Groups Geom. Dyn. 12 (2018), 571–613
DOI 10.4171/GGD/449

Groups, Geometry, and Dynamics

A dense geodesic ray in the Out(F_r)-quotient of reduced Outer Space

Yael Algom-Kfir[1] and Catherine Pfaff[2]

Abstract. In [16] Masur proved the existence of a dense geodesic in the moduli space for a surface. We prove an analogue theorem for reduced Outer Space endowed with the Lipschitz metric. We also prove two results possibly of independent interest: we show Brun's unordered algorithm weakly converges and from this prove that the set of Perron–Frobenius eigenvectors of positive integer $m \times m$ matrices is dense in the positive cone $\mathbf{R}_+^m$ (these matrices will in fact be the transition matrices of positive automorphisms). We give a proof in the appendix that not every point in the boundary of Outer Space is the limit of a flow line.

Mathematics Subject Classification (2010). 20E36, 20F65, 37E99.

Keywords. Outer automorphism group of the free group, Outer Space, geodesics.

Contents

[1] The research of Y. Algom-Kfir of was supported by THE ISRAEL SCIENCE FOUNDATION (grant no. 1941/14).

[2] C. Pfaff was supported first by the ARCHIMEDE Labex (ANR-11-LABX- 0033) and the A*MIDEX project (ANR-11-IDEX-0001-02) funded by the "Investissements d'Avenir," managed by the ANR. She was secondly supported by the CRC701 grant of the DFG, supporting the projects B1 and C13 in Bielefeld.

1. Introduction

1.1. Geodesics in Outer Space. One of the richest and most expansive methods for studying surfaces has been through the ergodic geodesic flow on Teichmüller space. As an example, it was used by Eskin and Mirzakhani [10] to count pseudo-Anosov conjugacy classes of a bounded length. For this reason, the papers of Masur [17] and Veech [21] independently proving the ergodicity of the Teichmüller flow were seminal in the field. The existence of an $\mathrm{Out}(F_r)$-invariant ergodic geodesic flow on Outer Space may similarly expand the tools for studying $\mathrm{Out}(F_r)$. Before giving the proof of the ergodicity theorem in Teichmüller space, Masur performed the following "litmus test" for its plausibility.

Theorem 1.1 ([16]). *Given a closed surface S_g of genus g, there exists a Teichmüller geodesic in the Teichmüller space T_g whose projection*

$$p\colon T_g \longrightarrow T_g / \mathrm{MCG}(S_g)$$

to moduli space is dense in both directions.

Our main theorem is an $\mathrm{Out}(F_r)$ analogue of the above theorem. Some of the terms in the theorem are defined directly below its statement.

Theorem A. *For each $r \geq 2$, there exists a geodesic fold ray in the reduced Outer Space $\mathcal{RX}_r$ whose projection to $\mathcal{RX}_r / \mathrm{Out}(F_r)$ is dense.*

Remark 1.2. (1) The reduced Outer Space $\mathcal{RX}_r$ is a subcomplex of the Outer Space $\mathcal{X}_r$, which consists of those graphs without separating edges (see Definition 2.12). It is an equivariant deformation retract of $\mathcal{X}_r$.

(2) The metric on $\mathcal{RX}_r$ with respect to which the ray in Theorem A is a geodesic is the Lipschitz metric (see Definition 2.25). It is an asymmetric metric (analogous to the Thurston metric on Teichmüller space [20]) that has proven to be very useful in the $\mathrm{Out}(F_r)$ context, e.g. [3].

(3) Because of the asymmetry of the metric, our geodesics will always be "directed geodesics," i.e. maps $\Gamma\colon [0,\infty) \to \mathcal{X}_r$ such that $d(\Gamma(t), \Gamma(t')) = t' - t$ for $t' > t \geq 0$, but not necessarily for $t > t'$.

(4) A fold line is a special kind of geodesic in $\mathcal{X}_r$ (explicitly described in Definition 3.14) that is analogous to a Teichmüller geodesic.

(5) Comparing between Theorem 1.1 and Theorem A, one may notice that our theorem declares the existence of a ray in contrast with Masur's theorem which declares the existence of a geodesic. We could easily extend our ray to a bi-infinite geodesic. However, the density of the image of the ray will follow from techniques that we cannot extend to the backwards direction.

Returning to Remark 1.2(1), we note that for proving algebraic properties of Out(F_r), one may usually replace Outer Space with $\mathcal{RX}_r$. However, on the geometric side, it is not known whether or not $\mathcal{RX}_r$ is convex in any coarse sense. That is, if x and y are points in is there a geodesic between them that stays in? If not, is there a geodesic that stays a distance R from?

We hence pose two questions:

Question 1.3. *For each $r \geq 2$, does there exist a geodesic fold line in $\mathcal{X}_r$ that is dense in both directions in $\mathcal{X}_r/$* Out(F_r)?

Question 1.4. *For each $r \geq 2$, is the reduced Outer Space $\mathcal{RX}_r$ coarsely convex? For example, given points $x, y \in \mathcal{RX}_r$ does there always exist a geodesic from x to y which is contained in reduced Outer Space?*

1.2. The unit tangent bundle. While Masur proved the existence of a dense geodesic first, this result also follows from his proof of the ergodicity of the Teichmüller flow. The unit tangent bundle of Teichmüller space is isomorphic to its unit cotangent bundle Q_0, which may be described explicitly as the space of unit area quadratic differentials on a closed surface S_g of genus g. The geodesic flow on Teichmüller space is a MCG(S_g) invariant action of $\mathbf{R}$ on Q_0. For $t \in \mathbf{R}$ we denote the flow by $T_t\colon Q_0 \to Q_0$. Given a quadratic differential $q \in Q_0$, the set of points $\{q_t\}_{t\in\mathbf{R}} = \{T_t q\}_{t\in\mathbf{R}}$ defines a geodesic in the Teichmüller space T_g.

Theorem 1.5 ([16]). *For a closed surface S_g of genus $g > 1$, there exists a quadratic differential $q \in Q_0$ so that the projection of $\{q_t\}$ for either $t > 0$ or $t < 0$ is dense in $Q_0/$* MCG(S_g).

The analogue of Theorem 1.5 is not obvious, as it is unclear what the unit tangent bundle should be. For example, two geodesics in Outer Space can meet for a period of time and then diverge from each other (or even alternately meet and diverge). This is an impediment to a "local" description of the tangent bundle. A more global approach would be to relate the tangent space at a point to the visual boundary of Outer Space. However, as of yet there is no description of the visual boundary of Outer Space. For example, we show in §9 that there are points on $\overline{\mathcal{X}_r} - \mathcal{X}_r$ (where $\overline{\mathcal{X}_r}$ denotes the set of very small F_r-trees) that are not ends of geodesic fold rays.

For the purposes of this paper, we propose the following analogue of the unit tangent bundle. Given a point $x \in \mathcal{RX}_r$ there are finitely many germs $[\alpha]$ of fold lines α in $\mathcal{RX}_r$ initiating at x. Define

$$\mathcal{URX}_r = \{(x, [\alpha]) \mid x \in \mathcal{RX}_r, \alpha \text{ is a fold line with } \alpha(0) = x \text{ and } \mathrm{Im}(\alpha) \subset \mathcal{RX}_r\}.$$

Given a geodesic ray $\gamma\colon [0, \infty) \to \mathcal{RX}_r$, for $t \in \mathbf{R}$ we denote by γ_t the path $\gamma_t(s) = \gamma(t + s)$. We may associate to γ a path in the unit tangent bundle $\tilde{\gamma}(t) = (\gamma(t), [\gamma_t])$. We prove:

Theorem B. *For each $r \geq 2$, there exists a Lipschitz geodesic fold ray $\tilde{\gamma}\colon [0,\infty) \to \mathcal{RX}_r$ so that the projection of $\tilde{\gamma}$ to $\mathcal{URX}_r / \operatorname{Out}(F_r)$ is dense.*

1.3. Geodesics in other subcomplexes. For each $r \geq 2$, we define the *theta subcomplex* $\mathcal{T}_r$ to be the subspace of $\mathcal{RX}_r$ consisting of all points in $\mathcal{X}_r$ whose underlying graph is either a rose or a theta graph, see Figure 1. This subcomplex carries the significance of being the minimal connected subcomplex containing the image of the Cayley graph under the natural map. Both as a warm-up, and for its intrinsic significance, we initially prove Theorem A in $\mathcal{T}_r$.

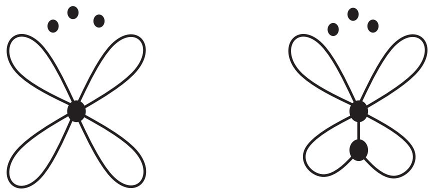

Figure 1. The underlying graphs of the simplices of $\mathcal{T}_r$. The graph on the left is called a *rose* and that on the right is called a *theta graph*.

Theorem C. *For each $r \geq 2$, there exists a dense fold ray in $\mathcal{T}_r$ that projects to a dense Lipschitz geodesic fold ray in $\mathcal{T}_r / \operatorname{Out}(F_r)$.*

1.4. Outline. We begin by outlining the proof of Theorem C. After proving Theorem C, we develop the topological machinery necessary to extend the more combinatorial arguments of the proof of Theorem C to the proof of Theorem B (and Theorem A as a corollary).

Recall that points in Outer Space are marked metric graphs equivalent up to homotopy, see Definition 2.2. As described in Definition 2.11, $\operatorname{Out}(F_r)$ acts on the right by changing the marking. To a point $x \in \mathcal{T}_r / \operatorname{Out}(F_r)$, one may associate a positive vector, the "length vector" recording the graph's edge lengths. The folding operation may be translated to a matrix recording the change in edge lengths from its initial point to its terminal point. In this dictionary, a fold ray in $\mathcal{T}_r / \operatorname{Out}(F_r)$ should correspond to an initial vector and a sequence of fold matrices. However, not every such sequence comes from a fold ray: a particular fold may or may not be allowed for a specific vector depending on whether its image is again a positive vector. Our challenge is to construct a sequence of fold matrices $\{T_k\}_{k=1}^{\infty}$ satisfying that, for some positive vector w_0, if we write $w_i = T_i \cdots T_1(w_0)$ for each $i \in \mathbf{N}$ then

I. the fold T_{i+1} is allowed in $w_i = T_i \cdots T_1(w_0)$ for each integer $1 \leq i < \infty$,

II. the set of vectors w_i is projectively dense in a simplex and,

III. the corresponding fold ray is a geodesic ray in the Lipschitz metric.

In order to prove Item II of the list, we prove the following fact:

Theorem D. *For each $r \geq 2$, let $S^r_{l_1}$ be the set of unit vectors according to the l_1 metric in $\mathbf{R}^r_+$. The set of Perron–Frobenius eigenvectors of matrices arising as the transition matrices of positive automorphisms in Aut(F_r) is dense in $S^r_{l_1}$.*

The proof of Theorem D, in §4, uses Brun's algorithm [7]. We also prove in §4 that Brun's (unordered) algorithm converges in angle, a result to our knowledge previously absent from the literature in dimensions higher than four. (Brun proved it in [7] for dimensions three and four.)

To construct the sequence $\{T_k\}$ of fold matrices we enumerate the powers of "Brun matrices" (see §4): $P_1, P_2, P_3, \ldots$. To each P_i, we can attach the following data:

- a positive Perron–Frobenius eigenvector v_i,
- a positive automorphism $g_i \in \mathrm{Aut}(F_r)$, also denoted g_{v_i}, so that the transition matrix of g_i is P_i,
- and a decomposition of P_i into fold matrices, arising from the decomposition of g_i into Nielsen generators which correspond to moves in Brun's algorithm.

We remark that this method is reminiscent of Masur's paper, where he proved the existence of a dense geodesic in $R_g = Q_0/\,\mathrm{MCG}(S_g)$ using the fact that closed loops in R_g are dense. The resemblance stems from the decomposition in the third item defining a loop in $\mathcal{RX}_r/\,\mathrm{Out}(F_r)$ based at v_i and the density, established by Theorem D, of the set of Perron–Frobenius eigenvectors $\{v_i\}$.

We concatenate the fold sequences associated to the matrices P_i together to form the sequence $\{T_k\}_{k=1}^{\infty}$. We address Item I on the list, i.e. the allowability of the sequence, in Lemma 5.1. We now have a fold ray through the points $\{w_j\}_{j=1}^{\infty}$.

For density of the geodesic ray we use the automorphisms g_i related to the P_i. By ensuring that arbitrarily high powers of these automorphisms (hence matrices) occur in the sequence, we ensure that the ray passes through points with length vectors arbitrarily close to the dense set of eigenvectors.

Finally, property (III) on the list, that the fold line is a Lipschitz geodesic, follows from the fact that every g_i is a positive automorphism (see Corollary 3.19).

To extend our proof of Theorem C to Theorems A and B, we prove that for a generic point y in reduced Outer Space, there exist roses x, z and a "positive" fold line $[x, z]$ remaining in reduced Outer Space and so that $y \in [x, z]$. Here by "positive" we mean that the change of marking from x to z is a positive automorphism. Moreover, if G is the underlying graph of y and E, E' are two adjacent edges in G, then one may choose $[x, z]$ so that it contains the fold of the turn $\{E, E'\}$ immediately after the point y. This construction is carried out in §6 and elevates the geodesic's density in $\mathcal{RX}_r/\,\mathrm{Out}(F_r)$ to density in $\mathcal{URX}_r/\,\mathrm{Out}(F_r)$. Additionally, the geodesic $[x, z]$ varies continuously as a function of y, as we prove in §7. Thus, one may adjust the previous argument to prove Theorems A and B, which we do in §8.

Acknowledgements. We wish to thank Pierre Arnoux for bringing our awareness to Brun's algorithm and its possible use in proving that the Perron–Frobenius eigenvectors are dense, as well as other helpful discussions. On this subject we are also heavily indebted to Jon Chaika who pointed out to us that Brun's algorithm is ergodic and that we could use this to elevate our proof of the Perron–Frobenius eigenvectors being dense in a simplex to additionally prove that Brun's algorithm converges in angle (also somewhat simplifying our proof that the Perron–Frobenius eigenvectors are dense in a simplex). We would further like to thank Jayadev Athreya for posing the question and helpful discussions. Inspirational ideas given by Pascal Hubert were also particularly valuable. We are indebted to Lee Mosher for pointing out that Keane has a paper on a complication we were facing with regard to Brun's algorithm. And we are indebted to Fritz Schweiger for his generosity in helping us understand arguments of his books. Finally, we are indebted to Valérie Berthé, Mladen Bestvina, Kai-Uwe Bux, Albert Fisher, Vincent Guirardel, Ilya Kapovich, Amos Nevo, and Stefan Witzel for helpful discussions.

2. Definitions and background

2.1. Outer Space and the action of Out(F_r). Culler and Vogtmann introduced Outer Space in [8]. Points of Outer Space are "marked metric graphs."

Definition 2.1 (graph, positive edges). A *graph* will mean a connected 1-dimensional cell complex. $V(G)$ will denote the vertex set and $E(G)$ the set of unoriented edges. The degree of a vertex $v \in V(G)$ will be denoted $\deg_G(v)$, or $\deg(v)$ when G is clear.

For each edge $e \in E(G)$, one may choose an orientation. Once the orientation is fixed, that oriented edge E will be called *positive* and the edge with the reverse orientation $\bar{E}$ will be called *negative*. Given an oriented edge E, $i(E)$ will denote its initial vertex and $\mathrm{ter}(E)$ its terminal vertex. A *directed* graph is a graph G with a choice of orientation on each edge $e \in E(G)$, we call this choice an *orientation* on G.

Given a free group F_r of rank $r \geq 2$, we choose once and for all a free basis $A = \{X_1, \dots, X_r\}$. Let $R_r = \bigvee_{i=1}^{r} \mathbf{S}^1$ denote the graph with one vertex and r edges, we call this graph an *r-petaled rose*. We choose once and for all an orientation on R_r and identify each positive edge of R_r with an element of the chosen free basis. Thus, a cyclically reduced word in the basis corresponds to an immersed loop in R_r.

Definition 2.2 (marked metric F_r-graph). For each integer $r \geq 2$ we define a *marked F_r-graph* to be a pair $x = (G, m)$ where

- G is a graph such that $\deg(v) \geq 3$ for each vertex $v \in V(G)$;
- $m\colon R_r \to G$ is a homotopy equivalence, called a *marking*.

A *marked metric graph* is a triple (G, m, ℓ) so that (G, m) is a marked graph and

- the map $\ell\colon E(G) \to \mathbf{R}_+$ is an assignment of lengths to the edges. We require that $\sum_{e\in E(G)} \ell(e) = 1$. The quantity $\mathrm{vol}(G) = \sum_{e\in E(G)} \ell(e)$ is called the *volume* of G.

Remark 2.3 (a metric graph as a metric space). The assignment of lengths to the edges does not quite determine a metric on G, but a homeomorphism class of metrics. This choice is inconsequential in this paper.

Define an equivalence relation on marked metric F_r-graphs by $(G, m, \ell) \sim (G', m', \ell')$ when there exists an isometry $\varphi\colon (G, \ell) \to (G', \ell')$ so that m' is homotopic to $\varphi \circ m$.

Definition 2.4 (underlying set of Outer Space). For each $r \geq 2$, as a set, the (*rank-r*) *Outer Space* $\mathcal{X}_r$ is the set of equivalence classes of marked metric F_r-graphs.

Remark 2.5. On occasion we may think of graphs with valence-2 vertices as living in Outer Space by considering them equivalent to the graphs obtained by unsubdividing at their valence-2 vertices.

Definition 2.6. The *simplex* σ in $\mathcal{X}_r$ corresponding to the marked graph (G, m) is

$$\sigma_{(G,m)} = \{(G, m, \ell) \in \mathcal{X}_r \mid \mathrm{vol}(G) = 1\}.$$

By enumerating the edges of G, we can identify $\sigma_{(G,m)}$ with the open simplex

$$\mathcal{S}_{|E|} = \left\{\vec{v} \in \mathbf{R}_+^{|E|} \;\middle|\; \sum_{i=1}^{|E|} v_i = 1\right\}.$$

Here $E = E(G)$. We denote this identification by $n\colon \sigma_{(G,m)} \to \mathcal{S}_{|E|}$. We call the open simplex corresponding to (R_r, id) the *base simplex* and denote it by σ_0.

Outer Space has the structure of an ideal simplicial complex built from open simplices (see [22]), faces of $\sigma_{(G,m)}$ arise by letting the edges of a tree in G have length 0.

Definition 2.7 (simplicial metric). Given an open simplex $\sigma_{(G,m)}$ in $\mathcal{X}_r$, the *simplicial metric* on $\sigma_{(G,m)}$ is defined by $d_s(\ell, \ell') = \sqrt{\sum_{e\in E(G)} (\ell(e) - \ell'(e))^2}$, for $\ell, \ell' \in \sigma_{(G,m)}$. We also denote by d_s the extension of this metric to a path metric on $\mathcal{X}_r$.

Remark 2.8. In §2.4 we define the Lipschitz metric on $\mathcal{X}_r$. The simplicial metric and Lipschitz metric on Outer Space differ in important ways. However, open balls with respect to the Lipschitz metric (in either direction, see Remark 2.26) contain open balls of the simplicial metric. Therefore, a set dense with respect to the simplicial topology will also be dense with respect to the Lipschitz topology.

Definition 2.9 (unprojectivized Outer Space). [8] The (*rank-r*) *unprojectivized Outer Space* $\widehat{\mathcal{X}}_r$ is the space of metric marked F_r-graphs where $\mathrm{vol}(G)$ is not necessarily 1.

There is a map from $\widehat{\mathcal{X}}_r$ to $\mathcal{X}_r$ normalizing the graph volume, i.e.

$$\begin{aligned} q\colon \mathbf{R}^m_+ &\longrightarrow \mathcal{S}_m, \\ (x_1, \dots, x_m) &\longmapsto \Big(\frac{x_1}{\sum_{i=1}^m x_i}, \dots, \frac{x_m}{\sum_{i=1}^m x_i}\Big), \end{aligned} \tag{1}$$

and taking the point (G, μ, ℓ) to the point $(G, \mu, q(\ell))$.

We call the full preimage under q of a simplex in $\mathcal{X}_r$ an *unprojectivized simplex*.

Definition 2.10 (topological Outer Space). [8] The topological space consisting of the set of equivalence classes of marked metric F_r-graphs, endowed with the simplicial topology, is called the (*rank-r*) *Outer Space* and is also denoted by $\mathcal{X}_r$.

Definition 2.11 ($\mathrm{Out}(F_r)$ action). If $\Phi \in \mathrm{Aut}(F_r)$ is an automorphism, let $f_\Phi\colon R_r \to R_r$ be a homotopy equivalence corresponding to Φ via the identification of $E(R_r)$ with the chosen free basis A of F_r. We define a *right action of* $\mathrm{Out}(F_r)$ on $\mathcal{X}_r$. An outer automorphism $[\Phi] \in \mathrm{Out}(F_r)$ acts by $[G, m, \ell] \cdot [\Phi] = [G, m \circ f_\Phi, \ell]$.

Definition 2.12 (reduced Outer Space $\mathcal{RX}_r$ and $\mathcal{M}_r$). For each integer $r \geq 2$, the (*rank-r*) *reduced Outer Space* $\mathcal{RX}_r$ is the subcomplex of $\mathcal{X}_r$ consisting of precisely those simplices $\sigma_{(G,m)}$ such that G contains no separating edges. This space is connected and an $\mathrm{Out}(F_r)$-equivariant deformation retract of $\mathcal{X}_r$.

Let $\mathcal{M}_r$ denote the quotient space of $\mathcal{RX}_r$ by the $\mathrm{Out}(F_r)$ action. Hence, $\mathcal{M}_r$ contains a quotient of a simplex for each graph (no longer marked). Note that, as a result of graph symmetries, simplices in $\mathcal{X}_r$ do not necessarily project to simplices in $\mathcal{M}_r$. Thus, $\mathcal{M}_r$ is no longer a simplicial complex but a union of cells which are quotients of simplices in $\mathcal{X}_r$.

2.2. Train track structures. Much of the following definitions and theory can be found in [5] or [4], for example. However, it should be noted that some of our definitions, including that of an illegal turn, are somewhat nonstandard.

Definition 2.13 (regular maps). We call a continuous map $f\colon G \to H$ of graphs *regular* if for each edge $e \in E(G)$, we have that $f|_{int(e)}$ is locally injective and that f maps vertices to vertices.

Definition 2.14 (paths and loops). Depending on the context an *edge-path* in a graph G will either mean a continuous map $[0,n] \to G$ that, for each $1 \leq i \leq n$, maps $(i-1, i)$ homeomorphically to the interior of an edge, or if the graph G is directed, a sequence of oriented edges $e_1, \ldots, e_n$ such that $\mathrm{ter}(e_i) = i(e_{i+1})$ for each $1 \leq i \leq n-1$. We may on occasion also allow for e_1 and e_n to be partial edges. Given any path $\gamma = e_1 \cdots e_n$, we will denote its initial vertex, i.e. $i(e_1)$, by $i(\gamma)$ and its terminal vertex, i.e. $\mathrm{ter}(e_n)$, by $\mathrm{ter}(\gamma)$.

A *loop* α in G is the image of an immersion $\alpha: \mathbf{S}^1 \to G$. We will associate to each loop an edge-path unique up to cyclic ordering.

In a directed graph G, we will call a path *directed* that either crosses all edges in a positive direction (a *positive path*) or crosses all edges in a negative direction (a *negative path*). The operation of path concatenation will be denoted by $*$.

Definition 2.15 (illegal turns and gates). Let $f: G \to H$ be a regular map. Let $e, e' \in E(G)$ be oriented edges with the same initial vertex. We call $\{e, e'\}$ a *turn*. We say a turn $\{e, e'\}$ is an *illegal turn* for f if the first edge of the edge-path $f(e)$ equals the first edge of the edge-path $f(e')$. The property of forming an illegal turn is an equivalence relation and the equivalence classes are called *gates*.

Definition 2.16 (train track structures). Let $f: G \to H$ be a regular map. If every vertex of G has ≥ 2 gates, then we call the partition of the turns of G into gates a *train track structure* and say that f induces a *train track structure* on G. An immersed path $\alpha: I \to G$ will be considered *legal* with respect to a given train track structure if it does not contain a subpath $e_i e_j$ where $\{\overline{e_i}, e_j\}$ is an illegal turn.

Remark 2.17. The image of a legal path is locally embedded.

Definition 2.18 (transition matrix). The *transition matrix* of a regular self-map $f: G \to G$ is the square $|E(G)| \times |E(G)|$ matrix (a_{ij}) such that a_{ij}, for each i and j, is the number of times $g(e_i)$ passes over e_j in either direction.[1]

We define the *transition matrix* for an element $\Phi \in \mathrm{Aut}(F_r)$ to be the transition matrix of f_Φ (see Definition 2.11).

2.3. Perron–Frobenius theory. We are particularly interested in positive matrices (defined below) because of their known properties (due to Perron–Frobenius theory) of contracting the *positive cone* $\mathbf{R}^d_+ = \{v \in \mathbf{R}^d \mid v_i > 0,\ i = 1, \ldots d\}$.

Definition 2.19 (positive matrices, Perron–Frobenius eigenvalues and eigenvectors). We call a matrix $A = [a_{ij}]$ *positive* if each entry of A is strictly positive. By Perron–Frobenius theory, we know that each such matrix has a unique eigenvalue

[1] This matrix is the transpose of the transition matrix as Bestvina and Handel define it in [5], but this definition will have a stronger relationship with the change-of-metric matrix we define later.

of maximal modulus and that this eigenvalue is real. This eigenvalue is called the *Perron–Frobenius* (*PF*) *eigenvalue* of A. It has an associated eigenvector whose entries are each strictly positive. We call the eigenvector with strictly positive entries and such that all entries sum to one the *Perron–Frobenius* (*PF*) *eigenvector*.

Definition 2.20 (weak convergence). A sequence $\{A_k\}_{k=1}^{\infty}$ of $d \times d$ matrices, restricted to vectors in $\mathbf{R}_+^d$, *converges weakly* if the sequence $\{A_k(\mathbf{R}_+^d)\}_{k=1}^{\infty}$ converges projectively to a point.

Remark 2.21. Perron–Frobenius theory also tells us that, for a positive matrix M, the sequence $\{M^k\}_{i=1}^{\infty}$ weakly converges to the line spanned by the PF eigenvector.

2.4. The Lipschitz metric

Definition 2.22 (difference in markings). Let $x = (G, m, \ell)$ and $y = (G', m', \ell')$ be two points in $\mathcal{X}_r$. Denote by $h\colon G \to R_r$ a homotopy inverse of m. *A difference in markings* is a linear map $f\colon G \to G'$ homotopic to $m' \circ h$.

Definition 2.23 (stretch). Let α be a conjugacy class in F_r, equipped with a free basis A. By abuse of notation we may think of α as an immersed loop $\alpha\colon \mathbf{S}^1 \to R_r$ in R_r via the identification of the edges of R_r with A. For $x = (G, m, \ell) \in \mathcal{X}_r$, let α_x denote the unique immersed simplicial loop in G homotopic to $m(\alpha)$.

Given a conjugacy class α in F_r and $x \in \mathcal{X}_r$, we define $l(\alpha, x)$ as the length of α_x. (Notice that since α_x is a simplicial loop in x, its length does not depend on the particular metric structure chosen for x, see Remark 2.3). For $x, y \in \mathcal{X}_r$ define the *stretch* of α from x to y as $\mathrm{st}_\alpha(x, y) := \frac{l(\alpha,y)}{l(\alpha,x)}$.

The following theorem is attributed to either White or Thurston. It can be found in [1] or [12].

Theorem 2.24. *Given a continuous map f of metric spaces, let $Lip(f)$ denote the Lipschitz constant for f. Then for each pair of points $x, y \in \mathcal{X}_r$, we have*

$$\inf\{\mathrm{Lip}(f) \mid f\colon x \to y \text{ is a difference in marking }\} = \sup\{\mathrm{st}_\alpha(x, y) \mid \alpha \in F_r\}. \tag{2}$$

Moreover, both the infimum and supremum are realized.

Definition 2.25 (Lipschitz metric). The *Lipschitz metric* $d(x, y)$ is defined as the log of either of the quantities in Equation 2. This function is not symmetric but satisfies the other axioms of a metric [4].

A difference in marking that achieves the minimum Lipschitz constant of (2) is called an *optimal map*. A loop that achieves the maximum stretch is called a *witness*. For each $x, y \in \mathcal{X}_r$ there exist optimal maps and witnesses.

Remark 2.26. An open ball based at x with radius R is either

$$B_{\rightarrow}(x,R) = \{y \in \mathcal{X}_r \mid d(x,y) < R\} \text{ or } B_{\leftarrow}(x,R) = \{y \in \mathcal{X}_r \mid d(y,x) < R\}.$$

In either case, the simplicial topology is equal to the Lipschitz topology.

For a given difference of marking f, the subgraph of G where the Lipschitz constant is achieved is called the *tension graph*, usually denoted Δ_f. Notice that f induces a train track structure on Δ_f. Proposition 2.27 gives one way to identify witnesses.

Proposition 2.27 ([3]). *Let x and y be two points in $\mathcal{X}_r$, let $f\colon x \to y$ be a map, and let Δ_f be the tension graph of f. If Δ_f contains a legal loop, then f is an optimal map and any legal loop in Δ_f is a witness. Conversely, if $\alpha \subset x$ is a witness, then it is a legal loop in Δ_f.*

Proposition 2.28. *Let $f\colon x \to y$ and $g\colon y \to z$ be difference in markings maps. Let α be a conjugacy class in F_r satisfying that α_x is f-legal and contained in Δ_f and that α_y is g-legal and contained in Δ_g. Then $d(x,z) = d(x,y) + d(y,z)$.*

Proof. Since α_x is contained in Δ_f, the map f stretches each edge of α_x by $\lambda_f = \mathrm{Lip}(f)$. Moreover, since α_x is f-legal, $l(\alpha,y) = \lambda_f \cdot l(\alpha,x)$. Similarly, if $\lambda_g = \mathrm{Lip}(g)$, then $l(\alpha,z) = \lambda_g \cdot l(\alpha,y) = \lambda_f\lambda_g \cdot l(\alpha,x)$. Thus, $\mathrm{st}_\alpha(x,z) = \lambda_f\lambda_g$ and therefore $d(x,z) \geq \log(\lambda_f\lambda_g)$. By Proposition 2.27, α_x and α_y are both witnesses, hence $d(x,y) = \log(\mathrm{st}_\alpha(x,y)) = \log(\lambda_f)$ and $d(y,z) = \log(\mathrm{st}_\alpha(y,z)) = \log(\lambda_g)$. Thus, we have $d(x,z) \geq \log(\lambda_f) + \log(\lambda_g) = d(x,y) + d(y,z)$. The triangle inequality gives us an equality. □

3. Fold paths and geodesics

In this section we introduce fold lines and prove results that will allow us to construct Lipschitz geodesics from certain infinite sequences of nonnegative matrices (unfolding matrices).

Definition 3.1 (unparametrized geodesic). Let $I \subset \mathbf{R}$ be a generalized interval in $\mathbf{R}$. An *unparametrized geodesic* in $\mathcal{X}_r$ is a map $\Gamma\colon I \to \mathcal{X}_r$ satisfying that

(1) for each $s < r < t$, we have $d(\Gamma(r),\Gamma(t)) = d(\Gamma(r),\Gamma(s)) + d(\Gamma(s),\Gamma(t))$ and

(2) there exists no nontrivial subinterval $I' \subset I$ and point $x_0 \in \mathcal{X}_r$ such that $\Gamma(t) = x_0$ for each $t \in I'$.

Remark 3.2. If Γ is an unparametrized geodesic then there exists a generalized interval I' and a homeomorphism $h\colon I' \to I$ so that $\Gamma \circ h$ is an honest directed geodesic, i.e. for all $s < t$, we have that $d(\Gamma \circ h(s), \Gamma \circ h(t)) = t - s$.

Again $A = \{X_1, \dots, X_r\}$ will denote a fixed free basis of F_r.

Definition 3.3 (fold automorphism). By a *fold automorphism* we will mean a "left Nielsen generator," i.e. an automorphism of the following form ($i \neq j$):

$$\Phi_{ij}(X_k) = \begin{cases} X_j X_k & \text{for } k = i, \\ X_k & \text{for } k \neq i. \end{cases} \tag{3}$$

To a fold automorphism one can associate a matrix.

Definition 3.4 ((un)folding matrix). Let $i \neq j \in \{1, \dots, r\}$. Then the (i, j) *folding matrix* T_{ij} has entries t_{kl}, where

$$t_{kl} = \begin{cases} 1 & \text{if } k = l, \\ -1 & \text{if } (k,l) = (i,j), \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Notice that T_{ij} is *not* the transition matrix of Φ_{ij}, though it will relate to the change-of-metric matrix coming from a folding operation. The matrix T_{ij} is invertible and we call $M_{ij} := T_{ij}^{-1}$ the (i, j) *unfolding matrix*. The entries of M_{ij} are m_{kl}, where

$$m_{kl} = \begin{cases} 1 & \text{if } k = l \text{ or } (k,l) = (i,j), \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Notice that the nonnegative *matrix* M_{ij} is the transition matrix of Φ_{ij}. We hence sometimes write $M(\Phi_{ij})$ for this matrix.

Definition 3.5 (combinatorial fold). Let G be a graph whose oriented edges are numbered. Let (e_k, e_j) be a pair of distinct oriented edges with $i(e_k) = i(e_j)$. A combinatorial fold is a tuple $(G, (e_k, e_j), G', f)$ where G' is a graph and $f\colon G \to G'$ is a quotient map that identifies an initial segment of e_k with an initial segment of e_j.

(1) When f identifies part of e_k with all of e_j we will call it a *proper full fold*. We will call this tuple "folding e_k over e_j."

(2) When f identifies all of e_k with all of e_j we will call it a *full fold* and say that "e_k and e_j are fully folded."

(3) When f identifies a proper segments of e_k and e_j we will call it a *partial fold* and say that "e_k and e_j are partially folded."

Notation We sometimes suppress some of the data depending on the context. We will denote the combinatorial fold by f or $G \rightharpoonup G'$ or $(G,(e_i,e_j))$ depending on which data we want to emphasize.

Definition 3.6 (combinatorial folds of fold automorphisms and the induced enumeration). Let G be the r-rose with an enumeration on its oriented edges, then the fold automorphism Φ_{ij} induces a combinatorial fold, which by abuse of notation we also denote Φ_{ij}. The target graph G' is also a rose and Φ_{ij} induces an enumeration of the edges of G' by declaring the edge $\Phi_{ij}(e_k)$ to be the k^{th} edge for each $k \neq i$ and calling the remaining edge of G' the i^{th} edge.

Definition 3.7 (direction matching folds). Let G be an oriented graph. A combinatorial fold $(G,(e_i,e_j))$ is *direction matching* if e_1 and e_2 are either both negative edges (we then call f *negative*) or both positive edges (we then call f *positive*).

Observation 3.8. Let G be an oriented graph and $f\colon G \rightharpoonup H$ a direction matching combinatorial fold, then f induces an orientation on the edges of H. Moreover, f maps each positive edge of G to a positive edge-path in H (of simplicial length ≤ 2).

Definition 3.9 (allowable folds). Let $x_0 = (G,m,\ell)$ be a point in Outer Space. The fold $(G,(e,e'))$ is *allowable* in x_0 if the following two conditions hold:

(1) $\ell(e) \geq \ell(e')$;

(2) if $\mathrm{ter}(e) = \mathrm{ter}(e')$, then $\ell(e) > \ell(e')$. In this case this is a *proper full fold*.

Let G be a rose with its edges enumerated. This enumeration induces a homeomorphism $n_\tau\colon \sigma_{(G,m)} \to \mathcal{S}_{|E|}$. We write $\sigma^{(i,j)}_{(G,m)}$ for the set where Φ_{ij} is allowable.

Lemma 3.10. *Let $\hat\sigma_0$ be the unprojectivized base simplex and let Φ be a fold automorphism. Assume the edges of R_r have been enumerated, let $n_0\colon \hat\sigma_0 \to \mathbf{R}^r_+$ be the induced identification, and let n_1 be the identification induced from n_0 by Φ (See Definition 3.6). Then for each $x \in \hat\sigma_1$ we have $n_1(x) = n_0(x\cdot\Phi^{-1})$ (as defined in Definition 2.11).*

If a fold $(G,(e,e'))$ is allowable at a point x_0 in Outer Space, one can construct a path $\{\hat x_t\}_{t\in[0,1]}$ in unprojectivized outer space, called a "fold path." This is done by identifying initial segments of e and e' of length $t\ell_0(e')$, for $0 \leq t \leq 1$. The quotient map $f_{t,0}\colon \hat x_0 \to \hat x_t$ is a homotopy equivalence, as are the quotient maps $f_{t,s}$ for $0 \leq s \leq t \leq 1$. By projectivizing we get a path $\{x_t\}$ in Outer Space.

Definition 3.11 (basic fold paths). Given an allowable fold as above, the path $\mathcal{F}\colon [0,1] \to \mathcal{X}_r$ defined by $t \mapsto x_t$ is the *fold path in $\mathcal{X}_r$ starting at x_0 and defined by folding e over e'*. The path $t \to \hat{x}_t$ will be called the (*unprojectivized*) *basic fold path*. We will not always use the hat notation when discussing unprojectivized paths but will mention whether the image lies in $\mathcal{X}_r$ or $\widehat{\mathcal{X}}_r$, if it is otherwise unclear.

Definition 3.12 (change-of-metric matrix). Let G, G' be graphs and assume we have numbered their oriented edges. Let Ψ be a linear map from a subset of the unprojectivized simplex $\hat{\sigma}_{(G,\mu)}$ to the unprojectivized simplex $\hat{\sigma}_{(G',\mu')}$. Then Ψ may be represented by an $|E(G')| \times |E(G)|$ matrix. This matrix will be called the *change-of-metric* matrix.

Lemma 3.13. *Let Φ be an allowable fold automorphism on the point $x_0 = (R_r, m, \ell)$ and suppose the edges of R_r have been numbered so that $\Phi = \Phi_{ij}$. Let x_1 be the folded graph and suppose $\hat{\sigma}_0$ and $\hat{\sigma}_1$ are unprojectivized open simplices containing respectively x_0 and x_1. The change-of-metric matrix for the fold operation from $n_0(\hat{\sigma}^{(i,j)}_{(R,m)})$ to $n_1(\hat{\sigma}_{(R,\Phi\circ m)})$ is the matrix T_{ij} of Definition 3.4.*

Definition 3.14 (fold paths). A *fold path* $\mathcal{F}\colon [0,k] \to \mathcal{X}_r$ is a path in $\mathcal{X}_r$ that may be divided into a sequence of basic fold paths $\{\mathcal{F}_i\}_{i=1}^k$ as in Definition 3.11, so that $\mathcal{F}_i(1) = \mathcal{F}_{i+1}(0)$ for each i. As above, we will denote by $x_t = (G_t, m_t, \ell_t)$ the points of the path in $\mathcal{X}_r$. The maps $f_{t,s}$ for $s, t \in [0,k]$ will be defined similarly as above.

If $\mathcal{F}$ is a fold path from $\mathcal{F}(0) = x$ to $\mathcal{F}(k) = y$ we sometimes denote it by $\mathcal{F}\colon x \rightharpoonup y$.

The following Lemma follows from Lemma 3.13 and its proof is left to the reader.

Lemma 3.15. *Let $f_1, \dots, f_k$ be a sequence of combinatorial proper full folds, with associated change-of-metric matrices $T_1, \dots, T_k$ having respective inverse matrices $M_1, \dots, M_k$. Suppose that $v \in M_1 \cdots M_k(\mathbf{R}^r_+)$. Then the combinatorial fold f_l is allowable in the metric graph $n^{-1}_{l-l}(T_{l-l} \cdots T_1(v))$, for each $2 \le l \le k$ (and any marking). Furthermore, applying $f_1, \dots, f_k$ to $x_0 \in \sigma_0$ will result in the point $n_k^{-1}(T_k \cdots T_1(n_0(x_0)))$ of the simplex $\sigma_{(R_r, f_k \circ \cdots \circ f_1 \circ m)}$.*

Lemma 3.16. *Let $\{D_i\}_{i=1}^\infty$ denote a sequence of nonnegative invertible matrices such that, for each $i \in \mathbf{N}$, there exist some integer $n > i$ such that $D_i D_{i+1} \cdots D_{n-1} D_n$ is strictly positive. Then there exists a vector $w_0 \in \mathbf{R}^r_+$ so that, if we define $w_{l+1} := D^{-1}_{l+1} w_l$, for each l, then each w_l is a positive vector.*

Proof. Let $\overline{\mathbf{R}^r_+}$ denote the set of vectors with nonnegative entries, and recall the map q from Equation 1. Let i_1 be such that $D_1 \cdots D_{i_1}$ is strictly positive and let $M_1 = D_1 \cdots D_{i_1}$. Recursively define i_k so that $D_{i_{k-1}+1} \cdots D_{i_k}$ is strictly positive and let $M_k = D_{i_{k-1}+1} \cdots D_{i_k}$. Let $\mathcal{I} = \bigcap_{k=1}^{\infty} M_1 \cdots M_k(\mathbf{R}^r_+)$. Note that

$$\mathcal{I} \supset \bigcap_{k=1}^{\infty} M_1 \ldots M_k(\overline{\mathbf{R}^r_+}) \supset \bigcap_{k=1}^{\infty} q(M_1 \ldots M_k(\overline{\mathbf{R}^r_+})).$$

The right-most intersection is nonempty since it is an intersection of nested compact sets. Moreover, $\mathcal{I} \subset M_1(\overline{\mathbf{R}^r_+}) \subset \mathbf{R}^r_+$. Choose $w_0 \in \mathcal{I}$. Given $l \in \mathbf{N}$, let $k \in \mathbf{N}$ be such that $i_{k-1} < l \leq i_k$. Since $w_0 \in \mathcal{I} \subset M_1 \cdots M_k(\mathbf{R}^r_+)$, we have that $M_k^{-1} \cdots M_1^{-1} w_0$ is a strictly positive vector. Hence, since each D_i is a nonnegative invertible matrix, $w_l = D_{l+1} \cdots D_{i_k} M_k^{-1} \cdots M_1^{-1} w_0 \in D_{l+1} \cdots D_{i_k}(\mathbf{R}^r_+) \subset \mathbf{R}^r_+$. □

Corollary 3.17. *Let $\{f_i\}_{i=0}^{\infty}$ be a sequence of combinatorial proper full folds of the r-rose, with associated change-of-metric matrices $\{T_i\}_{i=0}^{\infty}$ having respective inverse matrices $\{D_i\}_{i=0}^{\infty}$. Suppose that for each $i \in \mathbf{N}$, there exist some integer $n > i$ such that $D_i D_{i+1} \cdots D_{n-1} D_n$ is strictly positive. Then there exists a vector $w_0 \in \mathbf{R}^r_+$ so that the infinite fold sequence $\{f_i\}_{i=0}^{\infty}$ is allowable in the rose $x_0 = n_0^{-1}(w_0)$.*

Proposition 3.18. *Let $\{\mathcal{F}_i : x_i \rightharpoonup x_{i+1}\}_{i=0}^{k}$ be a sequence of fold paths with fold maps $\{f_{s,t}\}_{s \geq t \geq 0}$. Suppose there is a conjugacy class α in F_r satisfying that, for each i, the realization α_{x_i} of α in x_i is legal with respect to the train track structure induced by $f_{i+1,i}$. Then the corresponding fold path $\mathrm{Im}(\mathcal{F}) = \{x_t\}_{t \in [0,k]}$ is an unparametrized geodesic, i.e. for each $r \leq s \leq t$ in $[0,k]$, we have $d(x_r, x_t) = d(x_r, x_s) + d(x_s, x_t)$.*

Proof. The proof uses Propositions 2.27 and 2.28 and is left to the reader. □

Suppose y_0 is a graph and $f_0 : y_0 \rightharpoonup y_1$ is direction matching, then by Observation 3.8, we have that y_1 inherits an orientation such that the image of each edge is a positive edge-path (see Definition 2.14).

Corollary 3.19. *Let y_0 be a directed metric r-rose graph with length vector v and let $\{f_i : y_i \rightharpoonup y_{i+1}\}_{i=0}^{i=k}$ be an allowable sequence of proper full folds. Suppose that for each $i = 0, \ldots, k-1$ the fold f_{i+1} is direction matching with respect to the orientation of y_{i+1} inherited by f_i. For each $0 \leq i \leq k$, let $\{\mathcal{F}_i : y_i \rightharpoonup y_{i+1}\}$ denote the fold path determined by $\{f_i\}$. Then the corresponding fold path $Im(\mathcal{F}) = \{y_t\}_{t \in [0,k]}$ is an unparametrized geodesic.*

Proof. Without generality loss we can assume that the marking on R_r is the identity. Then, by Proposition 3.18, it suffices to show that there exists a conjugacy class α in F_r satisfying that, for each i, the realization α_{y_i} of α in y_i is legal with respect to the train track structure induced by the map f_i. We claim that this holds for the conjugacy class α of a positive generator X_1. By induction suppose that α_{y_i} is a positive loop. Since f_i maps positive edges to positive edge-paths, $\alpha_{y_{i+1}}$ is also a positive loop. Since f_{i+1} is direction matching, $\alpha_{y_{i+1}}$ is f_{i+1}-legal and positive. □

4. Brun's algorithm and density of Perron–Frobenius eigenvectors

We introduce fibered systems so that we can use an algorithm of Brun to prove, in Theorem 4.14, that the sequence of Brun matrices converge weakly for a full measure set of points. These matrices are unfolding matrices, which will be significant for proving Theorems A, B, and C in the next sections.

4.1. Fibered systems. The following definitions are taken from [19].

Definition 4.1 (fibered systems). A pair (B, T) is called a *fibered system* if B is a set, $T: B \to B$ is a map, and there exists a partition $\{B(i) \mid i \in I\}$ of B such that I is countable and $T|_{B(i)}$ is injective. The sets $B(i)$ are called *time-1 cylinders*.

Definition 4.2 (time-s cylinder). For each $x \in B$, one can define a sequence $\Phi(x) = (i_1(x), i_2(x), \dots) \in I^{\mathbf{N}}$ by letting $i_s(x) = i \iff T^{s-1}x \in B(i)$. In other words, $i_s(x)$ tells us which cylinder $T^{s-1}x$ lands in. Then a *time-s cylinder* is a set of the form

$$B(i_1, \dots, i_s) = \{x \in B \mid i_1(x) = i_1, \dots, i_s(x) = i_s\}.$$

From the definitions we have $B(i_1, \dots, i_{s+1}) = B(i_1, \dots, i_s) \cap T^{-s}B(i_{s+1})$.

4.2. The unordered Brun algorithm in the positive cone. The following algorithm (commonly referred to as *Brun's algorithm*) was introduced by Brun [7] as an analogue in dimensions 3 and 4 of the continued fractions expansion of a real number. It was later extended to all dimensions by Schweiger in [18].

Definition 4.3 (Brun's (unordered) algorithm). *Brun's unordered algorithm* is the fibered system (C^n, T) defined on

$$C^n = \mathbf{R}^n_+ = \{(x_1, \dots, x_n) \in \mathbf{R}^n_+ \mid x_i \geq 0 \text{ for all } 1 \leq i \leq n\}$$

by

$$T: C^n \longrightarrow C^n,$$
$$(x_1, \dots, x_n) \longmapsto (x_1, \dots, x_{m-1}, x_m - x_s, x_{m+1}, \dots, x_n),$$

where

$$m(x) := \min\{i \mid x_i = \max_{1\leq j\leq n} x_j\},$$
$$s(x) := \min\{i \neq m(x) \mid x_i = \max_{1\leq j\neq m(x)\leq n} x_j\}.$$

In other words, for each $x = (x_1, \dots, x_n) \in C^n$, we have that $m(x)$ is the first index that achieves the maximum of the coordinates and $s(x)$ is the first index that achieves the maximum of all of the coordinates except for $m(x)$.

Notice that, letting $(i, j) = (m(x), s(x))$, the transformation T is just left multiplication by the matrix $T_{i,j}$ from Definition 3.4. Then, given a vector $v_0 = (x_1, \dots, x_n) \in C^n$ with rationally independent coordinates, one obtains an infinite sequence

$$\{v_k = (x_1^k, \dots, x_n^k)\}_{k=1}^{\infty} \subset C^n,$$

where v_{k+1} is recursively defined by $v_{k+1} = T_{m(v_k),s(v_k)} v_k$.

Definition 4.4 (Brun sequence). In light of the above, the sequence for Brun's unordered algorithm (which we call the *Brun sequence*) will consist of the ordered pairs $\vec{k}_s = (i_s, j_s)$, where $(i_s, j_s) = (m(v_s), s(v_s))$. Further, the sequence $\Theta(x) = \{\vec{k}_s\}_{s=1}^{\infty}$ will determine a sequence of folding matrices, which we denote by $\{T_s^x\}_{s=1}^{\infty}$, where $T_s^x = T_{\vec{k}_s(x)}$. We let $\{M_s^x\}_{s=1}^{\infty}$ denote the corresponding inverses, i.e. the unfolding matrices. We denote their finite products by

$$A_s^x = M_1^x M_2^x \cdots M_s^x \text{ for each } s \in \mathbf{N}. \tag{6}$$

Then, as above, for each $v_0 \in C^n$, we have a sequence of vectors $\{v_s\}_{s=1}^{\infty} \subset C^n$, where $v_{s+1} = T v_s = T_{s+1}^{v_0} v_s$. Thus, $v_s = M_{s+1}^{v_0} v_{s+1}$, and hence $v_0 = A_s^{v_0} v_s$. So

$$v_0 \in \bigcap_{s=1}^{\infty} A_s^{v_0}(C^n). \tag{7}$$

This fact will become particularly important in the proof of Theorem 4.14, hence Theorem D.

Definition 4.5 (Brun matrix). For each vector $v_0 \in C^r$, we call each matrix

$$A_n^{v_0} = M_1^{v_0} \cdots M_n^{v_0} \tag{8}$$

a *Brun matrix*. We let $\mathcal{B}_r$ denote the set of $r \times r$ Brun matrices.

Remark 4.6. When it is clear from the context, we may leave out reference to the starting vector v_0 and simply write A_s, T_s, M_s, etc.

Proposition 4.7. *Let Y_r be the set of rationally independent vectors in C^r, then for each $x \in Y_r$ there exists an $N \in \mathbf{N}$ so that A^x_n is a positive matrix for each $n > N$.*

Proof. We fix v_0 and omit the index v_0 from the notation below. We first prove that, for each i, and for each $h \in \mathbf{N}$, there exists some $m > h$ and some j such that $\vec{k}_m = (i, j)$. Starting with v_h and until $(v_m)_i$ becomes the largest coordinate, at each step one subtracts a number $\geq (v_m)_i = (v_h)_i$ from some coordinate $\geq (v_m)_i$. This can only happen a finite number of times before each coordinate apart from $(v_m)_i$ becomes less than $(v_m)_i = (v_h)_i$.

Consider A_n as in (6). To prove the proposition, it suffices to show that, for each (i, j), there exists a large enough N so that for all $n > N$ the (i, j)-th entry of A_n is positive. In fact, it is enough to show that this entry is positive for some A_n. Indeed A_{n+1} is obtained from A_n by adding one of its columns to another one of its columns, so that an entry positive in A_n, will be positive in A_{n+1}.

Fix i, j. Let

$$a = \min\{t \mid \vec{k}_t = (i, c) \text{ for some } c\},$$
$$b = \min\{t > a \mid \vec{k}_t = (j, d) \text{ for some } d\}.$$

Let c_1 be such that $\vec{k}_a = (i, c_1)$. Observe that, since c_1 is the second largest coordinate in v_a, in the next vector v_{a+1}, either i is still the largest coordinate or c_1 becomes the largest coordinate. There is some $N_1 \geq 1$ and some index c_2 so that

$$A_{a+N_1} = A_{a-1} M^{N_1}_{(i,c_1)} M_{(c_1,c_2)}.$$

We continue in this way, $A_n = A_{a-1} M^{N_1}_{(i,c_1)} M^{N_2}_{(c_1,c_2)} M^{N_3}_{(c_2,c_3)} \cdots M^{N_t}_{(c_t,c_{t+1})}$. When $n = b$, we have that $c_t = j$. Thus, for $n = b - 1$, we have

$$A_{b-1} = A_{a-1} M^{N_1}_{(i,c_1)} M^{N_2}_{(c_1,c_2)} M^{N_3}_{(c_2,c_3)} \cdots M^{N_t}_{(c_t,j)}.$$

It is elementary to see that the (i, j)-th entry of $M^{N_1}_{(i,c_1)} M^{N_2}_{(c_1,c_2)} M^{N_3}_{(c_2,c_3)} \cdots M^{N_t}_{(c_t,j)}$ is positive. This implies that the (i, j)-th entry of A_{b-1} is positive. □

4.3. Other versions of Brun's algorithm. To use the results of Schweiger's books, we must give two other different, but related, versions of Brun's algorithm.

Definition 4.8 (Brun's ordered algorithm). *Brun's ordered algorithm* is the fibered system (Δ^n, T') defined on $\Delta^n := \{x \in C^n \mid x_1 \geq \cdots \geq x_n\}$ by

$$T' \colon \Delta^n \longrightarrow \Delta^n,$$
$$(x_1, \ldots, x_n) \longmapsto (x_2, \ldots, x_{i-1}, x_1 - x_2, x_i, \ldots, x_n),$$

where $i = i(x) \geq 2$ is the first index so that $x_1 - x_2 \geq x_i$ and, if there is no such index, we let $i(x) = n$. The time-1 cylinders are

$$\Delta^n(i-1) = \{x \mid x_{i-1} > x_1 - x_2 \geq x_i\}.$$

Notice that the transformation T' is just left multiplication by the matrix $T_{1,2}$, followed by a permutation matrix that we denote P_i, determined by the cylinder $\Delta^n(i)$. (We denote $P_i T_{1,2}$ by T'_i and its inverse by M'_i.) Hence, given a sequence of indices $\omega(x) = (i_1, i_2, \dots)$ with $2 \leq i_j \leq n$ for each i_j, one obtains a sequence of matrices $\{T'_{i_j}\}_{j=1}^{\infty}$. Given $v_0 \in \Delta^n(i_1, \dots, i_m)$, this gives a sequence of points $v_0, \dots, v_m \in \Delta^n$ such that $v_{k+1} = T'_{i_j} v_k$ for each $1 \leq k \leq m-1$. The fibered system sequence here will be $\omega(x) = (i_1, i_2, \dots)$ when

$$x \in \bigcap_{m=1}^{\infty} \Delta^n(i_1, i_2, \dots, i_m).$$

If $\omega(x) = (i_1, i_2, \dots)$, we define

$$A_k^x = M'_{i_1} \cdots M'_{i_k} \tag{9}$$

for each $k \in \mathbf{N}$.

Definition 4.9 (Brun's homogeneous algorithm). *Brun's homogeneous algorithm* is the fibered system $(B_n, \overline{T'})$ defined on

$$B_n = \{x \in \mathbf{R}^n \mid 1 \geq x_1 \geq x_2 \geq \cdots \geq x_n \geq 0\}$$

and where $\overline{T'}\colon B_n \to B_n$ is such that the following diagram commutes:

$$\begin{array}{ccc} \Delta^{n+1} & \xrightarrow{T'} & \Delta^{n+1} \\ \downarrow p & & \downarrow p \\ B_n & \xrightarrow{\overline{T'}} & B_n \end{array}$$

where $p\colon \Delta^{n+1} - \{0\} \to B_n$ is defined by

$$p(x_1, \dots, x_{n+1}) = \left(\frac{x_2}{x_1}, \dots, \frac{x_{n+1}}{x_1}\right). \tag{10}$$

We denote the time-1 cylinders by $B_n(i)$.

4.4. Relating the algorithms

Definition 4.10. Let $\mathcal{O}\colon C^n \to \Delta^n$ be defined by

$$(x_1, \dots, x_n) \mapsto (x_{i_1}, \dots, x_{i_n})$$

where $x_{i_1} \geq x_{i_2} \geq \cdots \geq x_{i_n}$. Note that for a particular x, $\mathcal{O}(x)$ is a permutation.

Lemma 4.11. *For each $x \in C^n$ and for each $m \in \mathbf{N}$, there exist permutation matrices P_{i_1}, P_{i_2} so that*

$$A^x_m = P_{i_1}(A^{\mathcal{O}(x)}_m)' P_{i_2}.$$

Proof. This follows from the following commutative diagram:

$$\begin{array}{ccc} C^r & \xrightarrow{T} & C^r \\ \downarrow{\scriptstyle \mathcal{O}} & & \downarrow{\scriptstyle \mathcal{O}} \\ \Delta^r & \xrightarrow{T'} & \Delta^r \end{array} .$$

□

Corollary 4.12. *For each $x \in C^n$ and for each $m \in \mathbf{N}$, we have that A^x_m is a positive matrix if and only if $(A^{\mathcal{O}(x)}_m)'$ is a positive matrix.*

Corollary 4.13. *For each irrational $x \in \Delta^n$ there exists an N so that $(A^x_n)'$ is positive for all $n > N$.*

Proof. This follows from Proposition 4.7 and Corollary 4.12. □

4.5. Weak convergence and consequences. Recall the definitions of the $(r-1)$-dimensional simplex $\mathcal{S}_r$ in Definition 2.6 and the projection map $q\colon C^r \to \mathcal{S}_r$ of Equation 1. We will show that for each $r \geq 2$, the set of Perron–Frobenius eigenvectors for the transition matrices of positive automorphisms in $\mathrm{Aut}(F_r)$ is dense in the simplex $\mathcal{S}_r$.

It is proved in [19, Theorem 21, p. 5] that Brun's ordered algorithm is ergodic, conservative, and admits an absolutely continuous invariant measure. The proof uses Rényi's condition, which further says that the measure is equivalent to Lebesgue measure. We are very much indebted to Jon Chaika for pointing out to us that we could use the ergodicity of Brun's algorithm to prove the following theorem.

Theorem 4.14. *There exists a set $K \subset \mathcal{S}_r$ of full Lebesgue measure such that for each $x \in K$*

$$\bigcap_{j=1}^{\infty} A^x_j(\mathbf{R}^r_+) = \mathrm{span}_{\mathbf{R}_+}\{x\}. \tag{11}$$

Remark 4.15. Before proceeding with the proof, we explain what we saw as an impediment to proving that the PF eigenvectors are dense in a simplex. It is possible to have a sequence of invertible positive integer d x d matrices $\{M_i\}_{i=1}^{\infty}$ so that

$$\bigcap_{k=1}^{\infty} M_1 \cdots M_k(\mathbf{R}^d_+)$$

is more than just a single ray. The existence of such sequences of matrices was proved in the context of non-uniquely ergodic interval exchange transformations. There are several papers on the subject (including [15], [14], [21], [17]). Because it may not be straightforward to the reader outside of the field, we briefly explain how [14] implies the existence of such a sequence.

We consider a sequence of pairs of positive integers $\{(m_k, n_k)\}_{k=1}^{\infty}$ satisfying the conditions of [14, Theorem 5]. We look at

$$\bigcap_{k=1}^{\infty} A_{m_1,n_1} A_{m_2,n_2} \cdots A_{m_k,n_k} (\mathbf{R}_+^d),$$

as defined on p. 191. Keane explains on p. 191 that the product of any two successive A_{m_i,n_i} is strictly positive and that always $det(A_{m_k,n_k}) = 1$. We let $B_k = A_{m_1,n_1} A_{m_2,n_2} \cdots A_{m_k,n_k}$. Keane [14] projectivizes B_k to $\tilde{B}_k$, so that $\tilde{B}_k$ is a map of the 3-dimensional simplex $\mathcal{S}_4$. By Lemma 4, $\{\tilde{B}_k(0,1,0,0)\}_{k=1}^{\infty}$ is a sequence of vectors converging to a vector whose 2^{nd} entry is at least $\frac{1}{3}$. By Lemma 3 (when Theorem 5(ii) holds), $\{\tilde{B}_k(0,0,1,0)\}_{k=1}^{\infty}$ is a sequence of vectors converging to a vector whose 3^{rd} entry is at least $\frac{7}{10}$. But $\frac{1}{3} + \frac{7}{10} > 1$ and we have assumed that we are in $\mathcal{S}_4$. So these limits must be distinct vectors.

Proof of Theorem 4.14. Choose any N-cylinder $\Omega := \Delta^r(i_1, \dots i_N)$ such that the corresponding matrix $Z := A'_N$ is positive (see Corollary 4.13). Let $\bar{\Omega} := p(\Omega)$. Then $\mu(\bar{\Omega}) > 0$, where μ is the Lebesgue measure on B_r. Since $\bar{T}' : B_r \to B_r$ is ergodic with respect to the Lebesgue measure, by Birkhoff's Theorem, there exists a set $\bar{K} \subset B_r$ such that $\mu(\bar{K}) = 1$ and so that for each $\bar{x} \in \bar{K}$ the set $J(\bar{x}) := \{n \in \mathbf{N} \mid \bar{T}'^n(\bar{x}) \in \bar{\Omega}\}$ is infinite.

We let $K' := p^{-1}(\bar{K})$. Then for each $x' \in K'$ the set $I(x') := \{n \in \mathbf{N} \mid (T')^n(x') \in \Omega\}$ is infinite, as $n \in I(x')$ if and only if $n \in J(p(x'))$.

Let $K'' := \mathcal{O}^{-1}(K') \subset \mathbf{R}_+^r$. If $x \in K''$, then $\mathcal{O}(x) \in K'$, and hence for each $n \in I(\mathcal{O}(x))$ we have

$$(T')^n(\mathcal{O}(x)) \in \Omega.$$

Let $s \in \mathbf{N}$ be arbitrary. Consider the first s integers $\{j_1, \dots, j_s\}$ in $I(\mathcal{O}(x))$ satisfying that any difference between two of these numbers is $> N$ (where N came from the original N-cylinder we started with). Let $N_1 \geq j_s + N + 1$. Then, for each $m > N_1$,

$$(A_m^{\mathcal{O}(x)})' = D_1 \cdots D_{j_1-1} Z D_{j_1+N} \cdots D_{j_2-1} \cdots D_{j_s-1} Z D_{j_s+N} \cdots D_m,$$

where Z is the positive matrix that we started with and the D_i are the M_i''s of Brun's ordered algorithm. The matrix Z appears in this product s times. By Lemma 4.11, for each $x \in K''$ and each $m \in \mathbf{N}$, there exist permutation matrices P_{i_1}, P_{i_2} so that

$$A_m^x = P_{i_1} (A_m^{\mathcal{O}(x)})' P_{i_2}.$$

Hence, for this arbitrary s we have found an $N_1(s) \in \mathbf{N}$ so that for all $m > N_1$,

$$A_m^x = P_1 D_1 \cdots D_{j_1-1} Z D_{j_1+N} \cdots D_{j_2-1} \cdots D_{j_s-1} Z D_{j_s+N} \cdots D_m P_2,$$

where Z appears in this product at least s times and the other matrices in this product are all invertible nonnegative integer matrices. Then, by [6] (see also [11, Corollary 7.9]), Equation 11 holds true for each $x \in K''$.

Let $q\colon \mathbf{R}_+^r \to \mathcal{S}_r$ be the projection to the simplex in the positive cone (see Equation 1) and let μ' be the Lebesgue measure on $\mathcal{S}_r$. Define $K := q(K'')$. Then, since $\mu(\bar{K}) = 1$, we arrive at $\mu'(K) = 1$, as desired. □

Definition 4.16. Recall $\mathcal{B}_r$ from Definition 4.5, we let $\mathcal{P}_r$ be defined as

$$\mathcal{P}_r := \{v_{PF} \in \mathcal{S}_r \mid v_{PF} \text{ is the PF eigenvector for some positive Brun matrix } M \in \mathcal{B}_r\}.$$

Definition 4.17. Suppose

$$A = M_{i_1,j_1} \cdots M_{i_n,j_n} \in \mathcal{B}_r. \tag{12}$$

Then each M_{i_k,j_k} is an unfolding matrix as in (5) and we can associate to it the fold-automorphism $f_{i_k j_k}$ (see Definition 3.3). Notice that M_{i_k,j_k} is in fact the transition matrix for $f_{i_k j_k}$. To each $A \in \mathcal{B}_r$ as in (12) we associate the automorphism

$$g_A = f_{i_n,j_n} \circ \cdots \circ f_{i_1,j_1}, \tag{13}$$

whose transition matrix is A. We also call this automorphism g_v, where v is the PF eigenvector of A.

Theorem D. *For each $r \geq 2$, let $S_{l_1}^r$ denote the set of unit vectors according to the l_1 metric in $\mathbf{R}_+^r$. The set of Perron–Frobenius eigenvectors of matrices arising as the transition matrices of positive automorphisms in* $\mathrm{Aut}(F_r)$ *is dense in $S_{l_1}^r$.*

Proof. It will suffice to show that $\mathcal{P}_r$ is dense. By Theorem 4.14, we know that Brun's algorithm weakly converges on a dense set of points $K \subset \mathcal{S}_r$. Thus, given any $x \in \mathcal{S}_r$ and $\varepsilon > 0$, there exists some $x' \in B(x, \frac{\varepsilon}{2}) \cap K$ on which Brun's algorithm weakly converges. Hence, there exists some N such that, for each $n \geq N$, we have that $q(A_n^{x'} C^r) \subset B(x', \frac{\varepsilon}{2})$. By possibly replacing N with a larger integer, we can further assume that the $A_n^{x'}$ are positive (see Proposition 4.7), so have PF eigenvectors. And the PF eigenvector v_n for each $A_n^{x'}$ is contained in $A_n^{x'} C^r$ and hence is in $B(x', \frac{\varepsilon}{2})$. Hence, there exists a $v_i \in \mathcal{P}_r$ such that $d(v_i, x) < \varepsilon$. □

5. Dense Geodesics in theta complexes

In this section we construct a geodesic ray dense in the theta $\mathcal{T}_r$ subcomplex whose top-dimensional simplex has underlying graph as in the right-hand graph of Figure 1. One could consider this a warm-up to the proof of Theorem B or interesting in its own right, as this is the minimal connected subcomplex containing the projection of the Cayley graph for Out(F_r).

5.1. Construction of the fold ray. We enumerate the vectors in $\mathcal{P}_r$ from Definition 4.16 as $\{v_i\}_{i=1}^{\infty}$. For each i there exists a positive matrix A_{v_i} in $\mathcal{B}_r$ so that v_i is the PF eigenvector of A_{v_i}. Further, there exists an automorphism g_{v_i} corresponding to A_{v_i} (Display 13). We also enumerate all possible fold automorphisms, as in (3), by $h_1, \dots, h_{r(r-1)}$.

We construct a sequence that contains each $g_{v_i}^k \circ h_j$ with $i, k, j \in \mathbf{N}$:

$$g_{v_1} \circ h_1,\ g_{v_2} \circ h_1,\ g_{v_1}^2 \circ h_2,\ g_{v_1}^3 \circ h_3,\ g_{v_2}^2 \circ h_2,\ g_{v_3} \circ h_1,\ g_{v_4} \circ h_1 \dots \tag{14}$$

Decompose each g_{v_i} in (14) according to (13), to obtain an infinite sequence of fold automorphisms

$$\{\Phi_k\}_{k=1}^{\infty}. \tag{15}$$

For example, if $g_{v_1} = f_{i_n,j_n} \circ \dots \circ f_{i_1,j_1}$ then $\Phi_1 = h_1$, $\Phi_2 = f_{i_1,j_1}$, $\Phi_3 = f_{i_2,j_2}$, etc. We denote by G_1 the r-rose whose positive edges are identified with the chosen basis. This identification induces an enumeration of the edges. Using Definition 3.6, Φ_1 can be represented by the combinatorial fold f_1. The target graph G_2 is topologically an r-rose and it inherits an orientation and an enumeration of edges. Thus we may continue inductively to define the combinatorial fold $f_k \colon G_{k-1} \to G_k$. The enumeration of edges induces the homeomorphism $n_k \colon \hat{\sigma}_{(G_k, f_k \circ \dots \circ f_1)} \to \mathcal{S}_r$.

Lemma 5.1. *Let σ_0 be the base simplex, then there exists a point $x_0 \in \sigma_0$ so that, for each $k \geq 0$, the fold f_{k+1} is allowable in the folded rose after performing the sequence of folds $f_1, \dots, f_k$. Moreover, this folded rose is $x_k = n_k^{-1}(T_k \circ \dots \circ T_1(n_0(x_0)))$ in the cone $\hat{\sigma}_{(G, f_k \circ \dots \circ f_1 \circ m)}$.*

Proof. This follows from Corollary 3.17. Note that the positivity condition follows from the positivity of the matrices A_{v_i}. □

Remark 5.2. [11, Corollary 7.9] implies the metric on x_0 is unique, as the same positive matrix occurs in infinitely many of the products $M_{i_1} \cdots M_{i_k}$.

Definition 5.3 ($\mathcal{R}$). We let $\mathcal{R}$ denote the infinite fold ray (Definition 3.14) initiating at x_0 and defined by the sequence of folds $\{f_i\}$ as constructed above.

Theorem C. *For each $r \geq 2$, there exists a dense fold ray in $\mathcal{T}_r$ that projects to a dense Lipschitz geodesic fold ray in $\mathcal{T}_r/\operatorname{Out}(F_r)$.*

Proof. We recall $\mathcal{R}$ from Definition 5.3. It is clear that $\mathcal{R}$ is contained in $\mathcal{T}_r$. $\mathcal{R}$ is a geodesic ray by Corollary 3.19.

First recall that the simplicial and Lipschitz metrics on $\mathcal{X}_r$ induce the same topology on $\mathcal{X}_r$. Hence, it suffices to prove density in the simplicial metric.

Let $\bar{a} \in \mathcal{T}_r/\operatorname{Out}(F_r)$ and let $\varepsilon > 0$ be arbitrary. Lift $\bar{a}$ to a point $a \in \mathcal{T}_r$ in the interior of a top dimensional simplex τ. Let $y \in \tau$ be a point such that $d_s(a, y) < \varepsilon$, and so that its coordinates are rationally independent. The point y lies on a fold line $\mathcal{F}_{i,j}$ from a point x on one face of τ to a point z in another face. Without generality loss assume $z \in \sigma_0$, the base simplex. Moreover, without generality loss assume the combinatorial fold is $f_{1,2}$. Let $e'_1, \dots, e'_r$ denote the edges of R_r as numbered in σ_0. Enumerate the edges of G_τ, the underlying graph of τ, as $e_1, \dots e_{r+1}$, so that if $c\colon G_\tau \to R_r$ is the map collapsing e_2 in G_τ, then $c(e_1) = e'_1$ and $c(e_i) = e'_{i-1}$ for each $i \geq 3$. We parameterize the path $[x, z]$ as an unfolding path $\hat{\gamma}_{1,2}(z, t)$ in unprojectivized Outer Space as follows. Let $\zeta = n_0(z) \in \mathcal{S}_r$, then

$$n_\tau(\hat{\gamma}_{1,2}(z, t)) = (\zeta_1 + t, t, \zeta_2 - t, \zeta_3, \dots, \zeta_r).$$

We note that $z = \hat{\gamma}_{1,2}(z, 0)$, $x = \frac{1}{1-\zeta_2}\hat{\gamma}_{1,2}(z, \zeta_2)$, and for some $0 < t_0 \leq \zeta_2$ we have $y = \frac{1}{1-t_0}\hat{\gamma}_{1,2}(z, t_0)$. Since the function $\gamma_{1,2}(z, t) = \frac{1}{1-t}\gamma_{1,2}(z, t)$ is continuous when t is bounded away from 1, there exists an $\varepsilon' > 0$ so that for each $w \in B(z, \varepsilon')$, the path $\gamma_{1,2}(w, t)$ passes through $B(a, \varepsilon)$.

Now, by the density of PF eigenvectors (Theorem D), there exists a vector $v_i \in \mathcal{P}_r$ such that $v_i \in n_0(B(z, \varepsilon'))$. Let $\varepsilon'' > 0$ be such that $n_0(B(z, \varepsilon'))$ contains $B(v_i, \varepsilon'')$. Let K be large enough so that $q(A^k_{v_i}(\mathbf{R}^r_+)) \subset B(v_i, \varepsilon'')$ for all $k > K$.

Let Ψ_k be the composition of $\Phi_1, \Phi_2, \dots$ from Display 15 up to the first fold automorphism in the decomposition of $g^k_{v_i}$. Choose $k > K$ so that the last fold automorphism in Ψ_k is $\Phi_{1,2}$. Let z_k be the rose-point in $\mathcal{R}$ directly after performing the fold sequence of Ψ_k. Let x_k be the rose-point directly before z_k.

We claim that $\mathcal{R} \cdot \Psi_k^{-1}$ is ε-close to a. To see this recall that directly after z_k in $\mathcal{R}$ we perform the folds corresponding to $g^k_{v_i}$. Let w_k be the rose-point in $\mathcal{R}$ directly after these folds and let n_s, n_m be the appropriate identifications of the simplices containing z_k and w_k with $\mathcal{S}_r$. Then $n_s(z_k) = A^k_{v_i}(n_m(w_k))$. Hence $n_s(z_k)$ is ε''-close to v_i. Thus $n_s(z_k)$ is ε'-close to $n_0(z)$.

The point z_k is in $\sigma_0 \cdot \Psi_k$. Hence, by Lemma 3.10, $n_0(z_k \cdot \Psi_k^{-1}) = n_s(z_k)$. Thus the fold path $[x_k, z_k] \cdot \Psi_k^{-1}$, which is a path induced by the fold $\Phi_{1,2}$, satisfies that its endpoint, $z_k \cdot \Psi_k^{-1}$, is ε''-close to z. Thus, this fold line intersects $B(a, \varepsilon)$, as desired. □

6. Finding a rose–to–graph fold path terminating at a given point

This section is the first step in our expansion of our methods of Section 5 to obtain a dense geodesic ray in the full quotient of reduced Outer Space.

In this section and in the next one, we find, for a dense set of points y in reduced Outer Space, roses x, z so that $y \in [x, z]$ and the difference in marking map $x \to z$ is positive. The path $[x, z]$ will be called a positive rose–to–rose fold line. It will replace our basic fold lines $\mathcal{F}_{i,j}$ in the proof of Theorem C.

A rose–to–rose fold path will have two parts: a rose–to–graph part $[x, y]$ and a graph–to–rose part $[y, z]$. We begin in (Subsection 6.1) with decomposing the graph of y into a union of positive loops (Lemma 6.6). This allows us to find the rose x.

6.1. Decomposing a top graph into a union of directed loops

Definition 6.1 (paths and distance in trees). Let T be a tree. Then for each pair of points p, q in T there is a unique (up to parametrization) path from p to q. We denote its image by $[p, q]_T$ and, when there is no chance for confusion, we drop the subscript T. Given a tree T, let $d_T(\cdot, \cdot)$ denote the combinatorial distance in T.

Definition 6.2 (rooted trees). By a *rooted tree* we will mean a finite tree T with a preferred vertex v_0 called a *root*. A rooted tree can be thought of as a finite set with a partial order that has a minimal element - which is the root. We will refer to the partial ordering induced by the pair (T, v_0) as $\leq_T$, i.e. $w \leq_T w'$ when $w \in [v_0, w']_T$.

Remark 6.3. In the figures to follow the root will always appear at the bottom.

We use special spanning trees to guide us in finding the loop decomposition of G:

Definition 6.4 (good tree). Let T be a rooted tree in G and $e = (v, w)$ an edge of G. We call e *bad* if $v \not\leq_T w$ and $w \not\leq_T v$. Let $B(T)$ be the number of bad edges in G with respect to T. When $B(T) = 0$ we call T *good* (sometimes elsewhere called *normal*).

We prove a somewhat stronger version of [9, Proposition 1.5.6].

Proposition 6.5. *For each $G \in \mathcal{X}_r$, and for each $E \in E(G)$ that is not a loop, there exists a rooted spanning tree (T, v_0) so that $B(T) = 0$ and $E \in E(T)$ and $v_0 = \mathrm{ter}(E)$. Moreover, when G is trivalent with no separating edges, T can be chosen so that* $\deg_T(v_0) = 1$.

Proof. Given an edge E in G, we let $v_0 = \mathrm{ter}(E)$. Let $G_1, \dots, G_N$ denote the components of $G - \{v_0\}$ with $\{v_0\}$ added back to each component separately. Thus $G = \cup_{i=1}^N G_i/\{v_0\}$. We will construct a good tree T_i (rooted at v_0) in each G_i. Then $T = \cup_{i=1}^N T_i$ will also be a good tree rooted at v_0, since all edges $e \in E(G - T)$ have endpoints inside some G_i.

Let G_1 be the component containing the edge E. For $i \neq 1$ we choose a spanning tree T_i in G_i arbitrarily. For $i = 1$ we construct a spanning tree T_1 in G_1 such that $\deg(v_0)_{T_1} = 1$. Denote by v_1 the endpoint of E distinct from v_0. Since E is non-separating in G it is non-separating in G_1 so each vertex in G_1 can be connected to v_0 via E. We choose T_1 to be a spanning tree that does not include any edges of G_1 adjacent to v_0 other than E. We now define a complexity of a tree and show how to decrease it so that in the end of the process we get a good spanning tree.

Let S be a tree. For an edge $e = (v, w)$ in G, the union $[v_0, v]_S \cup [v_0, w]_S$ forms a tripod. Denote the middle vertex of this tripod by q_e, i.e. q_e satisfies $[v_0, q_e]_S = [v_0, v]_S \cap [v_0, w]_S$. Let $d(e) = d_S(v_0, q_e)$. We define the complexity $C(S) = (n(S), m(S))$ of S by defining $n(S)$ and $m(S)$ as follows:

$$n(S) = \begin{cases} -\min\{d(e) \mid e \text{ is bad}\} & \text{if } B(S) \neq 0, \\ -\infty & \text{otherwise,} \end{cases}$$

and let

$$m(S) = \#\{e \in G \mid e \text{ is bad and } d(e) = -n(S)\}.$$

Note that the complexity is always a pair of integers (or $-\infty$). Moreover, $m(S) = 0$ if and only if there are no bad edges, so that, if there are no bad edges, $C(S) = (-\infty, 0)$.

The complexity is ordered by the lexicographical ordering of the pairs. Note that for S bad, the complexity is bounded from below. Indeed, when $G \in \mathcal{X}_r$, we have $|V(G)| \leq 2r - 2$, so $n(S) \geq 2 - 2r$, hence $C(S) \geq (2 - 2r, 1)$.

We will modify each T_i separately to make it a good tree in G_i. Suppose $e = (v, w)$ is a bad edge in G_i realizing the minimal distance $-n(T_i)$. Let e_1 be the first edge of $[v, q_e]_{T_i}$ and let e_2 be the last edge of $[v, q_e]_{T_i}$. Let $T_i' = T_i \cup \{e\} \setminus \{e_2\}$. We claim that $e_2 \neq E$. If $G_i \neq G_1$, then this is obvious. Otherwise, $\deg_{T_1}(v_0) = 1$ and $\deg(q_e) \geq 2$, so that $v_0 \neq q_e$ and hence $e_2 \neq E$. Therefore, $E \in E(T_1')$ after the move and still $\deg_{T_1'}(v_0) = 1$. Next, notice that some bad edges of T_i have become good in T_i', for example e is no longer bad, as is any edge from a vertex in $[v, q_e]_{T_i}$ to a vertex in $[q_e, w]_{T_i}$. Some bad edges remain bad. But the only edges that were good and became bad are edges with one endpoint in $[v, q_e]_{T_i}$ and one endpoint in a component of $T_i \setminus \{v\}$ that does not contain e_1 or w. For such an edge f, we have that $q_f = v$, so $d(f) = d_{T_i'}(v, v_0) > d_{T_i'}(q_e, v_0) = d(e)$ and the complexity has decreased.

If G is trivalent with no separating edges, then no edge is a loop. This implies that there are three edges E, E', E'' incident at v_0. If the tripod $E \cup E' \cup E''$, was separating then each of its edges would be separating - a contradiction. Therefore, $G = G_1$ in this case, and the proof is complete. □

Lemma 6.6. *If G is a trivalent graph with no separating edges and a turn at v_0 involves the unoriented edges E and E', then there exists an orientation on the edges of G so that,*

(1) *$G = \cup_{i=1}^{r} \alpha_i$ where each α_i is a positive embedded loop,*

(2) *$\alpha_i \cap \left(\cup_{j=1}^{i-1} \alpha_j \right)$ is a connected arc containing v_0 for each i, and*

(3) *v_0 is the terminal point of both E and E'.*

We will in fact prove:

Lemma 6.7. *Let G be a trivalent graph with no separating edges and let (T, v_0) be a good spanning tree in G. Let $e_1 \in E(G - T)$ be so that $i(e_1) = v_0$. Then one can enumerate $E(G - T) - \{e_1\}$ as $e_2, \dots, e_r$ and orient all of the edges of G so that*

(1) *there exist positive embedded loops $\alpha_1, \dots, \alpha_r \subset G$,*

(2) *for each i, we have that i is the smallest index such that α_i contains e_i, and*

(3) *for each i, we have that $\alpha_i \cap (\cup_{j=1}^{i-1} \alpha_j)$ is a connected arc containing e_1.*

Proof of Lemma 6.6 (from Lemma 6.7). To prove Lemma 6.6, we use Proposition 6.5 to find a good tree (T, v_0) containing E and so that v_0 has valence 1 in T. Denote by e_1 the third edge at v_0 distinct from E, E'. Since $E \in E(T)$, we have $E', e_1 \notin E(T)$. By Lemm 6.7(2), α_1 is a loop containing e_1. Since E is the only edge of T adjacent to v_0, α_1 contains E. E is oriented towards v_0, so e_1 is oriented away from v_0. E' is not in T, so $E' = e_m$ for some m. Now by Lemma 6.7(3), α_m contains e_1, so E' must be oriented towards v_0. This proves Lemma 6.6(3). Lemma 6.6(2) follows form Lemma 6.7(3). By 6.7(3), $(\cup_{i=1}^{k-1} \alpha_i) \cap \alpha_k$ is an arc so the rank of $\cup_{i=1}^{k} \alpha_i$ is k (by Van-Kampen's Theorem). Therefore, the subgraph $G' = \cup_{i=1}^{r} \alpha_i$ is connected and has rank r. Since G has no valence-1 vertices, a subgraph G' that has rank r is in fact all of G. This proves item Lemma 6.6(1). □

We will need the following definitions in our proof of Lemma 6.7.

Definition 6.8 ($v_-(e)$ and $v_+(e)$). Let T be a good tree in G. Given an edge $e \in E(G)$, one of its endpoints is closer (in T) to v_0 than the other. We denote by $v_-(e)$ the vertex closer to v_0 and by $v_+(e)$ the vertex further from v_0.

Definition 6.9. Let α be an embedded oriented path in the graph G and let x, y be two points in the image of α, i.e $x = \alpha(s)$ and $y = \alpha(t)$, for some s, t. If $s < t$, then we denote by $[x, y]_\alpha$ the image of the subpath of α initiating at x and terminating at y, i.e. $\alpha([s, t])$.

Definition 6.10 (left–right splitting). Let v_0 be a vertex of a graph G and let α be an embedded directed loop based at v_0. Let e be an edge of α, and m the midpoint of e. Then the *left–right* splitting of α at e is

$$L^e_\alpha = [v_0, m]_\alpha, \quad R^e_\alpha = [m, v_0]_\alpha.$$

Definition 6.11 (aligned edges). Let G be a graph and (T, v_0) a spanning good tree in G. Suppose $f_1, f_2 \in E(T)$ satisfy that the vertices $\{v_+(f_1), v_+(f_2), v_0\}$ span a line in T. Then we call the pair of edges f_1, f_2 *aligned*. If $v_+(f_1) < v_+(f_2)$ we say f_1 *lies below* f_2.

Definition 6.12 (highlighted subpaths). Let G be a graph and (T, v_0) a spanning good tree in G. Let f_1, f_2 be aligned edges in T such that f_1 lies below f_2. For each i, let α_{f_i} be an embedded loop containing f_i and v_0. We define $H^{f_1,f_2}_{f_1}$ and $H^{f_1,f_2}_{f_2}$, the *highlighted* subpaths of α_{f_1} and α_{f_2}, respectively, as follows:

$$H^{f_1,f_2}_{f_1} = L^{f_1}_{\alpha_{f_1}} \quad \text{and} \quad H^{f_1,f_2}_{f_2} = R^{f_2}_{\alpha_{f_2}} \quad \text{when } i(f_1) <_T \text{ter}(f_1), \tag{16}$$
$$H^{f_1,f_2}_{f_1} = R^{f_1}_{\alpha_{f_1}} \quad \text{and} \quad H^{f_1,f_2}_{f_2} = L^{f_2}_{\alpha_{f_2}} \quad \text{when } i(f_1) >_T \text{ter}(f_1). \tag{17}$$

Proof of Lemma 6.7. Enumerate $E(G - T) - \{e_1\}$ so that $k < j$ implies that $v_-(e_k) \not\geq v_-(e_j)$. We prove this lemma by induction on i for $1 \leq i \leq r$. We will define a loop α_i and orient its previously unoriented edges so that Items (1)–(3) of the lemma hold and moreover the following Items 4 and 5 hold. Denote by α_e the first loop that contains e (for example when $e = e_j \notin E(T)$ then $\alpha_e = \alpha_j$).

(4) If $f \in E(T)$ and $f' \in E(G)$ are such that $v_+(f) \leq v_+(f')$ and f' is oriented, then f is oriented.

(5) Let $f_1, f_2 \in E(T)$ be aligned and $H^{f_1,f_2}_{f_1}, H^{f_1,f_2}_{f_2}$ the corresponding highlighted paths with respect to $\alpha_{f_1}, \alpha_{f_2}$, then $H^{f_1,f_2}_{f_1} \cap H^{f_1,f_2}_{f_2}$ contains no half-edges.

We include Item 5 to ensure that the loop in the induction step is embedded.

The base case. We begin the base of the induction with the edge $e_1 \in E(G-T)$, which is adjacent with v_0. Define the directed circle $\alpha_{e_1} = e_1 * [t(e_1), v_0]_T$, this is clearly an embedded loop. We orient the edges of α_{e_1} accordingly, i.e. e_1 is directed away from v_0, and the edges $f \in [t(e_1), v_0]_T$ are directed toward v_0. Clearly Items 1–5 hold at this stage, i.e. in the subgraph consisting of precisely α_1.

The induction hypothesis. We now assume that we have oriented some subset of $E(G)$ so that for each $j \leq i-1$ we have that e_j is oriented. We call $G' = \cup_{j=1}^{i-1} \alpha_j$ the oriented subgraph. We let $e := e_i$ and note that $e \notin G'$.

The induction step. Let $I = [v_-(e), v_+(e)]_T$. Let t_1 be the edge of I adjacent at $v_-(e)$. We claim as follows that t_1 is oriented, see Figure 2. Indeed, if $v_-(t_1) = v_0$ this follows from the base case. Otherwise, there is an edge $t_3 \in E(T)$ adjacent to $v_-(t_1)$. Since t_3 is nonseparating, there is an edge $e' \in E(G-T)$ so that $v_-(e') \leq v_-(t_3)$ and $v_+(e') \geq v_+(t_3)$. But since the graph is trivalent, $v_+(e') \geq v_+(t_1)$. Now, since $v_-(e') < v_-(e)$, we have that e' is oriented. Hence, by Item 4 in the induction hypothesis, t_1 is oriented.

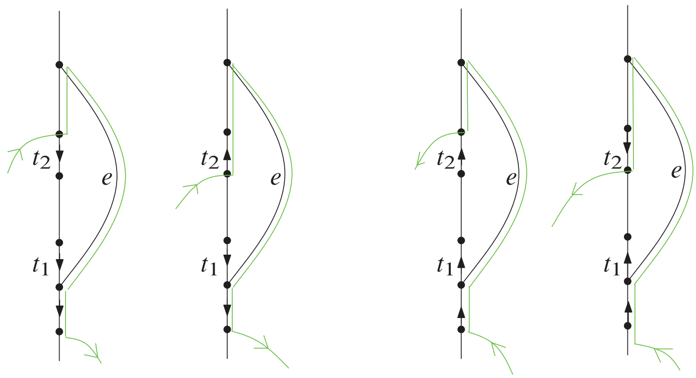

Figure 2. The green line indicates the path α_e takes near e. The edge t_1 is the first edge in I and t_2 is the last oriented edge in I.

Let t_2 be the last oriented edge of I. Denote $J = [v_+(t_2), v_+(e)]_T$. Note that we allow $t_1 = t_2$ and leave the adjustments of this case, from Cases A and C of Figure 2, to the reader. The loop α_e is constructed as in Figure 2 from the following segments by adding or removing a half-edge of t_2 or removing a half-edge of t_1 or t_2:

$$H_{t_2}^{t_1,t_2}, \quad J, \quad e, \quad H_{t_1}^{t_2,t_1}. \tag{18}$$

The orientation of α_e is chosen according to the direction of t_1. When $i(t_1) >_T t(t_1)$, Cases A and B of Figure 2, we orient α_e so that $H_{t_2}^{t_1,t_2}$ is the left (first) segment and when $i(t_1) <_T t(t_1)$, Cases C and D, we orient so that $H_{t_1}^{t_2,t_1}$ is the first segment.

We observe that α_e is indeed embedded as follows. The paths $H_{t_2}^{t_1,t_2}, H_{t_1}^{t_2,t_1} \subset G'$ while $J, e \subset G - G'$, hence they are edge-disjoint. Moreover, by Item 5 for G', the paths $H_{t_2}^{t_1,t_2}, H_{t_1}^{t_2,t_1}$ are half-edge disjoint so that even if we add t_1 or t_2 they remain edge disjoint. Therefore the path α_e does not self-intersect in an edge. It cannot self-intersect at a vertex since the graph is trivalent. Moreover, the loop is positive. This proves Item 1.

Note that α_e is contained in $G' \cup T \cup \{e\}$, hence it does not contain e_j for $j > i$. This proves Item 2. Item 3 is also clear. Item 4 is satisfied since for $f' = e$ and for all $f \in E(T)$ so that $v_+(f) \leq v_+(e)$ we have that f is oriented.

We are left with proving item 5 for each pair of aligned edges $f_1, f_2 \in E(T)$ such that at least one of them is oriented in the i^{th} step, i.e. $\alpha_f = \alpha_i$. It is less difficult to check that the claim holds when both edges are newly oriented. We leave this case to the reader and check the case where $f_1 \subset G' \cap T$ and $f_2 \subset T$ is newly oriented, i.e. $f_2 \subset J$. We illustrate the different cases in Figure 3.

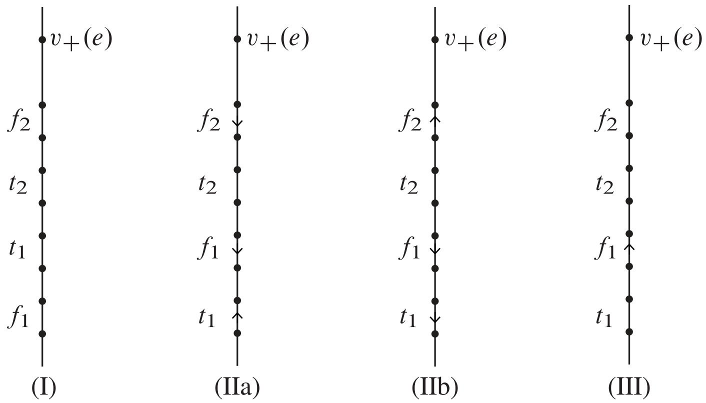

Figure 3. Checking Claim 5 for f_1, f_2 depending on their location. In Case I, f_1 lies below t_1, and in the other cases t_1 lies below f_1. In Cases IIa, IIb, f_1 is pointing down and in Case III, f_1 is pointing up.

Suppose $f_1 \notin I$ (Case I in Figure 3), then for each $f \in \{f_2, t_1, t_2\}$, we have $f_1 <_T f$. Thus, $H_{f_1}^{f_1,f} = H_{f_1}^{f_1,f'}$ for $f, f' \in \{f_2, t_1, t_2\}$.

Moreover, by our construction of $\alpha_{f_2} = \alpha_e$, we have $H_{f_2}^{f_1,f_2} \subset H_{t_k}^{f_1,t_k} \cup J \cup \{t_2\} \cup \{e\}$ for either $k = 1$ or $k = 2$. Thus, for the same k we have (see Figure 2)

$$H_{f_1}^{f_1,f_2} = H_{f_1}^{f_1,t_k} \quad \text{and} \quad H_{f_2}^{f_1,f_2} \subset H_{t_k}^{f_1,t_k} \cup \{e\} \cup J \cup \{t_2\}.$$

Since J is newly oriented, $H_{f_1}^{f_1,t_k}$ and J are edge-disjoint. By the induction hypothesis, $H_{f_1}^{f_1,t_k} \cap H_{t_k}^{f_1,t_k}$ contains no edges or half-edges. Moreover, if $H_{f_2}^{f_1,f_2} \not\subset H_{t_k}^{f_1,t_k} \cup J \cup \{e\}$ then $k = 2$ and a half of t_2 must be in $H_{t_2}^{f_1,t_2}$. Since $H_{f_1}^{f_1,t_k} \cap H_{t_k}^{f_1,t_k}$ contains no half-edges, then t_2 is not contained in $H_{f_1}^{f_1,t_k}$. This implies that $H_{f_1}^{f_1,f_2} \cap H_{f_2}^{f_1,f_2}$ contains no edges or half-edges.

If $f_1 \in I$, then there are two classes of cases: $i(f_1) \geq_T \text{ter}(f_1)$ (Cases IIa, IIb of Figure 3) and $i(f_1) \leq_T \text{ter}(f_1)$ (Case III). We will prove Case IIa and leave the others to the reader. If f_1 is pointing down, then $H_{f_1}^{f_1,f_2} = R_{\alpha_{f_1}}^{f_1}$, $H_{f_2}^{f_1,f_2} = L_{\alpha_{f_2}}^{f_2}$. There are two subcases: f_2 is pointing down (Case IIa) or up (Case IIb). If f_2 is pointing down, then t_1 is pointing up (see Figure 2), thus $L_{\alpha_e}^{f_2} \subset L_{\alpha_{t_1}}^{t_1} \cup e \cup J$. In this configuration,

$$\begin{aligned} H_{f_1}^{f_1,t_1} &= R_{\alpha_{f_1}}^{f_1} & H_{t_1}^{f_1,t_1} &= L_{\alpha_{t_1}}^{t_1} \\ H_{f_1}^{f_1,f_2} &= R_{\alpha_{f_1}}^{f_1} & H_{f_2}^{f_1,f_2} &= L_{\alpha_e}^{f_2} \subset L_{\alpha_{t_1}}^{t_1} \cup \{e\} \cup J \end{aligned}$$

Hence, the fact that $H_{f_1}^{f_1,t_1} \cap H_{t_1}^{f_1,t_1}$ contains no half-edges implies that $H_{f_1}^{f_1,f_2} \cap H_{f_2}^{f_1,f_2}$ contains no half-edges. The other cases are similar. □

6.2. Rose–to–graph fold line.

Given a point x whose underlying graph is trivalent with no separating edges, we wish to find a rose-point x_0 and a line in $\mathcal{X}_r$ from x_0 to x. This is done by simultaneous folding as defined below.

Definition 6.13 (rose–to–graph fold line $\mathcal{F}(x_0, \{s_{ij}\})$). Let $x_0 = (R, \mu, \ell_0)$ be a point in $\widehat{\mathcal{X}}_r$ whose underlying graph is a rose with r petals. There are $K = r(r-1)$ turns and we enumerate them in any way $\{\tau_i\}_{i=1}^{K}$. Let $\vec{s} \in \mathbf{R}^K$ be a nonnegative vector so that s_i is no greater than the length of each edge in the turn τ_i. Given the data $(x_0, \vec{s})$ we construct a continuous family of graphs $\{x_t\}$ for $0 \leq t \leq T = \sum s_i$, and maps $f_{t,0}\colon x_0 \to x_t$ as follows. In the i^{th} step let $\tau_i = \{e_j, e_m\}$ and fold initial segments of length s_i in $f_{t,0}(e_j)$ and $f_{t,0}(e_m)$. We caution that these $f_{t,0}$ are not always homotopy equivalences. However, if $f_{T,0}$ is a homotopy equivalence, then for each $t < T$ the map $f_{t,0}$ is a homotopy equivalence. In this case we get a path

$$\begin{aligned} \mathcal{F}(x_0, \vec{s})\colon [0, T] &\longrightarrow \widehat{\mathcal{X}}_r, \\ t &\longmapsto x_t. \end{aligned}$$

We denote its projectivization by $\overline{\mathcal{F}} = q(\mathcal{F})$.

Lemma 6.14. *Let G be a trivalent graph such that $\pi_1(G) \cong F_r$. Then for each $x \in \widehat{\mathcal{X}}_r$ with underlying graph G, there exists a point x_0 whose underlying graph is a rose and there exists a nonnegative vector $\vec{s}$ in $\mathbf{R}^K$, where $K = r(r-1)$, so that*

$$x = \mathcal{F}(x_0, \vec{s})(T).$$

Additionally, $x_0, \vec{s}$ are linear functions of the lengths of $E(G)$ as they vary throughout the unprojectivized simplex. (As above, $T = \sum s_i$.)

Proof. Let $E_1, \dots E_{3r-3}$ be the edges in G. Lemma 6.6 provides a decomposition of G as $G = \cup_{i=1}^r \alpha_i$. Let $R = \sqcup_{i=1}^r \alpha_i / \{v_0\}$, then R is an r-petaled rose. Let $e_1, \dots e_r$ be the edges of R. Let $\ell_R(e_i) = \ell_G(\alpha_i)$. We write α_i as a sequence of edges $\alpha_i = E_{m(i,1)} \cdots E_{m(i,k_i)}$. Then

$$\ell_R(e_i) = \ell_G(\alpha_i) = \sum_{j=1}^{k_i} \ell_G(E_{m(i,j)}). \tag{19}$$

Let x_0 denote R with these edge lengths. There is a natural map $\theta \colon x_0 \to x$ defined by the inclusion of α_i in G. This is a quotient map and, moreover, $\theta|_{e_i}$ is an isometry.

Recall that the intersection of α_i and α_j is an arc containing v_0. Let a, b be the endpoints of the arc. For $\tau_k = (e_i, e_j)$ and $\tau_m = (\overline{e_i}, \overline{e_j})$ we define

$$s_k = l_G([v_0, a]_{\alpha_i}) \quad \text{and} \quad s_m = l_G([b, v_0]_{\alpha_i}). \tag{20}$$

Consider the folding line $\mathcal{F}(x_0, \vec{s})$. By the definitions, θ is precisely the map $f_{T,0} \colon x_0 \to x_T$, since the points identified by θ are precisely those that are identified in the folds. Therefore, x_T equals the point x that we started with.

Moreover, θ is a homotopy equivalence by Lemma 6.6(2). Therefore, $\mathcal{F}(x_0, \vec{s})$ is a path in unprojectivized Outer Space. Equations 19 and 20 show that the dependence of l_i and s_k on edge lengths in x is linear. □

Lemma 6.15. *Suppose that $x = \mathcal{F}(x_0, \vec{s})(T)$ and the underlying graph of x is a trivalent graph. Then there exists a neighborhood of $(x_0, \vec{s})$ so that for each $(y_0, \vec{u})$ in this neighborhood, the endpoint $y := \mathcal{F}(y_0, \vec{u})(T')$ of the fold line $\mathcal{F}(y_0, \vec{u})$ lies in the same unprojectivized open simplex in $\widehat{\mathcal{X}}_r$ as x.*

Additionally, the edge lengths of y are linear combinations of the edge lengths of y_0 and $\vec{u}$.

Proof. Consider the positive edges $E_1, \dots E_m$ in G and let $G = \cup_{i=1}^r \alpha_i$ be the decomposition guaranteed by Lemma 6.6. The edge E_i is contained in a loop $\alpha_{j(i)}$. Since G is a trivalent graph, at each endpoint $\{v, w\}$ of E_i there is an edge, E_k, E_d resp., not contained in α_i. The edges E_k, E_d are contained in $\alpha_{j(k)}, \alpha_{j(d)}$ resp.

Now $v_0, v \in \alpha_{j(i)} \cap \alpha_{j(k)}$. Thus, by Lemma 6.6(3), either $[v_0, v]_{\alpha_{j(i)}} \subseteq \alpha_{j(i)} \cap \alpha_{j(k)}$ or $[v, v_0]_{\alpha_{j(i)}} \subseteq \alpha_{j(i)} \cap \alpha_{j(k)}$. This situation is similar for v_0, w. Therefore,

$$\ell_G(E_i) = \begin{cases} s_m - s_n & \text{when } \tau_m = (e_{j(d)}, e_{j(i)}) \text{ and } \tau_n = (e_{j(k)}, e_{j(i)}), \\ s_n - s_m & \text{when } \tau_m = (\overline{e_{j(d)}}, \overline{e_{j(i)}}) \text{ and } \tau_n = (\overline{e_{j(k)}}, \overline{e_{j(i)}}), \\ |\ell(\alpha_{j(i)}) - (s_n + s_m)| & \text{otherwise.} \end{cases} \tag{21}$$

Note, this dependence of the $\ell_G(E_i)$ on the variables s_m will be the same for all points in the same unprojectivized simplex, as they only depend on the loop decomposition. We also get that the dependence of $\ell_G(E_i)$ on s_{ij} and $\ell(e_i)$ is linear. Let U be the open subset of the unprojectivized rose simplex cross $\mathbf{R}_+^{r(r-1)}$ so that each expression in the right-hand side of (21) is positive for each i. This is an open set containing $(x_0, \vec{s})$. For any point y in the unprojectivized simplex of x one can use (19) and (20) to get y_0 and $\vec{u}$ so that $y = \mathcal{F}(y_0, \vec{u})$. Hence for each (y_0, u) in this neighborhood the point $\mathcal{F}(y_0, \vec{u})(T')$ is in the same unprojectivized open simplex as x. □

6.3. Folding a transitive graph to a rose

Definition 6.16. A *transitive* graph G is a directed graph G with the following property: for any two vertices w, w' there exists a directed path from w to w'.

Note that it is enough to check that for any choice of preferred vertex v, there exists a directed path to and from each other vertex v'.

The proof of the following is left to the reader.

Observation 6.17. Let G be a directed graph and let $f : G \to G'$ be a direction matching fold of two oriented edges e_1, e_2 in G, starting at a common vertex v. Then

(1) if G is transitive, then G' is transitive;

(2) if the lengths of the edges of G are rationally independent, then the lengths of the edges of G' are rationally independent.

Lemma 6.18. *Let G be any transitive graph with rationally independent edge lengths and let $\{E, E'\}$ be either a positive or negative turn. Then there exists a fold sequence $f_1, \ldots, f_k$ containing only direction matching folds and satisfying that $f_k \circ \cdots \circ f_1(G)$ is a rose. Further, assuming G is trivalent, we may choose f_1 so that it folds the turn $\{E, E'\}$.*

Proof. We perform the following steps.

Step 1. Let $c(G)$ denote the number of directed *embedded* paths in G between (distinct) vertices. For each pair of vertices w, w' there exists a directed path from w to w', thus it follows that there exists an embedded directed path from w to w'. We decrease $c(G)$ by folding two directed embedded paths α, β so that $i(\alpha) = i(\beta)$ and $\mathrm{ter}(\alpha) = \mathrm{ter}(\beta)$ and $\alpha \cap \beta = \{i(\alpha), \mathrm{ter}(\alpha)\}$. Note that the edges in α and β are distinct and thus we are consecutively performing folds as in Definition 3.5. Also, if G is trivalent and we choose a decomposition as in Lemma 6.6, then we can choose the first α to contain E and the first β to contain E' and fold α, β so that the first combinatorial fold folds the turn $\{E, E'\}$. Denote the new graph by G'. Then, by Observation 6.17, G' is transitive and has rationally independent edge lengths. The complexity has decreased, i.e. $c(G') < c(G)$.

At the end of this step we may assume that we have a connected graph G so that $G = \coprod_{i=1}^{n} \gamma_i / \sim$, where γ_i is a circle and for each $i \neq j$: $\gamma_i \cap \gamma_j$ is either empty or a single point (see Figure 4). We call such a graph a *gear graph*. Notice that given a gear graph, such a decomposition into circles is unique up to reindexing.

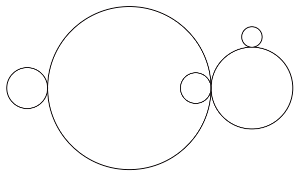

Figure 4. This graph is an example of a gear graph.

Step 2. For a gear graph, we define a new complexity. Let V be the set of vertices (of valence > 2). Our object is to remove one by one all vertices other than v_0.

Let $w \neq v_0$ be a vertex and let γ_1, γ_2 be circles so that $w = \gamma_1 \cap \gamma_2$ and suppose that $v \neq w$ is a vertex on γ_2 (see Figure 5). Let $\gamma_2 = \alpha\beta$ where $i(\alpha) = v = \mathrm{ter}(\beta)$ and $\mathrm{ter}(\alpha) = w = i(\beta)$. By folding, wrap β over γ_1 until v is on the image of γ_1 (we may have to wrap β multiple times over γ_1). Now there are two paths from v to w: α and γ_1', the remaining part of γ_1. Fold $\bar{\alpha}$ over $\overline{\gamma_1'}$. G' is a gear graph. Moreover, the valence of the image of w decreases. We continue this process for all loops based at w until it becomes a valence-2 vertex and we drop it out of the set of vertices. When all vertices other than v_0 have valence 2 then G is a rose. $\square$

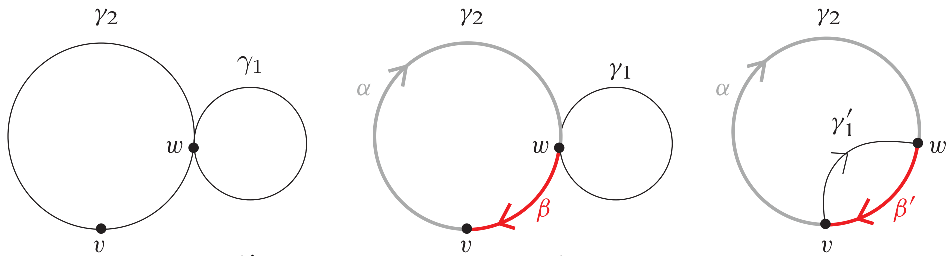

Figure 5. Step 2 (β' is the remaining portion of β, after it is wrapped around γ_1) .

7. Existence and continuity of rose–to–rose fold lines

In Subsection 6.2 we defined a rose–to–graph fold line $[x_0, x] = \mathcal{F}(x_0, \vec{s})$, given a point x and a loop decomposition of its underlying graph. We also had two continuity statements: (1) $x_0, \vec{s}$ vary continuously as a function of x by Lemma 6.14 and (2) $\mathcal{F}(x_0, \vec{s})$ vary continuously as a function of $x_0, \vec{s}$ by Lemma 6.15. We need similar existence and continuity statements for the full rose–to–rose fold line.

Proposition 7.1. *Let x be any point of $\widehat{\mathcal{X}}_r$ satisfying that $q(x)$ is in a top-dimensional simplex of reduced Outer Space with rationally independent edge lengths. Let $\{E, E'\}$ be a pair of adjacent edges in the underlying graph of x. Then there exists a positive rose–to–rose fold line, which we denote by $\mathcal{R}(x_0, \vec{s})\colon [0, L] \to \mathcal{X}_r$, containing x and containing the fold of $\{E, E'\}$.*

Proof. By Lemma 6.14 there exists a rose point x_0 and a vector $\vec{s}$ so that the rose–to–graph fold segment $\mathcal{F}(x_0, \vec{s})$ terminates at the point x. We consider the orientation on x given by the loop decomposition in Lemma 6.6. We may apply Lemma 6.18 to obtain a fold sequence $f_1, \ldots, f_k$ which terminates in a rose z' with some valence-2 vertices and so that f_1 folds the turn $\{E, E'\}$. Removing the valence-2 vertices gives a rose, which we denote by z. The line just described will be denoted $\mathcal{R}(x_0, \vec{s})$. It satisfies the statement in the theorem. □

Notation. Let x_0 be a rose, let $\vec{s} \in \mathbf{R}^{r(r-1)}$, and let $\mathcal{R}(x_0, \vec{s})$ be the fold line defined by these parameters. We will denote by x the trivalent metric graph at the end of the rose–to–graph fold segment, i.e. $x = \mathcal{R}(x_0, \vec{s})(\sum s_i)$, by z the end-point of the rose–to–rose fold segment, and by z' the point in the graph–to–rose fold segment from which z is obtained by removing the valence-2 vertices.

Definition 7.2 (proper fold line). Let $\mathcal{R}(x_0, \vec{s})$ be a rose–to–rose fold line and let $f_1, \ldots, f_k$ be the sequence of combinatorial folds from x to z'. If for each l the fold f_l is a proper full fold, then we will say that $\mathcal{R}(x_0, \vec{s})$ is a *proper fold line*.

For example, for each x as in Proposition 7.1, the fold line constructed in the proposition is a proper fold line since the edge lengths in x are rationally independent.

Proposition 7.3. *For each proper rose–to–rose fold line $\mathcal{R}(x_0, \vec{s})$ and $\varepsilon > 0$, there exists a neighborhood U of $(x_0, \vec{s})$ so that for any point $(y_0, \vec{u}) \in U$ the following holds.*

(1) *The endpoints of the rose–to–graph fold segments*

$$y := \mathcal{F}(y_0, \vec{u})(T'), \quad x := \mathcal{F}(x_0, \vec{s})(T)$$

lie in the same unprojectivized open simplex and are ε-close.

(2) *The sequence of combinatorial folds from x to z' appearing in the graph–to–rose fold segment is allowable in y.*

(3) *Let $\mathcal{R}(y_0, \vec{u})$ be the fold line defined by concatenating $\mathcal{F}(y_0, \vec{u})$ with the fold segments from (2), then the terminal points $w := \mathcal{R}(y_0, \vec{u})(L')$ and $z := \mathcal{R}(x_0, \vec{s})(L)$ are ε-close.*

Proof. By Lemma 6.15 there exists a neighborhood U of $(x_0, \vec{s})$ so that for each $(y_0, \vec{u}) \in U$ the fold line $\mathcal{F}(y_0, \vec{u})$ terminates at some point y lying in the same unprojectivized top-dimensional simplex as x. Since the edge lengths of y vary continuously with the edge lengths of y_0 and $\vec{s}$, we can make U smaller, if necessary, to ensure that y and x are ε-close. This proves Item 1.

Let $f_1, \dots, f_k$ be the combinatorial fold sequence from x to z', as described in Lemma 6.18. For each combinatorial proper full fold folding e_i over e_j, there corresponds a square $m \times m$ folding matrix (here $m > r$), which we denote $T'_{ij} = (a_{kl})$, so that $a_{kl} = 1$ for $k = l$ and $a_{ij} = -1$, but otherwise $a_{kl} = 0$ when $k \neq l$ (compare with Definition 3.4). Let $T'_1, \dots, T'_k$ be the fold matrices corresponding to $f_1, \dots, f_k$. Then just as in Lemma 3.15, the combinatorial fold sequence $f_1, \dots, f_k$ is allowable in the point y if and only if for each l the vector $T'_l \cdots T'_1(\ell(y))$ is positive for each $1 \le l \le k$. This defines an open neighborhood of x where Item 2 holds. Item 3 follows from the fact that matrix multiplication is continuous. □

Lemma 7.4. *Let $h\colon x_0 \to z$ be a homotopy equivalence representing the map from the initial rose x_0 to the terminal rose z in a rose–to–rose fold line. Let H be a matrix representing the change-of-metric from the rose z to the rose x_0. Then H is a nonnegative invertible integer matrix and it equals the transition matrix of h.*

Proof. Let x_0 be the initial rose with edges $e_1, \dots, e_r$ and let z be the terminal rose with edges $e'_1, \dots, e'_r$. Let z' be the graph on the fold line just before z, i.e. z is obtained from z' by unsubdividing at all of the valence-2 vertices. Let $h\colon x_0 \to z$ and $g\colon x_0 \to z'$ be the relevant homotopy equivalences. Since g is a subdivision followed by folding maps, each of which is a positive map, we have for each i that $g|_{e_i}$ is a local isometry. Thus, $h|_{e_i}$ is a local isometry. Moreover, h maps the unique vertex of x_0 to the unique vertex of z. Suppose $h(e_i)$ contains a part of an edge e'_i, then $h(e_i)$ contains a full appearance of e'_i (since there is no backtracking and the vertex maps to the vertex). Therefore, $h|_{e_i}$ is an edge-path in z. Thus, we may write

$$l(e_i) = \sum_j m(i,j) l(e'_j), \tag{22}$$

where the $m(i,j)$ are the nonnegative integer entries of the transition matrix of h. By Equation 22, the change-of-metric matrix from z to x_0 coincides with the transition matrix of the homotopy equivalence h. Therefore, H is nonnegative, and integer. Moreover, since all the folds in a rose–to–rose fold line are direction matching, h is a positive map. Thus, H is equal to $Ab(h)$, the map induced by h (viewed as an automorphism) by abelianization. Therefore, H is invertible. □

Lemma 7.5. *For each $\varepsilon > 0$ and proper fold line $\mathcal{R}(x_0, \vec{s})$, there exists a neighborhood U of the terminal rose z such that, for each $w \in U$, there exists a proper rose–to–rose fold line $\mathcal{R}(y_0, \vec{u})$ terminating at w satisfying that*

(1) *the top graphs x, y are ε-close,*

(2) *the combinatorial fold sequence corresponding to the graph–to–rose segments are the same in both lines, and*

(3) *the change-of-metric matrix for both fold lines is the same.*

Proof. We prove (1) and (2). By Lemma 7.4, the change-of-metric matrix H from z to x_0 is nonnegative. Thus, for any w in the same unprojectivized simplex as z, we have that $H\ell(w)$ is also positive. By Proposition 7.3, there exists a neighborhood V of x_0 so that for each $y_0 \in V$ there exists a vector $\vec{u}$ so that if y is the top graph of the fold line $\mathcal{F}(y_0, \vec{u})$, then x and y are ε-close and their combinatorial fold sequences are the same. The neighborhood U can be taken in $H^{-1}(V)$. This proves (1) and (2). Since the combinatorial folds are the same, the transition matrix $y_0 \to w$ is the same as the transition matrix $x_0 \to z$. (3) follows from Lemma 7.4. □

Definition 7.6 (rational rose–to–rose fold lines). A rose–to–rose fold line $\mathcal{R}(x_0, \vec{s})$ is called *rational* if for each edge e in G and for each loop α_i in the loop decomposition of the underlying graph of x, the quotient $\frac{l(e,x)}{l(\alpha_i,x)}$ is rational.

Proposition 7.7. *For each $\varepsilon > 0$ and for each x in an unprojectivized top-dimensional simplex of reduced Outer Space, with rationally independent edge-lengths, there exists a rational proper rose–to–rose fold line passing through some x' in the same open unprojectivized simplex as x, satisfying that $d(x', x) < \varepsilon$.*

Proof. Let x be a point in an unprojectivized top-dimensional simplex and having rationally independent edge lengths. Let $\mathcal{R}(x_0, \vec{s})$ be a proper rose–to–rose fold line containing x. Let V be a neighborhood of $(x_0, \vec{s})$ guaranteed by Proposition 7.3. Let U be a neighborhood of x in the same top-dimensional simplex and such that for each $y \in U$ there exists some $(y_0, \vec{u}) \in V$ with y the terminal point of $\mathcal{F}(y_0, \vec{u})$. This is possible by Lemma 6.14. Let x' be a point in U which is ε-close to x and such that the ratios $\frac{l(e',x')}{l(\alpha_i,x')}$ are rational. Then the resulting fold line through x' will have the required properties. □

8. Constructing the fold ray

Enumerate $\mathcal{P}_r$ by $\{v_i\}_{i=1}^{\infty}$ (see Section 4). There are countably many rational proper rose–to–rose fold lines with rational edge lengths. Thus there are countably many such rose–to–rose fold lines terminating arbitrarily close to the rose with length-vector v_i. For such lines that land sufficiently close to v_i there is a rose–to–rose fold line terminating at v_i that fellow-travels the rational one, i.e. follows the same fold sequence. For each such fold line $\mathcal{R}_{ij}$, let F_{ij} denote its folding matrix and H_{ij} its inverse – an invertible nonnegative matrix. Let U_{ij} denote the neighborhood of v_i from Lemma 7.5 (for $\varepsilon = 1$ will suffice), i.e. for each $w \in U_{ij}$ there exists a proper fold line terminating at w, passing through the same simplices as $\mathcal{R}_{ij}$ and satisfying that their fold matrices are the same.

Since each $A_i = A_{v_i}$ is a positive matrix, for each i, j there exists an integer $n(i, j)$ satisfying that for each $n > n(i, j)$ we have $q(A_i^n(\mathbf{R}_+^r)) \subset U_{ij}$.

Recall that g_{v_i} has a decomposition into fold automorphisms obtained from Brun's algorithm, this induces a decomposition of A_i into unfolding matrices and A_i^{-1} into folding matrices. We then create a sequence of pairs denoted $\{a_k\}$ which satisfies:

(1) if $a_k = (i, j)$ for an odd k then there exists some n such that $n > n(i, j)$ so that $a_{k+1} = (i, n)$;

(2) for each $i, N \in \mathbf{N}$ there exists an $n > N$ and an even k so that $(i, n) = a_k$;

(3) for each $i, j \in \mathbf{N}$, there exist infinitely many odd k's so that $a_k = (i, j)$.

To each a_k in this sequence we attach an (unfolding) matrix or sequence of matrices and automorphisms: if k is odd and $a_k = (i, j)$, we attach the matrix H_{ij} related to the rose–to–rose fold line $\mathcal{R}_{ij}$ and define f_k to be its change-of-marking automorphism. And if k is even and $a_k = (i, n)$, we attach a sequence of unfolding matrices according to a Brun's algorithm decomposition $A_i^n = (M_1^i \dots M_k^i)^n$ and a sequence of fold automorphisms. We get a sequence of matrices, which we denote by $\{D_l\}_{l=1}^{\infty}$. For each l, either $D_l = H_{ij}$ or $D_l = M_j^i$ for some i, j. We emphasize that we have not decomposed H_{ij}. Moreover, we get a sequence of automorphisms $\{f_l\}_{l=1}^{\infty}$. Let $\{Z_l\}_{l=1}^{\infty}$ denote the sequence of folding matrices, i.e. $Z_l = D_l^{-1}$ for each l.

Definition 8.1 (ray $\mathfrak{R}$). We construct as follows the geodesic fold ray $\mathfrak{R}$ that we later prove is dense. Let x_0 be a rose in the unprojectivized base simplex (i.e. having the identity marking) with the length vector w_0 provided by Lemma 3.16 for the matrix sequence $\{D_l\}$ defined in the paragraph above. For each k, let x_k be the rose with length vector inductively defined as $w_k = Z_k w_{k-1}$ and with the marking $f_k \circ \dots \circ f_1$. For each l, if Z_l is a single fold matrix coming from the matrix decomposition of A_i^n, let the line from x_l to x_{l+1} be the single proper full fold fold line corresponding to Z_l. This fold is allowable since $Z_l w_l = w_{l+1}$ is positive.

If $Z_l = H_{i,j}$ for some i, j coming from $a_k = (i, j)$, then $a_{k+1} = (i, n)$ for $n > n(i, j)$. Hence, there is a number s so that the following s matrices $\{D_d\}_{d=l+1}^{l+s}$ are the matrices of the decomposition of A_i^n. Therefore, $w_{l+1} = A_i^n(w_{l+s})$ for $n > n(i, j)$. Thus $x_{k+1} \in U_{ij}$, so the fold line corresponding to $\mathcal{R}_{ij}$ is allowable in $x_k = H_{ij}(x_{k+1})$. We insert this fold line between such x_k and x_{k+1}. This defines a fold ray connecting the x_k's.

Theorem B. *For each $r \geq 2$, there exists a geodesic ray $\tilde{\gamma}\colon [0, \infty) \to \mathcal{RX}_r$ so that the projection of $\tilde{\gamma}$ to $\mathcal{URX}_r/\operatorname{Out}(F_r)$ is dense.*

Proof. We will show that, for each $r \geq 2$, the fold ray $\mathfrak{R}$ of Definition 8.1 is contained in $\mathcal{URX}_r$ and projects densely into $\mathcal{U}_r$. The ray is geodesic by Corollary 3.19.

To prove that $\mathfrak{R}$ never leaves $\mathcal{RX}_r$, i.e. contains no graph with a separating edge, it will suffice to show that at each point $x \in \mathfrak{R}$, the underlying graph can be directed so that it is a transitive graph. This clearly holds for each proper full fold of a rose, hence for the fold sequences coming from the decompositions of the $g_{v_i}^k$. Moreover, for each i, j we have that $\mathcal{R}_{ij}$ consists of transitive graphs, since all folds are direction matching, see Observation 6.17(2).

Let x be a point in an unprojectivized top-dimensional simplex with rationally independent edge lengths. Let G be its underlying graph and $\{E, E'\}$ a turn. By Proposition 7.1 we can construct a proper positive rose–to–rose fold line $\mathcal{R} = \mathcal{R}(x_0, \vec{s})$ containing x and the combinatorial fold f_1 of $\{E, E'\}$ directly after x. We may assume that the terminal point of this rose–to–rose fold line z lies in σ_0. For each $\varepsilon > 0$, by Proposition 7.7, there exists a proper *rational* rose–to–rose fold line $\mathcal{R}'$ containing a point x' in the same unprojectivized open simplex as x, so that x, x' are ε-close, and so that the fold f_1 is the fold following x'. Let H be the unfolding matrix corresponding to $\mathcal{R}'$.

For each $\varepsilon > 0$, there exists, as in Lemma 7.4, an open neighborhood U of the terminal point of $\mathcal{R}'$ so that for each $w \in U$, there exists a proper rose–to–rose fold line $\mathcal{R}'(y_0, \vec{u})$, terminating at w and such that the top graphs x', y are ε-close, the combinatorial fold sequence in the graph–to–rose segments are the same, and the change-of-metric matrix for $\mathcal{R}'(y_0, \vec{u})$ is H. Since the set of PF eigenvectors is dense, there exists an i so that the PF eigenvector v_i is contained in U. Hence, there exists a rose–to–rose fold line R_{ij} passing through the same unprojectivized simplices as $\mathcal{R}'$ and having the same change-of-metric matrix $H_{ij} = H$. Moreover, there exists a point $x''_{ij} \in R_{ij}$ in the same unprojectivized top-dimensional simplex as x' and ε-close to x'. By Definition 8.1, there exist infinitely many k's so that the fold line between x_k and x_{k+1} is the one passing through the same unprojectivized simplices as R_{ij} (hence $\mathcal{R}'$). In fact, these occur before arbitrarily high powers of g_{v_i}, so that they terminate arbitrarily close to a rose with length vector v_i. Let k be such a number and let Ψ_k be the composition of the automorphisms $f_1, f_2, \ldots$ up to x_k. Thus, by Lemma 7.4, there exists a point $\xi \in [x_{k-1}, x_k] \cdot \Psi_k^{-1}$ in the same unprojectivized top-dimensional simplex as x''_{ij} and ε-close to x''_{ij}. Hence, ξ is the point on the ray defined in Definition 8.1 that is 3ε-close to our original point x and the fold immediately after ξ is the one folding the turn $\{E, E'\}$. □

Theorem A. *For each $r \geq 2$, there exists a geodesic fold ray in the reduced Outer Space $\mathcal{RX}_r$ whose projection to $\mathcal{RX}_r / \mathrm{Out}(F_r)$ is dense.*

Proof. This is an immediate corollary of Theorem B. □

9. Appendix: Limits of fold geodesics

In many cases, as in the case of the geodesic that we construct in this paper, a concatenation of fold segments $\{\gamma_i : [i, i+1] \to \widehat{\mathcal{X}}_r\}_{i=1}^{\infty}$ that glue together to a ray $\gamma : [0, \infty) \to \widehat{\mathcal{X}}_r$, projecting under q to a Lipschitz geodesic, satisfies the properties of a semi-flow line below.

Definition 9.1 (semi-flow line; cf. [13, p. 3], definition of a "fold line"). A *semi-flow line* in unprojectivized Outer Space is a continuous, injective, proper function $\mathbf{R} \to \widehat{\mathcal{X}}_r$ defined by a continuous 1-parameter family of marked graphs $t \to G_t$ for which there exists a family of homotopy equivalences $h_{ts} : G_s \to G_t$ defined for $s \leq t \in \mathbf{R}$, each of which preserves marking, such that the following hold.

(1) *Train track property.* For all $s \leq t \in \mathbf{R}$, the restriction of h_{ts} to the interior of each edge of G_s is locally an isometric embedding.

(2) *Semiflow property.* $h_{ut} \circ h_{ts} = h_{us}$ for all $s \leq t \leq u \in \mathbf{R}$ and $h_{ss} : G_s \to G_s$ is the identity for all $s \in \mathbf{R}$.

Handel and Mosher ([13, §7.3]) prove each semi-flow line converges (in the axes or Gromov-Hausdorff topologies) to a point T_∞ in $\overline{\mathcal{X}_r}$, the direct limit of the system.

Theorem 9.2. *For any semi-flow line, its direct limit T_∞ is an F_r-tree that has trivial arc stabilizers. Hence, in particular, not every point in the boundary is the direct limit of a semi-flow line.*

Remark 9.3. This theorem could also be deduced from the proof of Proposition 3.15 of [2]. It is shown there that if $\{G_i\}_{i=1}^{\infty} \subset \widehat{\mathcal{X}}_r$, $\{\widetilde{G}_i\}$ limits to T as an $\mathbf{R}$-tree and T has a nontrivial arc-stabilizer, then $\mathrm{qvol}(T) > \mathrm{qvol}(\widetilde{G}_i) = \mathrm{vol}(G_i)$ for each i. However, if G_{i+1} is obtained from G_i by folding, then $\mathrm{vol}(G_{i+1}) < \mathrm{vol}(G_i)$, which is a contradiction. In order to avoid the definitions above we give a more direct proof.

Proof. We lift the maps $h_{ts}: G_s \to G_t$ to $f_{ts}: T_s \to T_t$ to a direct system on trees $\{f_{ts}: T_s \to T_t\}$ that are F_r-equivariant, restrict to isometric embeddings on edges, and form a direct system. In [13, §7.3] it is shown that the maps $f_{\infty s}: T_s \to T_\infty$ that are given by the direct limit construction are also edge-isometries and F_r-equivariant.

Assume $x, y \in T_\infty$ are such that $\gamma \in \mathrm{Stab}[x, y]$. Without generality loss, assume $d(x, y) = 1$ in T_∞. We will show this leads to a contradiction. For each $t \in [a, \infty)$, we denote by $A_t(\gamma)$ the axis of γ in T_t. Letting $\varepsilon = \frac{1}{5}$, there exists some $s \geq a$ so that $f_{\infty s}(x_s) = x$, $f_{\infty s}(y_s) = y$, and $d(x_s, \gamma x_s) < \varepsilon$, $\quad d(y_s, \gamma y_s) < \varepsilon$, and $|d(x_s, y_s) - 1| < \varepsilon$.

Since f_{ts} is distance non-increasing for all t, s, we have for all $t \geq s$ that

$$1 \leq d(f_{ts}(x_s), f_{ts}(y_s)) \leq 1 + \varepsilon.$$

Note that for all $f_{ts}(x_s)$ and $f_{ts}(y_s)$, they are at most a distance of $\frac{\varepsilon}{2}$ from $A_t(\gamma)$. Otherwise, for example, $d(f_{ts}(x_s), \gamma f_{ts}(x_s)) \geq \varepsilon$, which contradicts the fact that the maps are distance non-increasing.

Thus, for each $t \geq s$ there exist z_t, w_t such that $[z_t, w_t] \subset A_t(\gamma) \cap [f_{ts}(x_s)), f_{ts}(y_s))]$, and $d(z_t, f_{ts}(x_s)) < \frac{\varepsilon}{2}$ and $d(w_t, f_{ts}(y_s)) < \frac{\varepsilon}{2}$. Hence, $d(z_t, w_t) \geq 1 - \varepsilon$.

Let N be the number of $f_{\infty s}$-illegal turns in the path $[x_s, y_s]$. Thus, the number of $f_{\infty t}$-illegal turns in the path $[f_{ts}(x_s), f_{ts}(y_s)]$ is $\leq N$. Hence, the number of $f_{\infty t}$-illegal turns in the path $[z_t, w_t]$ is $\leq N$. Let us denote the points of $[z_t, w_t]$ where the illegal turns occur by $a_t^1, \dots a_t^N$, and we also denote $a_t^0 := z_t$ and $a_t^{N+1} := w_t$.

However, for sufficiently large t, the translation length of γ in T_t is $< \frac{1-\varepsilon}{3(N+1)}$. Note that since $[z_t, w_t]$ is on $A_t(\gamma)$, we have that $[z_t, w_t] \cap \gamma[z_t, w_t]$ is equal to $[z_t, w_t]$ with (possibly) segments of lengths $\leq \frac{1-\varepsilon}{3(N+1)}$ cut from either end. Since there are $\leq N + 1$ segments (a_t^i, a_t^{i+1}) in $[z_t, w_t]$, one of them has length $\geq \frac{1-\varepsilon}{N+1}$. Thus, for some $i = 0, 1, \dots, N$, we have $\gamma(a_t^i) \in (a_t^i, a_t^{i+1})$ or $\gamma^{-1}(a_t^{i+1}) \in (a_t^i, a_t^{i+1})$, which are both legal segments. Without generality loss suppose the former. Thus, the turn taken by $[z_t, w_t]$ at a_t^i is illegal, but the turn taken by $\gamma[z_t, w_t]$ at $\gamma(a_t^i)$ is legal, since it equals the turn taken by $[z_t, w_t]$ at $\gamma(a_t^i)$. This is a contradiction to the equivariance. □

References

[1] Y. Algom-Kfir, Strongly contracting geodesics in outer space. *Geom. Topol.* **15** (2011), no. 4, 2181–2233. Zbl 1250.20019 MR 2862155

[2] Y. Algom-Kfir, The metric completion of Outer Space. Preprint 2012. arXiv:1202.6392 [math.GR]

[3] M. Bestvina, A Bers-like proof of the existence of train tracks for free group automorphisms. *Fund. Math.* **214** (2011), no. 1, 1–12. Zbl 1248.20025 MR 2845630

[4] M. Bestvina, PCMI Lectures on the geometry of Outer space, 2012. http://www.math.utah.edu/pcmi12/lecture_notes/bestvina.pdf

[5] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups. *Ann. of Math.* (2) **135** (1992), no. 1, 1–51. Zbl 0757.57004 MR 1147956

[6] G. Birkhoff, Extensions of Jentzsch's theorem. *Trans. Amer. Math. Soc.* **85** (1957), 219–227. Zbl 0079.13502 MR 0087058

[7] V. Brun, Algorithmes euclidiens pour trois et quatre nombres. In *Treizième congrès des mathèmaticiens scandinaves.* (Helsinki, 1957.) Mercators Tryckeri, Helsinki, 1957, 45–64. Zbl 0086.03204 MR 0111735

[8] M. Culler and K. Vogtmann, Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84** (1986), no. 1, 91–119. Zbl 0589.20022 MR 0830040

[9] R. Diestel, *Graph theory.* 3$^{\text{rd}}$ revised and updated edition. Graduate Texts in Mathematics, 173. Springer-Verlag, Berlin, 2005. Zbl 1074.05001 MR 2159259

[10] A. Eskin and M. Mirzakhani, Counting closed geodesics in moduli space. *J. Mod. Dyn.* **5** (2011), no. 1, 71–105. Zbl 1219.37006 MR 2787598

[11] A. Fisher, Nonstationary mixing and the unique ergodicity of adic transformations. *Stoch. Dyn.* **9** (2009), no. 3, 335–391. Zbl 1231.28014 MR 2566907

[12] S. Francaviglia and A. Martino, Metric properties of outer space. *Publ. Mat.* **55** (2011), no. 2, 433–473. Zbl 1268.20042 MR 2839451

[13] M. Handel and L. Mosher, Axes in Outer Space. *Mem. Amer. Math. Soc.* **213** (2011), no. 1004. Zbl 1238.57002 MR 2858636

[14] M. Keane, Non-ergodic interval exchange transformations. *Israel J. Math.* **26** (1977), no. 2, 188–196. Zbl 0351.28012 MR 0435353

[15] H. B. Keynes and D. Newton, A "minimal", non-uniquely ergodic interval exchange transformation. *Math. Z.* **148** (1976), no. 2, 101–105. Zbl 0308.28014 MR 0409766

[16] H. Masur, Dense geodesics in moduli space. In I. Kra and B. Maskit (eds.), *Riemann surfaces and related topics.* (Stony Brook, N.Y., 1978.) Annals of Mathematics Studies, 97. Princeton University Press, Princeton, N.J., and University of Tokyo Press, Tokyo, 1981, 417–438. Zbl 0476.32027 MR 0624830

[17] H. Masur, Interval exchange transformations and measured foliations. *Ann. of Math.* (2) **115** (1982), no. 1, 169–200. Zbl 0497.28012 MR 0644018

[18] F. Schweiger, Ergodische Eigenschaften der Algorithmen von Brun und Selmer. *Österreich. Akad. Wiss. Math.-Natur. Kl. Sitzungsber.* II **191** (1982), no. 8-9, 325–329. Zbl 0518.10060 MR 0716668

[19] F. Schweiger, *Multidimensional continued fractions.* Oxford Science Publications. Oxford University Press, Oxford, 2000. Zbl 0981.11029 MR 2121855

[20] W. Thurston, Minimal stretch maps between hyperbolic surfaces. Preprint 1998. arXiv:math.GT/9801039 [math.GT]

[21] W. A. Veech, Gauss measures for transformations on the space of interval exchange maps. *Ann. of Math.* (2) **115** (1982), no. 1, 201–242. Zbl 0486.28014 MR 0644019

[22] K. Vogtmann, Automorphisms of free groups and outer space. *Geom. Dedicata* **94** (2002), 1–31. Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000). Zbl 1017.20035 MR 1950871

Received August 17, 2016

Yael Algom-Kfir, Department of Mathematics, University of Haifa, Mount Carmel, Haifa, 31905, Israel

home page: http://www.math.haifa.ac.il/algomkfir/

e-mail: yalgom@univ.haifa.ac.il

Catherine Pfaff, University of California at Santa Barbara, South Hall, Room 6607, Santa Barbara, CA 93106-3080, USA

home page: http://math.ucsb.edu/~cpfaff/

e-mail: cpfaff@math.ucsb.edu