \dots, R_0$ so that

- The length of α at each R_i is at least L.
- m is comparable with $\log L \, \mathrm{tw}_\alpha(R, R_0)$.

Example 3.7. Let R be a rose of rank 5 with edge labels

$$\big\langle (bc)^m a, d\, b, \phi^k(d)c, d, e \big\rangle$$

where $\phi \colon \langle d, e \rangle \to \langle d, e \rangle$ is a strongly irreducible automorphism of the free factor $\langle d, e \rangle$ with exponential growth. Let L be the word length of $\phi^{\lfloor k/2 \rfloor}(d)$ (and also the length of $\phi^{-\lfloor k/2 \rfloor}(d)$) in $\langle d, e \rangle$ which can be chosen to be arbitrarily large by choosing k large enough. Let $\alpha = [bc]$. Let R_0 be a rose with labels

$$\langle a, d\, b, \phi^k(d)c, d, e \rangle.$$

The shortest way to express α in R is

$$bc = (\overline{d}) \cdot (db) \cdot \phi^{-k}(d) \cdot (\phi^k(d)c),$$

where the terms in parentheses are labels of edges in R and is a word of length roughly L^2 in $\langle d, e \rangle$. We have

$$\ell_R(\alpha) \asymp L^2 > L \qquad \text{and} \qquad \mathrm{tw}_\alpha(R, R_0) = m.$$

We now start twisting around the loop $[bc]$, however in a somewhat un-natural way that always keeps the length of α larger than L, using the following steps:

1) Twist around the first half of α, that is to say, cancel b (which is half of α) from $(bc)^m a$ using 2 folds:

$$\overline{(db)} \cdot d \cdot (bc)^m a = c(bc)^{m-1} a.$$

2) Fold $\langle d, e \rangle$ to $\langle \phi^k(d), \phi^k(e) \rangle$; this takes

$$k \, \|\phi\|_{\langle d, e \rangle} \asymp \log L$$

many steps. Note that the immersed loop

$$bc = (\overline{d}) \cdot (db) \cdot \phi^{-k}(d) \cdot (\phi^k(d)c),$$

contains both $\overline{d}$ and $\phi^{-k}(d)$. At any point along this folding path, say after ϕ_i has been applied i times, we have

$$\max\left(\left|\overline{d}\right|_{\left\langle \phi^i(d), \phi^i(e)\right\rangle}, \left|\phi^{-k}(d)\right|_{\left\langle \phi^i(d), \phi^i(e)\right\rangle}\right) \geq |\phi^{\lfloor k/2 \rfloor}(d)| = L.$$

where $|\cdot|_{\left\langle \phi^i(d), \phi^i(e)\right\rangle}$ denotes the word length of an element in the group $\langle d, e\rangle$ in terms of $\phi^i(d)$ and $\phi^i(e)$. Thus the combinatorial lengths of α remains at least L.

3) Twist around the second half of α, that is, cancel c from $c(bc)^{m-1}a$ in 2 folds:

$$\overline{(\phi^k(d)c)} \cdot \phi^k(d) \cdot c(bc)^{m-1}a = (bc)^{m-1}a.$$

4) Fold $\langle \phi^k(d), \phi^k(e)\rangle$ to $\langle d, e\rangle$. Again, the number of steps is comparable to $\log L$ and the length of α remains larger than L.

We now repeat steps 1–4, m–times. Every time the relative twisting around α is reduced by 1. The path has a length of order $m \log L$ as desired.

4. Quasi-Geodesics in $\mathrm{Out}(\mathbb{F}_n)$

In this section we use Theorem B to prove Theorem C. Consider a path $p\colon [0, m] \to \mathcal{X}$ from an intervel $[0, m] \subset \mathbb{R}$ to a metric space $\mathcal{X}$ and, $for i \in [0, m]$, let $p_i = p(i)$. Recall that p is a (K, C)–quasi-geodesic if for any $i, j \in [0, m]$, we have

$$\frac{|i-j| - C}{K} \leq d_{\mathcal{X}}\big(p_i, p_j\big) \leq K\,|i-j| + C.$$

It follows that, any 3 points in the image of a quasi-geodesic satisfy a coarse reverse triangle inequality. Namely, for any $i, j, k \in [0, m]$, $i \leq j \leq k$, we have

$$d_{\mathcal{X}}(p_i, p_j) + d_{\mathcal{X}}(p_j, p_k) \leq K^2 d_{\mathcal{X}}(p_i, p_k) + (K+2)C. \tag{9}$$

We say p is a re-parametrized (K, C)–quasi-geodesic if there is a re-parametrization $\rho\colon [0, m] \to [0, m]$ so that $p \circ \rho$ is a (K, C)–quasi-geodesic. Since the image of p and $p \circ \rho$ are the same, if p is a re-parametrized quasi-geodesic, any 3 points in its image still satisfy the coarse-reverse-triangle inequality given in Equation (9). Often, it is convenient to consider maps from intervals $[0, m]_{\mathbb{Z}}$ in $\mathbb{Z}$ to a metric space $\mathcal{X}$. Then we say $p\colon [0, m]_{\mathbb{Z}} \to \mathcal{X}$ is a re-parametrized quasi-geodesic if it is a restriction of a re-paramterized quasi-geodesic from $[0, m] \to \mathcal{X}$.

We can now re-starte Theorem C explicitly as follows:

Theorem 4.1 (Quasi-geodesics back-track in sub-factors)**.** *For given constants K_1, C_1, K_2, $C_2 > 0$ there exists an automorphism, $\phi \in \mathrm{Out}(\mathbb{F}_n)$ such that, for any (K_1, C_1)–quasi-geodesic $p\colon [0, m] \to \mathrm{Out}(\mathbb{F}_n)$ with $p(0) = id$ and $p(m) = \phi$, the shadow $\Theta_{\mathcal{A}} \circ p$ of p in $\mathcal{F}(A)$ is not a (K_2, C_2)–reparameterized quasi-geodesic.*

Proof. Let $\langle a, b, c\rangle$ be a generating set for $\mathbb{F}_3$, and let R_0 be a rose with labels $\{a, b, c\}$. Let $\psi \in \mathrm{Out}(\mathbb{F}_3)$ be an automorphism defined as:

$$\begin{cases} a \longrightarrow aba \\ b \longrightarrow ab \\ c \longrightarrow c \end{cases}$$

Given the generating set introduced in the introduction, $\|\psi\| = \|\psi^{-1}\| = 2$ where $\|\cdot\|$ represents the word length. Let $\mathcal{A} = \langle a, b\rangle < \mathcal{F}_3$ be a rank 2 free factor. The automorphism

ψ fixes $\mathcal{A}$. We denote the restriction of ψ to $\mathcal{A}$ by $\psi_{\mathcal{A}}$. Then, $\psi_{\mathcal{A}}$ is an irreducible automorphism and acts loxodromicly on the free-factor graph $\mathcal{F}(\mathcal{A})$ of $\mathcal{A}$. Hence, there exists a constant $c_\psi > 0$ so that, for an integer $q > 0$

$$d_{\mathcal{A}}\big(R_0, \psi^q(R_0)\big) \geq c_\psi q. \tag{10}$$

Let α be the loop represented by the word $\psi^q(a)$ and, for a large integer $t > 0$, let $\phi \in \mathrm{Out}(\mathbb{F}_3)$ be the automorphism that twists the element c around α t–times, namely:

$$\begin{cases} a \longrightarrow a \\ b \longrightarrow b \\ c \longrightarrow c(\psi^q(a))^t \end{cases}$$

We first find an upper-bound for the $\|\phi\|$ by constructing a path connecting identity to ϕ. First apply ψ^q so α is represented by one edge in the rose $\psi^q(R_0)$, then twist c around α t–times, and then apply ψ^{-q}. We have

$$\|\phi\| \leq q(\|\psi\| + \|\psi^{-1}\|) + t = 4q + t.$$

Now consider the (K_1, C_1)–quasi-geodesic $p\colon [0, m] \to \mathrm{Out}(\mathbb{F}_3)$ connecting the identity to ϕ. We have

$$m \leq K_1\|\phi\| + C_1 \leq 4K_1\, q + K_1\, t + C_1. \tag{11}$$

If the shadow of p to $\mathcal{A}$ is a (K_2, C_2)–reparatmetrized quasi-geodesic then, the coarse-reserve-triangle-inequality (Equation (9)) holds. That is, for any index i and $R_i = p(i)(R_0)$, we have

$$d_{\mathcal{A}}(R_0, R_i) + d_{\mathcal{A}}(R_i, \phi(R_0)) \leq (K_2)^2 d_{\mathcal{A}}(R_0, \phi(R_0)) + (K_2 + 2)C_2.$$

But ϕ fixes a and b and thus R_0 and $\phi(R_0)$ have the same projection to $\mathcal{A}$. Hence,

$$2d_{\mathcal{A}}(R_0, R_i) \leq (K_2 + 2)C_2.$$

Now, using Equation (10), we get

$$d_{\mathcal{A}}(R_i, \psi^q(R_0)) \geq d_{\mathcal{A}}(R_0, \psi^q(R_0)) - d_{\mathcal{A}}(R_0, R_i) \geq c_\psi q - (K_2 + 2)C_2.$$

By Lemma 3.3 in [BF14], this implies

$$\ell_{R_i}(\alpha) \geq \frac{c_\psi\, q - (K_2 + 2)C_2 - 13}{6} =: L \tag{12}$$

Now, Theorem 3.6 implies that,

$$t = \mathrm{tw}_\alpha(R_0, R_m) \leq \frac{m}{\log_5 \frac{L}{50}}$$

and using Equation (11) we get,

$$t \leq \frac{4K_1\, q + K_1\, t + C_1}{\log_5 \frac{L}{50}}.$$

If we choose q large enough so that

$$\log_5 \frac{L}{50} > 2K_1$$

and then choose t large enough so that

$$\frac{t}{2} > \frac{4K_1\, q + C_1}{\log_5 \frac{L}{50}}$$

we get

$$t \leq \frac{4K_1\, q + C_1}{\log_5 \frac{L}{50}} + \frac{K_1\, t}{\log_5 \frac{L}{50}} < \frac{t}{2} + \frac{t}{2} = t$$

which is a contradiction. The contradiction proves that the shadow of p to $\mathcal{F}(\mathcal{A})$ is not a re-parametrized (K_2, C_2)–quasi-geodesic. □

5. Outer Space

Outer Space CV_n is metric space with $\mathrm{Out}(\mathbb{F}_n)$ action defined as an analogue of the Teichmüller space. See [CV86] for more details. Here, we introduce CV_n briefly and prove Theorem D.

We assume a *graph* is always simple and all vertices have degree 3 or more. A *marked metric graph* (x, f) is a metric graph x together with homotopy equivalence $f\colon R_0 \to x$. The space of all marked metric graphs whose edge lengths sum up to one is called the *Outer Space* [CV86] and is denoted by CV_n. The group $\mathrm{Out}(\mathbb{F}_n)$ acts on CV_n by precomposing the marking: for an element $\phi \in \mathrm{Out}(\mathbb{F}_n)$, $\phi(x, f) = (x, f \circ \phi)$.

Note that we can still think of x as a labeled graph. Recall that $\ell_x(\alpha)$ denotes the combinatorial length of α in x. Let $|e|_x$ denote the metric length of an edge e in x and $|\alpha|_x$ the metric length of α in x, which is the metric length of the immersed loop of the representative of α that realizes its combinatorial length in x.

For a fixed $\epsilon > 0$ define the *thick part* of CV_n to be the set of $x \in \mathrm{CV}_n$ such that

$$|\alpha|_x \geq \epsilon \text{ for every nontrivial conjugacy class } \alpha.$$

A map $h\colon (x, f_x) \to (y, f_y)$ is a *difference of markings* map if $h \circ f_x \simeq f_y$ (homotopy). We will only consider Lipschitz maps and we denote by L_h the Lipschitz constant of h. In many ways it is natural to consider the (asymmetric) Lipschitz metric on CV_n:

$$d(x, y) := \inf_h \log L_h$$

where the infimum is taken over all differences of markings maps. We refer the reader to [FM11, AKB12] for review for some metric properties of $d(\cdot,\cdot)$. In particular, there always exists a non-unique difference of markings map that realizes the infimum. Since a difference of markings map is homotopic rel vertices to a map that is linear on edges, we also use h to denote the representative that realizes the infimum and is linear on edges and refer to such a map as an *optimal map* from x to y. For this section we always assume h is an optimal difference of markings map. Since h is linear on edges, we define

$$\lambda(e) = \frac{|h(e)|_y}{|e|_x}$$

to be the *stretch factor* of an edge e and

$$\lambda(\alpha) = \frac{|\alpha|_y}{|\alpha|_x}$$

to be the the *stretch factor* of a shortest immersed loop that represents α. Define the *tension subgraph*, x_ϕ, or $\mathrm{stretch}(x, y)$, to be the subgraph of x consisting of maximally stretched edges.

Now we restate and prove part (i) of Theorem D:

Theorem 5.1 (Shadow of a geodesic in CV_n is not a quasi-geodesic in $\mathrm{Out}(\mathbb{F}_n)$)**.** *For given constants K and C, there are points $x, y \in \mathrm{CV}_n$ such that for every geodesic $[x, y]_{\mathrm{CV}_n}$ in CV_n connecting x to y, its image in $\mathrm{Out}(\mathbb{F}_n)$ is not $(K, C)-$quasi-geodesic.*

Proof. Consider the same example in rank 3 as in Theorem 4.1 where ϕ is defined as

$$\begin{cases} a \longrightarrow a \\ b \longrightarrow b \\ c \longrightarrow c(\psi^q(a))^t \end{cases}$$

Let $w = c\,(\psi^q(a))^t \in \mathbb{F}_3$ and let M be the length of w in the basis $\langle a, b, c\rangle$. Let $x \in \mathrm{CV}_n$ be a rose where the edges are labeled a, b and w and

$$\ell_x(a) = \ell_x(b) = \frac{1}{M+2} \qquad \text{and} \qquad \ell_x(w) = \frac{M}{M+2}$$

and let $y \in \mathrm{CV}_n$ be a rose where the edges are labeled a, b and c and have length $\frac{1}{3}$ each. Note that the length ratio of a, b and w from x to y are identical,

$$\frac{\ell_y(w)}{\ell_x(w)} = \frac{M/3}{M/(M+2)} = \frac{M+2}{3} = \frac{1/3}{1/(M+2)} = \frac{\ell_y(a)}{\ell_x(a)} = \frac{\ell_y(b)}{\ell_x(b)}.$$

In particular, we have

$$d_{\mathrm{CV}_n(x,y)} = \log \frac{M+2}{3}.$$

In fact, it follows from [QR18] that there is a unique geodesic $[x, y]_{\mathrm{CV}_n}$ in CV_n connecting x to y and it folds along the unique illegal turn. Namely, it folds the edge labeled w around the free factor $\mathcal{A}$ and if $p\colon [0, m] \to \mathrm{Out}(\mathbb{F}_n)$ is the shadow of $[x, y]_{\mathrm{CV}_n}$ in $\mathrm{Out}(\mathbb{F}_n)$, then the projection of p to $\mathcal{A} = \langle a, b\rangle$. Hence, as was seen in the proof of Theorem 4.1, for any K and C, we can choose q and t large enough so that p is not a (K, C)–quasi-geodesics. □

We can also modify the example in Theorem 4.1 to prove part (ii) of Theorem D. For brevity, we do not define greedy folding paths here. They are used in [BF14] in an essential way to prove the hyperbolicity of of the free-factor graph. What we need is that if $\mathrm{stretch}(x, y) = x$ then there is a greedy folding path connecting x to y.

Theorem 5.2. *There are point x and y are in the thick part of CV_n and are connected by a greedy folding path whose shadow in $\mathrm{Out}(\mathbb{F}_n)$ is not a quasi-geodesic.*

Proof. Let $\langle a, b, c\rangle$ be a generating set for $\mathbb{F}_3$, let R_0 be a rose with labels $\{a, b, c\}$. Let ψ and $\psi_{\mathcal{A}}$ be as before. Let $y \in \mathrm{CV}_n$ be the rose R_0 with edge-labels a, b and c and edge-lengths $\frac{1}{3}$.

Let $\mathfrak{g}$ be the axis of $\psi_{\mathcal{A}}$ in $\mathcal{F}(\mathcal{A})$, the free factor graph associated to $\mathcal{A}$. Also let $\alpha_1, \alpha_2, \dots, \alpha_k \in \mathcal{A}$ be primitive loops in $\mathcal{A}$ that, considered as vertices in $\mathcal{F}(\mathcal{A})$, are distance D or more far from $\mathfrak{g}$ for a large constant D. [1] For positive integers $n_1, n_2, \dots, n_k$,

[1] As we shall see, it is enough to have one such loop, however choosing many loops showcases how different the shadow of a quasi-geodesic in $\mathrm{Out}(\mathbb{F}_n)$ to $\mathcal{F}(\mathcal{A})$ could be from being a quasi-geodesic.

let ϕ be the following automorphism:

$$\begin{cases} a \longrightarrow \psi^n(a) \\ b \longrightarrow \psi^n(b) \\ c \longrightarrow c\,\alpha_1^{n_1} \dots \alpha_k^{n_k} \end{cases}$$

We observe that $|\psi^n(a)|_y$ and $|\psi^n(b)|_y$, as a function n, grow at a fixed exponential rate that is less than 3. Therefore for any given loops α_i and powers n_i, there is a power n so that

$$\max\big(|\psi^n(a)|_y, |\psi^n(b)|_y\big) \le |c\,\alpha_1^{n_1} \dots \alpha_k^{n_k}|_y \le 3\,\min\big(|\psi^n(a)|_y, |\psi^n(b)|_y\big), \tag{13}$$

Let x be a rose with edge labels, $\psi^n(a)$, $\psi^n(b)$ and $c\,\alpha_1^{n_1} \dots \alpha_k^{n_k}$ (note that these form a basis for $\mathbb{F}_3$) and edge lengths

$$\frac{|\psi^n(a)|_y}{T}, \qquad \frac{|\psi^n(b)|_y}{T} \qquad \text{and} \qquad \frac{|c\,\alpha_1^{n_1} \dots \alpha_k^{n_k}|_y}{T}.$$

where

$$T = |\psi^n(a)|_y + |\psi^n(b)|_y + |c\,\alpha_1^{n_1} \dots \alpha_k^{n_k}|_y.$$

Then $\text{stretch}(x, y) = x$ and Equation (13) implies that every edge length in x is larger than $\frac{1}{7}$ (x is $\frac{1}{7}$–thick part of CV_n).

Let $r_i = \|\phi_i\|$ be smallest word length of $\phi_i \in \text{Out}(\mathbb{F}_n)$ such that $\phi_i(\alpha_i) = a$. It follows from construction that,

$$\|\phi\| \le 3n + \sum_{i=1}^{k} (2r_i + n_i). \tag{14}$$

Let ℓ_i be the length of α_i in the $\langle a, b\rangle$ basis and let λ be the stretch factor of ϕ. To make Equation (13) holds, we need

$$\lambda^n \sim 1 + \sum_{k=1}^{n} n_i \ell_i.$$

Letting $n_{\max} = \max_i n_i$ we have, for constants c_1 and c_2 depending on k, r_i and ℓ_i, that

$$n \le c_1 \log n_{\max}, \qquad \text{and} \qquad \|\phi\| \le (k+1) n_{\max} + c_2.$$

Now let $p\colon [0, m] \to \text{Out}(\mathbb{F}_n)$ be the shadow of $[x, y]_{\text{gf}}$ to $\text{Out}(\mathbb{F}_n)$. Bestvina-Feighn[BF14] showed that the projection of $[x, y]_{\text{gf}}$ to $\mathcal{F}(\mathcal{A})$ is a quasi-geodesic. Hence the projection stays far from every α_i. And, again by Lemma 3.3 in [BF14], this implies that the combinatorial length of α_i at any point along $[x, y]_{\text{gf}}$ is large, say larger than some constant L depending linearly on D. By Theorem 3.3,

$$m \succ \max_i n_i \log L.$$

But, since p is a quasi-geodesic, $m \prec \|\phi\|$. For large enough L, the above inequality contradicts Equation (14). This implies p cannot be a quasi-geodesic. $\square$

References

[AK11] Yael Algom-Kfir. Strongly contracting geodesics in outer space. *Geom. Topol.*, 15(4):2181–2233, 2011.

[AKB12] Yael Algom-Kfir and Mladen Bestvina. Asymmetry of outer space. *Geom. Dedicata*, 156:81–92, 2012.

[BBF15] Bestvina, Bromberg and Fujiwara. Constructing group actions on quasi-trees and applications to mapping class groups. *Publ. Math. Inst. Hautes tudes Sci.* 122 (2015), 1-64.

[BF12] M. Bestvina and M. Feighn. Subfactor projections. arXiv:1211.1730, November 2012.

[BF14] Mladen Bestvina and Mark Feighn. Hyperbolicity of the complex of free factors. *Adv. Math.*, 256:104–155, 2014.

[BH92] Mladen Bestvina and Handel. Train tracks and automorphisms of free groups. *Ann. of Math.* (2) 135 (1992), no. 1, 1-51.

[CP12] Matt Clay and Alexandra Pettet. Relative twisting in Outer Space. *J. Topol. Anal.*, 4 (2012), 173-201.

[CQR17] Matt Clay, Yulan Qing and Kasra Rafi. Uniform fellow traveling between surgery paths in the sphere graph. *Alg. and Geom. Topol.* , 17 (2017) 3751-3778.

[Cul84] M. Culler. Finite groups of outer automorphisms of a free group. In *Contributions to group theory*, volume 33 of *Contemp. Math.*, pages 197–207. Amer. Math. Soc., Providence, RI, 1984.

[CV86] Marc Culler and Karen Vogtmann. Moduli of graphs and automorphisms of free groups. *Invent. Math.*, 84(1):91–119, 1986.

[DT14] S. Dowdall and S. J. Taylor. Hyperbolic extensions of free groups. *Geom. Topol. to appear*, June 2014.

[FBR16] M Fortier Bourque and K. Rafi. Non-convex balls in the teichmüller metric. to appear in Journal of Differential Geometry, arXiv:1606.05170, 2016.

[FM11] Stefano Francaviglia and Armando Martino. Metric properties of outer space. *Publ. Mat.*, 55(2):433–473, 2011.

[Gui05] Vincent Guirardel. Cur et nombre d?intersection pour les actions de groupes sur les arbres. *Ann. Sci. E?cole Norm. Sup.* (4), 38(6):847–888, 2005.

[HM13] Handel and Mosher. The free splitting complex of a free group, I: hyperbolicity. *Geom. Topol.* 17 (2013), no. 3, 1581-1672.

[Ker83] S.P. Kerckhoff. The Nielsen realization problem. *Ann. of Math. (2)*, 117(2):235–265, 1983.

[Khr85] D. G. Khramtsov. Finite groups of automorphisms of free groups. *Mat. Zametki*, 38(3):386–392, 476, 1985.

[Kra59] S. Kravetz. On the geometry of Teichmüller spaces and the structure of their modular groups. *Ann. Acad. Sci. Fenn. Ser. A I No.*, 278:35, 1959.

[KR14] Kapovich and Kasra Rafi. On hyperbolicity of free splitting and free factor complexes. *Groups Geom. Dyn.* 8 (2014), no. 2, 391-414.

[LR11] Anna Lenzhen and Kasra Rafi. Length of a curve is quasi-convex along a Teichmüller geodesic. *J. Differential Geom.*, 88(2):267–295, 2011.

[MM99] Masur, and Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.* 138 (1999), no. 1, 103-149.

[MM00] H. Masur. and Y. Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geom. Funct. Anal.* 10 (2000), 902-974.

[Nie24] J. Nielsen. Die isomorphism en gruppe der freien Gruppen. *Math. Ann.*; 1924; 91: 169-209.

[McC89] J. McCool. A faithful polynomial representation of OutF3. *Math. Proc. Cambridge Philos. Soc.* 106 (1989), no. 2, 207-213.

[QR18] Yulan Qing and Kasra Rafi. Convexity of Balls in Outer Space

[Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551-565, 1983.

[Thu86] W.P. Thurston. Minimal stretch maps between hyperbolic surfaces. preprint, `arXiv:math.GT/9801039`, 1986.

[Vog15] K. Vogtmann. On the geometry of outer space. *Bull. Amer. Math. Soc. (N.S.)*, 52(1):27–46, 2015.

[Whi93] T. White. Fixed points of finite groups of free group automorphisms. *Proc. Amer. Math. Soc.*, 118(3):681–688, 1993.

Department of Mathematics, University of Toronto, Toronto, ON
E-mail address: yulan.qing@gmail.com

Department of Mathematics, University of Toronto, Toronto, ON
E-mail address: rafi@math.toronto.edu