STABILITY IN OUTER SPACE

URSULA HAMENSTÄDT AND SEBASTIAN HENSEL

Abstract. We characterize strongly Morse quasi-geodesics in Outer space as quasi-geodesics which project to quasi-geodesics in the free factor graph. We define convex cocompact subgroups of $\mathrm{Out}(F_n)$ as subgroups such that an orbit map in the free factor graph is a quasi-isometric embedding, and we characterize such groups via their action on Outer space in a way which resembles the characterization of convex cocompact subgroups of mapping class groups.

1. Introduction

A quasi-geodesic γ in a geodesic metric space X is called *stable* if for every $L \geq 1$ there exists some $M(L) > 0$ such that any L-quasi-geodesic η with endpoints on γ is contained in the $M(L)$-neighborhood of γ.

In a hyperbolic geodesic metric space, any quasi-geodesic is stable in this sense, but many geodesic metric spaces which are not hyperbolic contain stable quasi-geodesics as well. We refer to [CS15] for examples.

Often the existence of stable quasi-geodesics in a space X is related to hyperbolicity of some graph associated to X, in particular in the presence of a large isometry group. In the case of the *Teichmüller space* $\mathcal{T}(S)$ of a surface S of finite type equipped with the Teichmüller metric, stability is characterized as follows. A quasi-geodesic is stable if and only if it projects to a quasi-geodesic in the *curve graph* $\mathcal{C}(S)$ of S [H10]. The curve graph is a hyperbolic geodesic metric graph equipped with an action of the *mapping class group.*

In the attempt to find similarities between the geometry of the mapping class group of a surface of finite type and the geometry of the *Outer automorphism group* $\mathrm{Out}(F_n)$ of a free group F_n with $n \geq 3$ generators, two natural hyperbolic $\mathrm{Out}(F_n)$-graphs were found.

The so-called *free factor graph* $\mathcal{FF}$ is the metric graph whose vertices are conjugacy classes of non-trivial free factors of F_n. Two vertices A, B are connected by an edge of length one if up to conjugation, either $A < B$ or $B < A$. It is a hyperbolic $\mathrm{Out}(F_n)$-graph [BF14]. There is also the *free splitting graph* which however is not relevant for this work.

Date: August 9, 2017.
AMS subject classification: 20M34
Both authors were partially supported by ERC grant "Moduli".

Our first goal is to show that the free factor graph characterizes stability in *Outer space* $cv_0(F_n)$ equipped with the *symmetrized Lipschitz metric* d in the same way the curve graph characterizes stability of Teichmüller space with the Teichmüller metric. Here Outer space is the space of all minimal free actions of F_n on simplicial trees T with quotient T/F_n of volume one. It is equipped with the symmetrized Lipschitz metric d which is invariant under the action of $\mathrm{Out}(F_n)$. There also is a natural coarsely well defined projection $\Upsilon : cv_0(F_n) \to \mathcal{FF}$.

The symmetrized Lipschitz metric on Outer space is not geodesic and therefore we will work instead with coarse geodesics. By definition, a *c-coarse geodesic* is a path $\gamma : \mathbb{R} \to cv_0(F_n)$ such that for all $s, t \in \mathbb{R}$ we have

$$|s - t| - c \leq d(\gamma(s), \gamma(t)) \leq |s - t| + c.$$

Note that a coarse geodesic need not be continuous.

Let d_L be the *one-sided Lipschitz metric* on $cv_0(F_n)$. By definition, $d_L(T, T')$ is the infimum of the logarithms of the Lipschitz constants among all marked homotopy equivalences $f : T/F_n \to T'/F_n$. For a number $K > 0$, a *K-quasi-geodesic* for d_L is a path $\gamma : [a, b] \to cv_0(F_n)$ so that for all $s < t$ we have

$$|t - s|/K - K \leq d_L(\gamma(s), \gamma(t)) \leq K|t - s| + K.$$

Definition 1. A coarse geodesic $\gamma \subset (cv_0(F_n), d)$ is called *strongly Morse* if for any $K \geq 1$ there is a number $M = M(K) > 0$ with the following property. Any K-quasi-geodesic for d_L or d with endpoints on γ is contained in the M-neighborhood of γ for the symmetrised Lipschitz metric.

In the sequel we talk about uniformly strongly Morse coarse geodesics to denote a collection of paths which fulfill the conditions in Definition 1 with the same constants. Similarly, we talk about collections of uniform quasi-geodesics.

For a number $\epsilon > 0$ let $\mathrm{Thick}_\epsilon(F_n) \subset cv_0(F_n)$ be the subspace of all F_n-trees whose volume one quotient does not contain a nontrivial loop of length less than ϵ. We show

Theorem 1. *Let $\gamma : (a, b) \to \mathrm{Thick}_\epsilon(F_n)$ be a c-coarse geodesic for the symmetrized Lipschitz metric. Then the following are equivalent.*

(1) The path $t \to \Upsilon(\gamma(t))$ is a uniform quasi-geodesic in $\mathcal{FF}$.
(2) γ is uniformly strongly Morse.

The precise meaning of the implication (1) ⇒ (2) in Theorem 1 is as follows. For all $c > 0, L > 1, K > 0$ there exists a constant $M = M(c, L, K) > 0$ with the following property. Let $\gamma : (a, b) \to cv_0(F_n)$ be a c-coarse geodesic for the symmetrized Lipschitz metric whose composition with Υ is an L-quasi-geodesic. Then any K-quasi-geodesic for the one sided Lipschitz metric or the symmetrized Lipschitz metric with endpoints on γ is contained in the M-neighborhood of γ with respect to the symmetrized Lipschitz metric.

Motivated by the theory of Kleinian groups, Farb and Mosher define in [FM02] the notion of a *convex cocompact subgroup* Γ of the mapping class group $\mathrm{Mod}(S)$ of a surface S of genus $g \geq 2$ via geometric properties of the action of Γ on Teichmüller space $\mathcal{T}(S)$ for S. Later, an equivalent characterization using the action of Γ on the curve graph of S was established in [H05] and [KeL08]. Our second goal is to develop a similar theory for subgroups of the outer automorphism group of a free group. We briefly recall the results in the mapping class group case. To this end denote by $\partial\mathcal{T}(S)$ the Thurston boundary of Teichmüller space.

Theorem 2 ([FM02, H05, KeL08])**.** *The following properties of a finitely generated subgroup Γ of $\mathrm{Mod}(S)$ are equivalent.*

(1) The orbit map on the curve graph $\mathcal{C}(S)$ of S is a quasi-isometric embedding.
(2) A Γ-orbit on $\mathcal{T}(S)$ is quasi-convex: For any $x \in \mathcal{T}(S)$ and all $g, h \in \Gamma$, the Teichmüller geodesic connecting gx, hx is contained in a uniformly bounded neighborhood of Γx.

Moreover, the group Γ is hyperbolic, and there is an equivariant homeomorphism $F : \partial\Gamma \to \Lambda(\Gamma)$ where $\partial\Gamma$ is the Gromov boundary of Γ and where $\Lambda(\Gamma) \subset \partial\mathcal{T}(S)$ is the set of accumulation points of a Γ-orbit on $\mathcal{T}(S)$.

Part (2) of Theorem 2 shows that such so-called convex cocompact subgroups of the mapping class group are precisely those groups which act on Teichmüller space equipped with the Teichmüller metric as convex cocompact groups: This is a characterization not by intrinsic properties of the group, but via its isometric action on Teichmüller space, a bounded domain in $\mathbb{C}^{3g-3}$. It then turns out that there is an equivalent characterization via the action of the group on the curve graph. An intrinsic characterization of convex cocompact subgroups of $\mathrm{Mod}(S)$ is due to Durham and Taylor [DuT15].

Definition 2. A finitely generated subgroup Γ of $\mathrm{Out}(F_n)$ is *convex cocompact* if one (and hence every) orbit map of its action on the free factor graph is a quasi-isometric embedding.

The following result is the analog of Theorem 2 for $\mathrm{Out}(F_n)$. For its formulation, we denote by $\mathrm{CV}(F_n)$ the projectivization of $cv_0(F_n)$ with its boundary $\partial\mathrm{CV}(F_n)$. There is a natural topology on $\overline{\mathrm{CV}(F_n)} = \mathrm{CV}(F_n) \cup \partial\mathrm{CV}(F_n)$ such that $\overline{\mathrm{CV}(F_n)}$ is compact and that the restriction of this topology to the open dense subset $\mathrm{CV}(F_n)$ is the topology inherited from $cv_0(F_n)$.

Theorem 3. *Let Γ be a finitely generated subgroup of $\mathrm{Out}(F_n)$. Then the following are equivalent.*

(1) Γ is convex cocompact.
(2) Let $T \in cv_0(F_n)$. Then for all $g, h \in \Gamma$, the points gT, hT are connected by a uniformly strongly Morse c-coarse geodesic which is contained in a uniformly bounded neighborhood of ΓT.

Moreover, the group Γ is hyperbolic, and there is an equivariant homeomorphism $F : \partial\Gamma \to \Lambda(\Gamma)$ where $\partial\Gamma$ is the Gromov boundary of Γ and where $\Lambda(\Gamma) \subset \partial\mathrm{CV}(F_n)$ is the set of accumulation points of a Γ-orbit on $\mathrm{CV}(F_n)$.

Corollary 5.3 contains another characterization of convex cocompact subgroups of $\mathrm{Out}(F_n)$ in the spirit of the work of Farb and Mosher [FM02]. This description uses lines of minima as defined in [H14] and is a bit more difficult to explain. However, it is the characterization of such groups which contains the most information.

Examples of convex cocompact groups are Schottky groups. In the case $n = 2g$ for some $g \geq 2$, convex cocompact subgroups of the mapping class group of a surface of genus g with one puncture, viewed as subgroups of $\mathrm{Out}(F_n)$, are convex cocompact. We discuss these examples in Section 7.

For mapping class groups of a closed surface S, there is yet another characterization of convex cocompact subgroups [FM02, H05]. Namely, Γ is convex cocompact if and only if the extension G of Γ given by the exact sequence

$$0 \to \pi_1(S) \to G \to \Gamma \to 0$$

is word hyperbolic.

In contrast, the F_n-extension of a convex cocompact subgroup of $\mathrm{Out}(F_n)$ need not be word hyperbolic. As an example, the extension of a convex cocompact subgroup of the mapping class group of a surface with a puncture is not hyperbolic. However, Dowdall and Taylor [DT14] showed that the F_n-extension of a convex cocompact subgroup Γ all of whose elements are non-geometric is hyperbolic. Moreover, they characterize hyperbolic F_n-extensions via the action of $\mathrm{Out}(F_n)$ on the so-called *cosurface graph* [DT16].

There is overlap of our work with the work of Dowdall and Taylor [DT14, DT15]. However, their arguments are very different from the arguments we use, and their main goal is different in flavor as well. There is also some overlap with the work [NPR14] whose main result is a key ingredient in the proof of Theorem 2.

Organization: Section 2 collects the basic tools and background. In Section 3 we relate lines of minima as introduced in [H14] to coarse geodesics in the thick part of Outer space whose shadows in the free factor graph are quasi-geodesics. Section 4 investigates strongly Morse coarse geodesics in Outer space which leads to the proof of the implication (2) $\Rightarrow$ (1) in Theorem 1. Section 5 uses lines of minima to show that a subgroup of $\mathrm{Out}(F_n)$ which has the properties stated in the second part of Theorem 3 is convex cocompact. In Section 6, the proof of Theorem 3 and the implication (1) $\Rightarrow$ (2) in Theorem 1 is completed. In Section 7 we discuss examples of convex cocompact subgroups of $\mathrm{Out}(F_n)$. The appendix collects some results from the work [NPR14] which are used in Section 4.

Acknowledgment: We thank the anonymous referee for useful suggestions that improved the article. We are also grateful to Spencer Dowdall and Samuel Taylor for drawing our attention to the work [DT14].

2. Geometric tools

2.1. **The boundary of the free factor graph.** The free factor graph $\mathcal{FF}$ is hyperbolic [BF14]. Its Gromov boundary $\partial\mathcal{FF}$ can be described as follows [BR15, H12].

Unprojectivized Outer space $cv(F_n)$ of simplicial minimal free F_n-trees equipped with the equivariant Gromov Hausdorff topology can be completed by attaching a boundary $\partial cv(F_n)$. This boundary consists of all minimal *very small* actions of F_n on $\mathbb{R}$-trees which either are not simplicial or which are not free [CL95, BF92]. Here an F_n-tree is very small if arc stabilizers are at most maximal cyclic and tripod stabilizers are trivial. We denote by $\mathrm{CV}(F_n)$ the projectivization of $cv(F_n)$, with its boundary $\partial\mathrm{CV}(F_n)$. Also, from now on we always denote by $[T] \in \overline{\mathrm{CV}(F_n)} = \mathrm{CV}(F_n) \cup \partial\mathrm{CV}(F_n)$ the projectivization of a tree $T \in \overline{cv(F_n)} = cv(F_n) \cup \partial cv(F_n)$.

Let ∂F_n be the Gromov boundary of F_n and denote by Δ the diagonal in $\partial F_n \times \partial F_n$. An element $u \in F_n$ is called *primitive* if it can be completed to a free basis of F_n. The set of Dirac measures on pairs of fixed points of all elements in some primitive conjugacy class α of F_n is a locally finite F_n-invariant Borel measure on $\partial F_n \times \partial F_n - \Delta$ which we call *dual* to α. The closure of all such measures in the space of all locally finite F_n-invariant Borel measures on $\partial F_n \times \partial F_n - \Delta$, equipped with the weak*-topology, is the space $\mathcal{ML}$ of *measured laminations* for F_n. It is invariant under the natural action of $\mathrm{Out}(F_n)$. The projectivization $\mathcal{PML}$ of $\mathcal{ML}$ is compact, and $\mathrm{Out}(F_n)$ acts on $\mathcal{PML}$ minimally by homeomorphisms [Ma95]. In the sequel we always denote by $[\mu] \in \mathcal{PML}$ the projectivization of a measured lamination $\mu \in \mathcal{ML}$.

By [KL09], there is a continuous length pairing

$$\langle , \rangle : \overline{cv(F_n)} \times \mathcal{ML} \to [0, \infty).$$

If $\xi \in \mathcal{ML}$ is dual to a primitive conjugacy class α in F_n and if $T \in cv(F_n)$ then $\langle T, \xi \rangle$ equals the shortest length of a representative of α on T/F_n. If $\mu \in \mathcal{ML}$ is arbitrary then $\langle T, \mu \rangle > 0$ for every tree $T \in cv(F_n)$.

Definition 2.1. A measured lamination $\mu \in \mathcal{ML}$ is *dual* to a tree $T \in \partial cv(F_n)$ if $\langle T, \mu \rangle = 0$.

Note that if μ is dual to T then any multiple of μ is dual to every tree obtained from T by scaling, so we can talk about a projective measured lamination which is dual to a projective tree. We note for later reference (see [H12])

Lemma 2.2. *Every projective tree $[T] \in \partial\mathrm{CV}(F_n)$ admits a dual measured lamination.*

We say that a measured lamination μ is *supported in a free factor H* of F_n if the support of μ is contained in the F_n-orbit of $\partial H \times \partial H - \Delta$. If $[T]$ has point stabilizers containing a free factor, then any measured lamination supported in the free factor is dual to T. If $[T] \in \partial\mathrm{CV}(F_n)$ is simplicial then the set of measured laminations dual to $[T]$ consists of convex combinations of measured laminations supported in a point stabilizer of $[T]$.

A (projective) tree $[T] \in \partial \mathrm{CV}(F_n)$ is called *indecomposable* if for any non-degenerate segments $I, J \subset T$ there are elements $u_1, \dots, u_n \in F_n$ with $I \subset u_1 J \cup \dots \cup u_n J$ and so that $u_i J \cap u_{i+1} J$ is a non-degenerate segment for all i.

Let $\sim$ be the smallest equivalence relation on $\partial \mathrm{CV}(F_n)$ with the following property. For a projective tree $[T] \in \partial \mathrm{CV}(F_n)$ and a projective measured lamination $[\mu] \in \mathcal{PML}$ dual to $[T]$, any tree $[S] \in \partial \mathrm{CV}(F_n)$ dual to $[\mu]$ is equivalent to $[T]$.

Theorem 2.3 ([BR15, H12])**.** *The Gromov boundary $\partial \mathcal{FF}$ of $\mathcal{FF}$ can be identified with the set of equivalence classes under $\sim$ of indecomposable projective trees $[T]$ with the following additional property. Either the F_n-action on T is free, or there is a compact surface S with non-empty connected boundary, and there is a minimal filling measured lamination μ on S so that T is dual to μ.*

In the sequel we call a (projective) tree with the properties stated in the theorem *arational*. We refer to the main result of [R12] for a characterization of arational trees which justifies this terminology.

At this point, we record a criterion for arationality that will be used later. To state it, we need the following definition. An *alignment preserving map* between two F_n-trees $T, T' \in \overline{cv(F_n)}$ is defined to be an equivariant map $\rho : T \to T'$ with the property that $x \in [y, z]$ implies $\rho(x) \in [\rho(y), \rho(z)]$. The map ρ is continuous on segments.

Lemma 2.4. *Let $[T] \in \partial \mathrm{CV}(F_n)$ be given. Suppose that T does not have point stabilizers containing free factors, and that there is no tree $T' \in \partial cv(F_n)$ which can be obtained from T by a one-Lipschitz alignment preserving map $\rho : T \to T'$ collapsing a nontrivial subtree of T to a point. Then $[T]$ is arational.*

Proof. By the results in Section 10 of [R12], a tree T which satisfies the assumption in the lemma on non-existence of interesting alignment preserving maps is indecomposable. By Proposition 10.1 of [H12], such a tree T is arational if it does not have point stabilisers containing nontrivial free factors (compare also [BR15]), proving the lemma. □

By continuity of the length pairing, the set of all projective trees $[S]$ which are dual to some fixed measured lamination μ is a closed subset of $\partial \mathrm{CV}(F_n)$. The topology on $\partial \mathcal{FF}$ is the quotient topology for the closed equivalence relation $\sim$ on the set of arational projective trees. It can be described as follows. A sequence of equivalence classes represented by trees S_i converges to the equivalence class represented by S if there is a sequence $(\nu_i) \subset \mathcal{ML}$ so that $\langle S_i, \nu_i \rangle = 0$ for all i and such that $\nu_i \to \nu$ in $\mathcal{ML}$ with $\langle S, \nu \rangle = 0$ [H12].

Definition 2.5. A pair of measured laminations $(\mu, \nu) \in \mathcal{ML} \times \mathcal{ML}$ is called a *positive pair* if for any tree $S \in \overline{cv(F_n)}$ we have $\langle S, \mu + \nu \rangle > 0$.

Positivity of a pair is invariant under scaling each individual component by a positive factor, so it is defined for pairs of projective measured laminations.

Lemma 2.6. *The set of positive pairs is an open subset of $\mathcal{PML} \times \mathcal{PML}$.*

Proof. By invariance under scaling, it suffices to show the following. Let $A \subset \overline{cv(F_n)}$ be a compact set which projects onto $\overline{CV(F_n)}$, and let $B \subset \mathcal{ML}$ be a compact set which projects onto $\mathcal{PML}$. Let $(\mu, \nu) \in B \times B$ be a positive pair; then there exists a neighborhood U of (μ, ν) in $B \times B$ such that $\langle S, \zeta + \xi \rangle > 0$ for all $(\zeta, \xi) \in U$ and all $S \in A$. However, this is immediate from continuity of the length pairing and compactness of A. $\square$

Corollary 10.6 of [H12] identifies positive pairs which are of particular significance for our purpose.

Lemma 2.7. *Let $[T] \neq [T']$ be arational trees which define distinct points in $\partial\mathcal{FF}$. Let $\mu, \mu' \in \mathcal{ML}$ be dual to $[T], [T']$; then (μ, μ') is a positive pair.*

2.2. **Lines of minima.** In this subsection we introduce the central tool used in this paper: *lines of minima* as defined in [H14].

For $\epsilon > 0$ define

$$\mathrm{Thick}_\epsilon(F_n)$$

to be the set of all trees $T \in cv_0(F_n)$ with volume one quotient so that the shortest length of any loop on T/F_n is at least ϵ. For the remainder of this paper we always choose $\epsilon > 0$ sufficiently small that $\mathrm{Thick}_\epsilon(F_n)$ is non-empty and path connected. Clearly $\mathrm{Thick}_\epsilon(F_n)$ is invariant under the action of $\mathrm{Out}(F_n)$.

For a tree $T \in cv_0(F_n)$ define

$$\Lambda(T) = \{\mu \in \mathcal{ML} \mid \langle T, \mu \rangle = 1\}. \tag{1}$$

Then $\Lambda(T)$ is a compact subset of $\mathcal{ML}$, and the projection $\Lambda(T) \to \mathcal{PML}$ is a homeomorphism. Let moreover

$$\Sigma(T) = \{S \in cv(F_n) \mid \sup\{\langle S, \mu \rangle \mid \mu \in \Lambda(T)\} = 1\}. \tag{2}$$

Let $(\mu, \nu) \in \mathcal{ML}^2$ be a positive pair. By Lemma 3.2 of [H14], the function $S \to \langle S, \mu + \nu \rangle$ on $\mathrm{Thick}_\epsilon(F_n)$ is proper. This means that this function assumes a minimum, and the set

$$\mathrm{Min}_\epsilon(\mu + \nu) = \{T \in \mathrm{Thick}_\epsilon(F_n) \mid \langle T, \mu + \nu \rangle = \min\{\langle S, \mu + \nu \rangle \mid S \in \mathrm{Thick}_\epsilon(F_n)\}\}$$

of all such minima is compact. Note that this set does not change if we replace $\mu + \nu$ by a positive multiple.

Definition 2.8. Let $([\mu], [\nu]) \in \mathcal{PML} \times \mathcal{PML} - \Delta$ be a positive pair. A *line of minima* for $([\mu], [\nu])$ is a map $\gamma : \mathbb{R} \to \mathrm{Thick}_\epsilon(F_n)$ which associates to $t \in \mathbb{R}$ a point $\gamma(t) \in \mathrm{Min}_\epsilon(e^{t/2}\mu + e^{-t/2}\nu)$.

A line of minima is by no means unique. If the measured laminations μ, ν are dual to some primitive conjugacy class then it may be of finite diameter. Moreover, a line of minima is in general not continuous.

We next introduce a class of positive pairs which define line of minima with particularly nice properties. To this end define for a positive pair $(\mu, \nu) \in \mathcal{ML} \times \mathcal{ML}$ the set

$$\mathrm{Bal}(\mu, \nu) = \{T \mid \langle T, \mu \rangle = \langle T, \nu \rangle\} \subset cv(F_n)$$

of balancing trees.

Call a primitive conjugacy class α *basic* for $T \in cv_0(F_n)$ if α can be represented by a loop of length at most two on the quotient graph T/F_n. Note that any $T \in cv_0(F_n)$ admits a basic primitive conjugacy class.

Definition 2.9. For $B > 1$, a positive pair of points

$$([\mu],[\nu]) \in \mathcal{PML} \times \mathcal{PML} - \Delta$$

is called *B-contracting* if for any pair $\mu, \nu \in \mathcal{ML}$ of representatives of $[\mu],[\nu]$ there is some "distinguished" $T \in \mathrm{Min}_\epsilon(\mu+\nu)$ with the following properties.

(1) $\langle T,\mu\rangle/\langle T,\nu\rangle \in [B^{-1}, B]$.
(2) If $\tilde{\mu},\tilde{\nu} \in \Lambda(T)$ are representatives of $[\mu],[\nu]$ then $\langle S, \tilde{\mu}+\tilde{\nu}\rangle \geq 1/B$ for all $S \in \Sigma(T)$.
(3) Let $\mathcal{B}(T) \subset \Lambda(T)$ be the set of all normalized measured laminations which are up to scaling dual to a basic primitive conjugacy class for a tree $U \in \mathrm{Bal}(\mu,\nu)$. Then $\langle S,\xi\rangle \geq 1/B$ for every $\xi \in \mathcal{B}(T)$ and every tree
$$S \in \Sigma(T) \cap \left(\bigcup_{s \in (-\infty,-B)\cup(B,\infty)} \mathrm{Bal}(e^s\mu, e^{-s}\nu)\right).$$

Remark 2.10. The requirement in part 3) of the definition is slightly stronger than stated in [H14] as in [H14] it was assumed that the tree U is contained in $\mathrm{Thick}_\epsilon(F_n)$. We will establish below that this stronger property serves our needs.

Each B-contracting pair $(\mu,\nu) \in \mathcal{ML}\times\mathcal{ML}$ (i.e. such that the projectivized pair $([\mu],[\nu])$ is B-contracting in the sense of Definition 2.9) defines a *contracting line of minima* γ by associating to each $t \in \mathbb{R}$ a point $\gamma(t) \in \mathrm{Min}_\epsilon(e^{t/2}\mu + e^{-t/2}\nu)$ which fulfills the above definition. Such a contracting line of minima γ is a line of minima in the sense of Definition 2.8. It is not unique, but its Hausdorff distance (for the two-sided Lipschitz metric introduced below) to any other choice defined by any pair $(\hat{\mu},\hat{\nu}) \in \mathcal{ML}\times\mathcal{ML}$ with $[\hat{\mu}] = [\mu]$ and $[\hat{\nu}] = [\nu]$ is uniformly bounded (see [H14] for details). Note that a line of minima as introduced in [H14] is a contracting line of minima in the above sense. We hope that this discrepancy of terminology will not cause confusion.

Definition 2.11. Let (μ,ν) be a positive pair. The *balancing function* $f_{\mu,\nu}$: $\overline{\mathrm{CV}(F_n)} \to [-\infty,\infty]$ associates to a projective tree $[T]$ the unique number $t \in \mathbb{R}$ so that $\langle T, e^{t/2}\mu\rangle = \langle T, e^{-t/2}\nu\rangle$ if such a t exists, and it associates to $[T]$ the value ∞ (or $-\infty$) if $\langle T,\mu\rangle = 0$ (or $\langle T,\nu\rangle = 0$).

Lemma 2.12. *The balancing function $f_{\mu,\nu}$ of a positive pair (μ,ν) is continuous.*

Proof. Choose a compact subset of $\overline{cv(F_n)}$ which projects onto $\overline{CV(F_n)}$ and use continuity of the length function. □

Definition 2.13. Let (μ,ν) be a positive pair and γ an associated line of minima. We define the *balancing projection* $\Pi_\gamma : cv_0(F_n) \to \gamma$ by

$$\Pi_\gamma(T) = \gamma(f_{\mu,\nu}([T])).$$

The *one-sided Lipschitz metric* between two trees $S, T \in cv_0(F_n)$ is defined as

$$d_L(S,T) = \log \sup \left\{ \frac{\langle T, \nu \rangle}{\langle S, \nu \rangle} \mid \nu \in \mathcal{ML} \right\}. \tag{3}$$

The one-sided Lipschitz metric satisfies $d_L(S,T) = 0$ only if $S = T$, moreover it satisfies the triangle inequality, but it is not symmetric. The definition of the one-sided Lipschitz metric we give here is not standard, and we refer to [FM11] for a discussion why our definition is equivalent to the definition found in the introduction and in other articles.

Define the *symmetrized Lipschitz metric*

$$d(S,T) = d_L(S,T) + d_L(T,S).$$

Proposition 5.2 of [H14] shows the following.

Proposition 2.14. *For every $B > 0$ there is a number $\kappa = \kappa(B) > 0$ with the following property. Let $([\mu],[\nu])$ be a B-contracting pair, let γ be a contracting line of minima defined by $([\mu],[\nu])$ and let $U \in cv_0(F_n)$.*

(1) If $S \in cv_0(F_n)$ is such that $d(\Pi_\gamma(U), \Pi_\gamma(S)) \geq \kappa$ then
$$d_L(U,S) \geq d_L(U, \Pi_\gamma(U)) + d_L(\Pi_\gamma(U), \Pi_\gamma(S)) + d_L(\Pi_\gamma(S), S) - \kappa \text{ and}$$
$$d(U,S) \geq d(U, \Pi_\gamma(U)) + d(\Pi_\gamma(U), \Pi_\gamma(S)) + d(\Pi_\gamma(S), S) - \kappa.$$
(2) If $S \in \gamma(\mathbb{R})$ is such that $d(U,S) \leq \inf_t d(U, \gamma(t)) + 1$ then $d(S, \Pi_\gamma(U)) \leq \kappa$.
(3) For all $s < t$,
$$|s - t| - \kappa \leq d(\gamma(s), \gamma(t)) \leq |s - t| + \kappa.$$

Proof. Proposition 5.2 as stated in [H14] requires that the tree U is contained in $\mathrm{Thick}_\epsilon(F_n)$. The proof uses the axioms in the definition of a B-contracting pair and applies axiom (3) to U (compare the remark after Definition 2.9). No other assumption on U is used. Since we are using a stronger notion of a B-contracting pair, the statement holds true for all trees $U \in cv_0(F_n)$. □

3. Lines of minima and their shadows

The goal of this section is to show that contracting lines of minima are strongly Morse coarse geodesics for the symmetrized Lipschitz distance whose shadows in the free factor graph are parametrized quasi-geodesics.

Fix once and for all a number $\epsilon > 0$, so that the ϵ–thick part $\mathrm{Thick}_\epsilon(F_n)$ of Outer space is nonempty and path connected. Let

$$\Upsilon : cv_0(F_n) \to \mathcal{FF} \tag{4}$$

be a map which associates to a tree T the free factor generated by some basic primitive element for T (i.e. a primitive element $\alpha \in F_n$ which can be represented by a loop on T/F_n of length at most two). To simplify the notations we always assume from now on that a line of minima defined by a B-contracting pair is contracting.

Lemma 3.1. *For every $B > 0$ there is a number $R = R(B) > 0$ with the following property. Let $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$ be a line of minima defined by a B-contracting pair and let α be a primitive conjugacy class in F_n.*

Suppose that $T, T' \in \mathrm{Thick}_\epsilon(F_n)$ are two trees for which α is basic. Then

$$d(\Pi_\gamma(T), \Pi_\gamma(T')) \le R.$$

Proof. Recall that both T and T' are normalized so that the volume of the quotient graph $T/F_n, T'/F_n$ is 1. Let μ, ν be the measured laminations defining the line of minima γ, normalized so that $T \in \mathrm{Bal}(\mu, \nu)$ and hence $\Pi_\gamma(T) = \gamma(0)$.

Suppose that there is a primitive conjugacy class α which is basic for both T and T' and let ξ_α be the measured lamination dual to α. Suppose furthermore that $d(\Pi_\gamma(T), \Pi_\gamma(T')) \ge R > 2B + \kappa$ where $\kappa > 0$ is as in Proposition 2.14. By (3) of Proposition 2.14 and the definition of a balancing projection, this implies that $T' \in \mathrm{Bal}(e^s\mu, e^{-s}\nu)$ for some $|s| > B$.

Let $c > 0$ be so that $cT' \in \Sigma(\gamma(0))$. By property (3) in Definition 2.9 and the requirement that α is basic for T, we have

$$\langle T', \xi_\alpha \rangle / \langle \gamma(0), \xi_\alpha \rangle \ge 1/Bc.$$

As $\gamma(0) \in \mathrm{Thick}_\epsilon(F_n)$ we have $\langle \gamma(0), \xi_\alpha \rangle \ge \epsilon$ and therefore since ξ_α is basic for T' we conclude that

$$2 \ge \langle T', \xi_\alpha \rangle \ge \epsilon/Bc.$$

In particular, $1/c \le 2B/\epsilon$.

On the other hand, from the definitions (see the detailed discussion in Section 4 of [H14]), we get

$$d_L(\gamma(0), T') = 1/c.$$

Now by Proposition 2.14, a large distance between the projections $\Pi_\gamma(T) = \gamma(0)$ and $\Pi_\gamma(T')$ implies a large distance between $\gamma(0)$ and T', and hence a small c. This contradicts that $1/c \le 2B/\epsilon$ and finishes the proof. □

Proposition 3.2. *For every $B > 0$ there is a number $L = L(B) > 0$ with the following property. If $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$ is a line of minima defined by a B-contracting pair, then the image of γ under the map Υ is an L-quasi-geodesic in $\mathcal{FF}$.*

Proof. Let $([\mu], [\nu]) \in \mathcal{PML}^2$ be a B-contracting pair with associated contracting line of minima γ. Let

$$\mathcal{P} \subset F_n$$

be the collection of all primitive conjugacy classes of F_n. Define a map

$$\Psi : \mathcal{P} \to \gamma$$

by associating to $\alpha \in \mathcal{P}$ a point $\Psi(\alpha) = \gamma(t)$ as follows. Choose a tree $T \in \mathrm{Thick}_\epsilon(F_n)$ such that α is basic for T. Define $\Psi(\alpha) = \Pi_\gamma(T)$ where $\Pi_\gamma : cv_0(F_n) \to \gamma$ is the balancing projection. By Proposition 2.9 and Lemma 3.1, this is a coarsely well-defined map. This means that there is a universal constant $C = C(B) > 0$ such that for any other choice Π'_γ of such a map we have $d(\Pi_\gamma(\alpha), \Pi'_\gamma(\alpha)) \le C$ for all $\alpha \in \mathcal{P}$.

Each element $\alpha \in \mathcal{P}$ generates the conjugacy class of a rank one free factor $\langle\alpha\rangle$ of F_n and hence $\mathcal{P}$ can be viewed as a subset of the vertex set of the free factor graph. Thus the map Ψ is a map between metric spaces where $\mathcal{P}$ is equipped with the restriction of the metric on $\mathcal{FF}$ and where γ is equipped with the restriction of the metric d. We claim that Ψ is $2R$-Lipschitz where $R = R(B) > 0$ is as in Lemma 3.1.

To this end let $\alpha, \beta \in \mathcal{P}$ be primitive conjugacy classes which generate rank one free factors $\langle\alpha\rangle$ and $\langle\beta\rangle$ of distance two in the free factor graph. Then up to conjugation, there is a proper free factor A of F_n so that $\langle\alpha\rangle < A, \langle\beta\rangle < A$. As a consequence, there are representatives a of α, b of β, and there is a primitive conjugacy class $\zeta \in \mathcal{P}$ and a representative u of ζ such that a, u and u, b can be completed to a free basis of F_n.

Choose a tree $T \in \mathrm{Thick}_\epsilon(F_n)$ so that both α, ζ are primitive basic for T, and choose a tree $S \in \mathrm{Thick}_\epsilon(F_n)$ so that both ζ, β are primitive basic for S. By Lemma 3.1, $\Psi(\alpha)$ and $\Psi(\zeta)$ as well as $\Psi(\zeta)$ and $\Psi(\beta)$ are R-close to each other. Hence, $\Psi(\alpha)$ and $\Psi(\beta)$ are $2R$-close.

Now any two points $\alpha, \beta \in \mathcal{P}$ can be connected in $\mathcal{P}$ by a sequence $(\alpha_i)_{0\leq i\leq n}$ with $\alpha_0 = \alpha, \alpha_n = \beta$ whose length n is not bigger than the distance between $\langle\alpha\rangle$ and $\langle\beta\rangle$ in $\mathcal{FF}$ and such that moreover $d_{\mathcal{FF}}(\langle\alpha_i\rangle, \langle\alpha_{i+1}\rangle) \leq 2$.

Namely, connect $\langle\alpha\rangle$ to $\langle\beta\rangle$ by a geodesic (A_i) in $\mathcal{FF}$. For each i choose a rank one free factor $\langle\alpha_i\rangle < A_i$. It now suffices to observe that $d_{\mathcal{FF}}(\langle\alpha_i\rangle, \langle\alpha_{i+1}\rangle) \leq 2$ for all i. But this follows from the fact that whenever A, B are free factors with $d_{\mathcal{FF}}(A, B) = 1$ then up to exchanging A and B and conjugation we have $A < B$. Thus if $\xi \in \mathcal{P}$ generates a free factor $\langle\xi\rangle$ of A, $\rho \in \mathcal{P}$ generates a free factor $\langle\rho\rangle$ of B, then $d_{\mathcal{FF}}(\langle\xi\rangle, \langle\rho\rangle) \leq 2$.

As the map Ψ maps two primitive conjugacy classes of distance two in $\mathcal{FF}$ to points on γ which are $2R$-close, the map Ψ expands distances at most by a factor of $2R$ which is what we wanted to show. Furthermore, the two-neighborhood of $\mathcal{P}$ is all of $\mathcal{FF}$ and hence Ψ can be extended to a coarsely well defined $4R$-Lipschitz map from $\mathcal{FF}$ into γ which we denote again by Ψ.

Another application of Lemma 3.1 shows that for every $t \in \mathbb{R}$ and any primitive basic element ξ for $\gamma(t)$ we have

$$d(\Psi(\xi), \gamma(t)) \leq R.$$

We conclude that

$$d(\Psi \circ \Upsilon(\gamma(t)), \gamma(t)) \leq R.$$

The map $\Upsilon : cv_0(F_n) \to \mathcal{FF}$ is coarsely M-Lipschitz for some number $M > 0$, i.e. we have $d_{\mathcal{FF}}(\Upsilon T, \Upsilon T') \leq Md(T, T') + M$ for all $T, T' \in cv_0(F_n)$ [BF14]. Together with the above properties of the map Ψ, this shows that $\Upsilon \circ \Psi$ is a coarse $4MR$-Lipschitz retraction of $\mathcal{FF}$ onto $\Upsilon(\gamma)$. In particular, it maps a point on $\Upsilon(\gamma)$ to a point of distance at most $4MR$.

As a consequence, $\Upsilon(\gamma)$ is a parametrized $4MR$-quasi-geodesic in $\mathcal{FF}$. Namely, for $s < t$ let $g : [0, N] \to \mathcal{FF}$ be a simplicial geodesic joining $\Upsilon(\gamma(s))$ to $\Upsilon(\gamma(t))$.

The retractions $\Upsilon(\Psi(g(i)))$ are points on $\Upsilon(\gamma)$ which are of distance at most $4MR$ apart, and the endpoints $\Upsilon(\Psi(g(0))), \Upsilon(\Psi(g(N)))$ are of distance at most $2MR$ from $\Upsilon(\gamma(s)) = g(0)$ and $\Upsilon(\gamma(t)) = g(N)$.

As a consequence, there is an edge path in $\mathcal{FF}$ connecting $\Upsilon(\gamma(s))$ to $\Upsilon(\gamma(t))$ whose image is contained in the $4MR$-neighborhood of $\Upsilon(\gamma)$ and whose length does not exceed $4MRd_{\mathcal{FF}}(\Upsilon(\gamma(s)), \Upsilon(\gamma(t)))$. Since $s < t$ was arbitrary this yields that indeed $\Upsilon(\gamma)$ is an $4MR$-quasi-geodesic in $\mathcal{FF}$. □

Recall from Section 2.1 that the Gromov boundary $\partial\mathcal{FF}$ of $\mathcal{FF}$ can be identified with the space of equivalence classes of arational projective trees in $\partial\mathrm{CV}(F_n)$. The equivalence relation $\sim$ on this set is such that two trees $[S], [T]$ are equivalent if there is a measured lamination μ dual to both $[T], [S]$.

Corollary 3.3. *Let (μ, ν) be a contracting pair defining a line of minima γ.*

(1) μ, ν are dual to arational projective trees defining the endpoints of $\Upsilon(\gamma)$ in $\partial\mathcal{FF}$.
(2) The limit of any convergent subsequence of $[\gamma(t)] \subset \mathrm{CV}(F_n)$ as $t \to \pm\infty$ is an arational tree defining an endpoint of $\Upsilon(\gamma)$ in $\partial\mathcal{FF}$.

Proof. Let γ be a line of minima, defined by the B-contracting pair (μ, ν). It follows from the defining properties of a contracting line of minima (compare [H14] for details) that

$$\langle \gamma(t), \mu \rangle \to 0 \quad (t \to \infty). \tag{5}$$

As the restriction to $cv_0(F_n)$ of the map $cv(F_n) \to \mathrm{CV}(F_n)$ which associates to an F_n-tree its projectivization is a homeomorphism, the map Υ defined in equation (4) induces a map $\mathrm{CV}(F_n) \to \mathcal{FF}$ which we denote again by Υ. Using this convention, in Section 10 of [H12] the following is shown. Let us assume that $[T_i] \subset \mathrm{CV}(F_n)$ is any sequence with the property that $\Upsilon([T_i]) \subset \mathcal{FF}$ converges to a point $\xi \in \partial\mathcal{FF}$. Assume furthermore that the sequence $[T_i]$ converges in $\overline{\mathrm{CV}(F_n)}$ to a tree $[S] \in \partial\mathrm{CV}(F_n)$; then $[S]$ is arational and represents ξ.

Thus by Proposition 3.2, any limit $[T]$ in $\overline{\mathrm{CV}(F_n)}$ of a sequence of projectivized trees $[\gamma(t_i)]$ $(t_i \to \infty)$ is an arational tree whose equivalence class represents the forward endpoint of $\Upsilon(\gamma)$ in $\partial\mathcal{FF}$. By continuity of the length pairing and (5), we have $\langle T, \mu \rangle = 0$. □

Recall from the introduction the definition of a strongly Morse coarse geodesic. The next observation shows that lines of minima are strongly Morse. For its formulation, for a subset A of $cv_0(F_n)$ we denote by $N_M(A)$ the M-neighborhood of A for the two-sided Lipschitz metric.

Proposition 3.4. *For all $B > 0, K > 1$ there is a constant $M = M(B, K) > 0$ with the following property. Let $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$ be a B-contracting line of minima. Then every K-quasi-geodesic $\sigma \subset cv_0(F_n)$ for the one-sided Lipschitz metric or for the symmetrized Lipschitz metric with endpoints on γ is contained in $N_M(\gamma)$.*

Proof. The argument is standard; we follow the clear proof in Lemma 3.3 of [S14]. We show the claim for quasi-geodesics for the one-sided Lipschitz metric d_L, the case of the two-sided Lipschitz metric follows from the same argument.

Let γ be a B-contracting line of minima and let $\sigma : [a,b] \to cv_0(F_n)$ be a K-quasi-geodesic for the one-sided Lipschitz metric with endpoints on γ. Assume without loss of generality that σ is continuous.

Let $\Pi_\gamma : \mathrm{Thick}_\epsilon(F_n) \to \gamma$ be the balancing projection. By Proposition 2.14 there is a number $\kappa = \kappa(B) > 1$ with the following property. If $d(\Pi_\gamma(x), \Pi_\gamma(y)) \geq \kappa$ then

$$d_L(x,y) \geq d_L(x, \Pi_\gamma(x)) + d_L(\Pi_\gamma(x), \Pi_\gamma(y)) + d_L(\Pi_\gamma(y), y) - \kappa. \tag{6}$$

Set $A = 4\kappa K^2$. Let $[s_1, s_2] \subset [a,b]$ be a maximal connected subinterval such that $d(\sigma(s), \gamma) \geq A$ for all $s \in [s_1, s_2]$. Let $s_1 = r_1 < \cdots < r_m < r_{m+1} = s_2$ be such that $d_L(\sigma(r_i), \sigma(r_{i+1})) = A$ for $i < m$ and $d_L(\sigma(r_m), \sigma(r_{m+1})) \leq A$. By the assumption on $\sigma|[s_1, s_2]$, by the estimate (6) and the fact that $d(x,y) \geq d_L(x,y)$ for all x, y, we have

$$d_L(\Pi_\gamma(\sigma(r_i)), \Pi_\gamma(\sigma(r_{i+1}))) \leq \kappa \quad \forall i. \tag{7}$$

Now σ is a K-quasi-geodesic and hence

$$A \leq K(r_{i+1} - r_i) + K$$

and $s_2 - s_1 \geq A(m-1)/K - (m-1)$. Then

$$d_L(\sigma(s_1), \sigma(s_2)) \geq (m-1)(A/K - 1)/K - K \geq 3(m-1)\kappa - K.$$

On the other hand, summing inequality (7) over all i and using (2) of Proposition 2.14 and the assumption that the symmetrized Lipschitz distance between $\sigma(s_1), \sigma(s_2)$ and γ equals A, we obtain

$$d_L(\sigma(s_1), \sigma(s_2)) \leq 2A + 2\kappa + (m-1)\kappa = 8\kappa K^2 + 2\kappa + (m-1)\kappa.$$

This shows $2(m-1)\kappa \leq 8\kappa K^2 + 2\kappa + K$ and therefore m is bounded by a number only depending on B and K. This is what we wanted to show. ☐

Remark 3.5. As the number M in Proposition 3.4 only depends on B and K, by local compactness the statement of the lemma is also valid for continuous two-sided infinite quasi-geodesics which connect the endpoints of γ. By such a quasi-geodesic we mean a map ζ which is a limit in the topology of uniform convergence on compact sets of a sequence of quasi-geodesics with endpoints $\gamma(a_i), \gamma(b_i)$ and such that $a_i \to -\infty$, $b_i \to \infty$.

4. Strongly Morse coarse geodesics

In Section 3 we showed that a contracting line of minima projects to a parametrized quasi-geodesic in the free factor graph. Furthermore, a B-contracting line of minima is uniformly strongly Morse.

The goal of this section is to study strongly Morse coarse geodesics in the thick part of Outer space and prove Theorem 1 from the introduction.

Our main tool to this end are *fast folding paths*. Following [FM11], we define such a fast folding path connecting two points $S, T \in cv_0(F_n)$ as follows. In the simplex defined by S (which consists of all normalized trees obtained from S by rescaling the edges), there is a point S' and an optimal train track map $S' \to T$. This train track map then induces a folding path connecting S' to S. Compose this path with a rescaling path connecting S to S' and call the resulting path a fast folding path from S to T.

There is a notion of parametrization by arc length of fast folding paths. Fast folding paths parametrized by arc length are geodesics for the one-sided Lipschitz metric [FM11] and hence by Proposition 3.4 we have

Lemma 4.1. *Let $\gamma : \mathbb{R} \to \mathrm{Thick}_\epsilon(F_n)$ be a uniformly strongly Morse c-coarse geodesic for the symmetrized Lipschitz metric. There exists a number $M > 0$ with the following property.*

1. *(1) For any two points $x, y \in \gamma$, the Hausdorff distance between a fast folding path connecting x to y and a subsegment of γ connecting x to y is at most M.*
2. *(2) There exists a fast folding path $\zeta : \mathbb{R} \to cv_0(F_n)$ whose Hausdorff distance to γ is at most M.*

Proof. Let $\gamma : \mathbb{R} \to \mathrm{Thick}_\epsilon(F_n)$ be a uniformly strongly Morse c-coarse geodesic. For some $s < t$ connect $\gamma(s)$ to $\gamma(t)$ by a fast folding path ζ. Since such a path is a geodesic for the one-sided Lipschitz metric and γ is uniformly strongly Morse, the image of the path is contained in a uniformly bounded neighborhood of γ for the two-sided Lipschitz metric. In particular, this image is contained in $\mathrm{Thick}_\delta(F_n)$ for a number $\delta > 0$ only depending on the constants in the definition of a Morse coarse geodesic.

Since γ is a c-coarse geodesic for the symmetrized Lipschitz metric and the symmetrized Lipschitz metric is proportional to the one-sided Lipschitz metric for points in $\mathrm{Thick}_\epsilon(F_n)$ [AB12], the path γ is a uniform quasi-geodesic for the one-sided Lipschitz metric. This implies that the Hausdorff distance for the two-sided Lipschitz distance between $\gamma[s,t]$ and ζ is unformly bounded. Namely, otherwise there is a long subsegment γ_0 of $\gamma[s,t]$ not contained in a uniformly bounded neighborhood of ζ. We may assume that the symmetrized distance between the endpoints of γ_0 and ζ is uniformly bounded. But ζ is a geodesic for the one-sided Lipschitz metric contained in a uniformly bounded (for the symmetrized Lipschitz distance) neighborhood of γ_0 and therefore the distance between the endpoints of γ_0 has to be uniformly bounded. As γ is a uniform quasi-geodesic for the one-sided Lipschitz distance, the length of γ_0 has to be uniformly bounded which is what we wanted to show.

Now let β_n be a fast folding path connecting $\gamma(-n)$ to $\gamma(n)$. Apply the Arzela Ascoli theorem for folding paths to the paths β_n (we refer to Proposition 3.7 of [H12] and its proof for an explanation why this is possible) and obtain a biinfinite fast folding path β whose Hausdorff distance to γ is bounded by a constant only depending on B. □

Call a fast folding path ζ in $cv_0(F_n)$ *stable* if it is entirely contained in $\text{Thick}_\delta(F_n)$ for some $\delta > 0$ and if furthermore ζ is strongly Morse. Lemma 4.1 shows that for any finite or infinite strongly Morse coarse geodesic $\gamma : [a, b] \to \text{Thick}_\epsilon(F_n)$ there exists a stable fast folding path whose Hausdorff distance to γ is uniformly bounded.

If $\gamma(t)$ is a fast folding path, then for sufficiently large T_0 and all $T_0 < s < t$ there exists a *morphism* $f_{s,t} : \gamma(s)/F_n \to \gamma(t)/F_n$, ie a homotopy equivalence which is a homothety on edges, with scaling factor independent on the edge. We next use morphisms to study fast folding paths which are not stable.

Lemma 4.2. *Let $\gamma(t) : \mathbb{R} \to \text{Thick}_\epsilon(F_n)$ be a fast folding path such that for each t, the graph $G_t = \gamma(t)/F_n$ contains a non-degenerate proper subgraph E_t with fundamental group F_k for some $1 \leq k < n$ such that $f_{s,t}(E_s^i) = E_t^i$ for all $s < t$; then $\gamma(t)$ is not stable.*

Proof. Since $\gamma(t) \in \text{Thick}_\epsilon(F_n)$, the volume of $G_t - E_t$ is bounded from below by ϵ independent of t.

Modify the path of graphs G_t by scaling E_t with a constant $a(t) \in [0, 1]$ and renormalizing the volume. If the function a is chosen in such a way that its value equals one at the endpoints, that it decreases very slowly to some chosen number $\delta/2 > 0$ and increases again slowly to one (here very slowly means that the derivative of a is required to be very small) then the modified path is a uniform quasi-geodesic for d_L (which can easily be checked using the definition (3) for d_L) which passes through $cv_0(F_n) - \text{Thick}_\delta(F_n)$. Thus γ is not stable. □

The following example was suggested to us by an anonymous referee; it shows that fast folding paths which are entirely contained in $\text{Thick}_\epsilon(F_n)$ need not be stable.

Example 4.3. Let R be a marked metric rose with two petals which is also a train track for an iwip $\varphi \in \text{Out}(F_2)$. Let G be two copies of R wedged together and let Φ be the induced outer automorphism (perform φ independently on each copy of R) which has G as a train track representative. Then Φ is reducible, but the obvious fast folding path γ from G to $G \circ \Phi^n$ stays uniformly in the thick part of Outer space.

On the other hand, Lemma 4.2 shows that γ is not stable.

We showed so far that strongly Morse coarse geodesics in $\mathrm{Thick}_\delta(F_n)$ are fellow-traveled by stable fast folding paths which are entirely contained in $\mathrm{Thick}_\epsilon(F_n)$ for some $\epsilon > 0$ only depending on the quality of the Morse coarse geodesic. Now if $\gamma : [0,\infty) \to \mathrm{Thick}_\epsilon(F_n)$ is a one-sided infinite fast folding path then $[\gamma(t)]$ converges as $t \to \infty$ in $\overline{\mathrm{CV}(F_n)}$ to a tree $[T] \in \partial\mathrm{CV}(F_n)$. We call $[T]$ the *endpoint tree* of the path. The tree is called *uniquely ergometric* (see [NPR14] for this notion) if it admits a unique non-atomic length measure up to scale. Here a length measure assigns to each non-degenerate segment of T a positive length, and this length is invariant under the action of F_n and additive with respect to concatenation.

Our next goal is to show that endpoints of stable fast folding paths are very special. The following lemma uses some technical result of [NPR14] as an essential ingredient. These results are summarized in the appendix.

Lemma 4.4. *The endpoint tree of a stable fast folding path $\gamma : [0,\infty) \to \mathrm{Thick}_\epsilon(F_n)$ is uniquely ergometric.*

Proof. Let $\gamma : [0,\infty) \to \mathrm{Thick}_\epsilon(F_n)$ be a stable fast folding path. For $s \geq 0$ write $G_s = \gamma(s)/F_n$. Then G_s is a metric graph with fundamental group F_n, volume one, no univalent vertices and no loop of length smaller than ϵ. Furthermore, by perhaps removing the initial segment of γ we may assume that for $s \leq t$ there exists a morphism $f_{s,t} : G_s \to G_t$ such that $f_{s,u} = f_{t,u} \circ f_{s,t}$ for $s < t < u$. These morphisms are homotopy equivalences which are homotheties on edges, and the scaling factor for the metric does not depend on the edge.

The number of abstract graphs without univalent vertices and fundamental group F_n is finite and therefore we can find a sequence n_i with $n_{i+1} - n_i \geq 1$ such that the graphs G_{n_i} are all isomorphic, ie they are isomorphic to a fixed graph G.

An *invariant sequence of subgraphs* is a sequence of non-degenerate proper subgraphs $E_{n_i} \subset G_{n_i}$ with the property that f_{n_i,n_j} restricts to a change of marking morphism $E_{n_i} \to E_{n_j}$ up to global scaling. The sequence is *stabilized* if for large enough i and $j > i$ the restriction of f_{n_i,n_j} to E_{n_i} is a permutation. Lemma 4.2 shows that fast folding paths containing stabilized proper subgraphs are not stable. Hence we may assume from now on that γ is *reduced*, ie it does not contain stabilized subgraphs (compare [NPR14]).

We now argue by contradiction and we assume that the endpoint tree T of γ is not uniquely ergometric. Since by the above remark the sequence G_{n_i} is reduced, we can apply the results from Section 6 of [NPR14].

Following [NPR14], by passing to a subsequence we may assume that the graphs G_{n_i} converge as $i \to \infty$ in the moduli space of metric graphs of volume one to a metric graph $\hat{G}$. Since the graphs G_{n_i} are all isomorphic to G as abstract graphs, the graph $\hat{G}$ can be obtained from G by collapsing a (perhaps empty) set E of edges to points. Thus $\hat{G}$ is abstractly isomorphic to the quotient graph G/E. As $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$, the graph G/E has fundamental group F_n. Let $\Pi : G \to G/E$ be the natural collapsing map.

A *transverse decomposition* of the graph G is a collection $H^0, H^1, \dots, H^k$ of subgraphs of G such that each edge e of G is contained in precisely one of the graphs H^i. By Theorem 5.6 of [NPR14] (see Theorem A.1 of the appendix), there exists a *transverse decomposition* of G into subgraphs $H^0, H^1, \dots, H^k$ with $k \geq 2$ which record the geometry of the different length measures on the limit tree $[T]$ in the following way. First, by Corollary 5.15 of [NPR14], we have $E = H^0$, and for $j \geq 1$ the graph $\Pi(H^j)$ does not have univalent vertices, and it contains a loop. Moreover, by Lemma 6.7 of [NPR14] (see Proposition A.2), for each ℓ there exists some i such that the subgraphs $\Pi(f_{n_i,n_{i+p}}H^j)$ $(j = 1, \dots, k, 1 \leq p \leq \ell)$ of G/E do not share any edge. Since $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$, the volumes of all graphs $f_{n_i,n_{i+p}}(H^j)$ are bounded from below by a universal constant.

The geometric setting is now very similar to the situation in Lemma 4.2. Namely, since $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$, by the volume renormalization procedure along a folding path, the graphs H^j are folded along the segment G_s $(n_i \leq s \leq n_{i+\ell})$ with roughly the same speed, and this speed is linear in s.

We now modify the path γ as follows. First collapse all edges of E in the graphs G_{n_i} to a point. Denote by $\hat{G}_{n_i}$ the resulting graph, normalized to have volume one. As the volume of E tends to zero along the sequence and as the path γ is entirely contained in $\mathrm{Thick}_\epsilon(F_n)$, by possibly removing an initial segment of γ we may assume that the distance in $cv_0(F_n)$ between the universal covers of $G_{n_i} = \gamma(n_i)/F_n$ and $\hat{G}_{n_i}$ is as small as we wish, uniformly for all i.

Fix a large number $\ell > 1$ and choose i so that the subgraphs $\Pi(f_{n_i,n_{i+p}}H^j)$ of G/E do not share any edge for $1 \leq p \leq \ell$. Then $i \to \hat{G}_{n_j}$ $(i \leq j \leq i + \ell)$ is a sequence of metric graphs which are all isomorphic to G/E, and they are connected by morphisms $\hat{f}_{n_i,n_i+p}$. Furthermore, as the graphs $\Pi(f_{n_i,n_{i+p}}H^j)$ do not share any edge, they define a transverse decomposition of G/E.

We can now change the metric on $\hat{G}_{n_{i+p}}$ by scaling $f_{n_i,n_{i+p}}(H^1)$ by a positive constant which slowly tends to a very small number as we follow the sequence, and gradually increase the scaling parameter again so that it equals one for $\hat{G}_{n_{i+\ell}}$. If ℓ is sufficiently large then as in Lemma 4.2, this deformation produces a uniform quasi-geodesic with endpoints on γ which passes through $cv_0(F_n) - \mathrm{Thick}_\delta(F_n)$ for an arbitrarily prescribed $\delta > 0$. This contradicts the assumption on stability and shows the lemma. □

Remark 4.5. It follows from the discussion in the proof of Lemma 4.4 (which can be thought of an interpretation of the results in [NPR14]) that folding paths limiting on a tree which is not uniquely ergometric share geometric properties with Teichmüller geodesics whose vertical measured foliation is not uniquely ergodic (and hence which are not stable). We refer to [NPR14] for a more comprehensive discussion.

Corollary 4.6. *A contracting line of minima has a pair of endpoints in $\partial \mathrm{CV}(F_n)$, and any such endpoint is arational and uniquely ergometric.*

Proof. By Proposition 3.4, a B-contracting line of minima γ is a uniformly strongly Morse coarse geodesic for the symmetrized Lipschitz metric. Lemma 4.1 shows that there is a stable fast folding path ζ which is contained in a uniformly bounded neighborhood of γ. This fast folding path limits on an arational tree $[T] \in \partial\mathrm{CV}(F_n)$ whose equivalence class defines the endpoint of the projection of ζ to $\mathcal{FF}$.

By Lemma 4.4, the tree $[T]$ is uniquely ergometric. Now if $[T'] \in \partial\mathrm{CV}(F_n)$ is another tree obtained as a limit of a stable fast folding path which fellow travels γ, then there exists an equivariant Lipschitz map $T \to T'$. Since T is uniquely ergometric, this implies that $[T] = [T']$. □

The following proposition finishes the proof of the implication (2) $\Rightarrow$ (1) in Theorem 1.

Proposition 4.7. *Uniformly strongly Morse coarse geodesics in* $\mathrm{Thick}_\epsilon(F_n)$ *project via the map* Υ *to uniform quasi-geodesics in* $\mathcal{FF}$.

Proof. Our goal is to show that stable fast folding paths project to uniform quasi-geodesics in the free factor graph.

Shadows of folding paths in $\mathcal{FF}$ are uniform unparametrized quasi-geodesics [BF14]. By Lemma 2.6 of [H10] and its proof (more precisely, the last paragraph of the proof which is valid in the situation at hand without modification), it therefore suffices to show the following. For $M > 0, p > 0, \epsilon > 0$ there is a number $k = k(p, M, \epsilon) > 0$ such that the endpoints of any M-stable fast folding path in $\mathrm{Thick}_\epsilon(F_n)$ whose d_L-length is at least k are mapped by the map $\Upsilon : cv_0(F_n) \to \mathcal{FF}$ to points of distance at least p. Here in the definition of a uniformly Morse path, we use the number M to quantify stability for d_L-one-quasigeodesics.

Assume to the contrary that this is not true. Then there are $M > 0, \epsilon > 0, p > 0$, and for each $i > 0$ there is an M-stable fast folding path β_i in $\mathrm{Thick}_\epsilon(F_n)$ of length i whose endpoints are mapped by Υ to points in $\mathcal{FF}$ of distance at most p. Since the images of folding paths under the map Υ are uniformly unparametrized quasi-geodesics in $\mathcal{FF}$, this then implies that $\mathrm{diam}(\Upsilon(\beta_i)) \leq q$ for all i and a universal constant $q > 0$.

Using as before invariance under the action of $\mathrm{Out}(F_n)$ and the Arzela-Ascoli theorem for folding paths, up to passing to a subsequence we may assume that the paths β_i converge as $i \to \infty$ to a one-sided infinite limiting folding path β. The path β is contained in $\mathrm{Thick}_\epsilon(F_n)$, and it connects a basepoint to a tree $[T] \in \partial\mathrm{CV}(F_n)$. Moreover, β is M-stable. By possibly removing the initial segment of the path, we may assume that it is guided by an *optimal train track map* $f : \beta(0) \to T$ (see [H12] for details). This map realizes the optimal Lipschitz constant for equivariant maps $\beta(0) \to T$, and it is an isometry on edges for a suitable representative T of $[T]$.

Our goal is to show that the tree T is arational. Once we have established this, we know that the shadow in $\mathcal{FF}$ of the fast folding path β has infinite diameter. Since β is a limit of the paths β_i for the topology of uniform convergence on compact sets and since the map $\Upsilon : cv_0(F_n) \to \mathcal{FF}$ is coarsely Lipschitz, we deduce that for every $k > 0$ and all sufficiently large i, there is a some $t_i > 0$ so that the distance between $\Upsilon(\beta_i(0))$ and $\Upsilon(\beta_i(t_i))$ is at least k. This violates the assumption that the diameter of $\Upsilon(\beta_i)$ is at most q.

The rest of the proof is concerned with showing that T is arational. Note first that by Lemma 4.4, the endpoint tree $[T]$ is uniquely ergometric. This implies that there is no tree $T' \in \partial cv(F_n)$ which can be obtained from T by a one-Lipschitz alignment preserving map $\rho : T \to T'$ collapsing a nontrivial subtree of T to a point.

Lemma 2.4 now yields that $[T]$ is arational provided that T does not have point stabilizers containing free factors. The existence of such point stabilizers is ruled out with an argument which is very similar to the proof of Lemma 4.4.

We argue by contradiction and we assume that there exists a primitive element $\alpha \in F_n$ which stabilizes a point in T. By Lemma 8.1 of [H12], there is a simplicial tree S_0 in $\partial cv(F_n)$ of covolume one such that the translation length of α on S_0 is trivial, and there is an optimal train track map $g : S_0 \to T$ which gives rise to a folding path entirely consisting of trees on which α fixes a point. We may assume that S_0/F_n is a rose with $n-1$ petals.

Modify S_0 slightly to a tree $S_1 \in cv(F_n)$ so that S_1/F_n is a rose with n petals, one very short petal corresponding to α, and such that S_0 is contained in the simplex defined by S_1. For each t connect S_1 to $\beta(t)$ by a fast folding path ξ_t. As β is stable and fast folding paths are geodesics for d_L, these fast folding paths are contained in a uniformly bounded neighborhood of β. In particular, they are contained in $\mathrm{Thick}_\nu(F_n)$ for some $\nu > 0$. Hence we can take a limit ξ of a subsequence of these folding paths as $t \to \infty$. The path ξ connects S_1 to a tree T' which is F_n-equivariantly bilipschitz to T. Since $[T]$ is uniquely ergometric, we have $[T'] = [T]$. In particular, α fixes a point on T'. Furthermore, ξ is stable.

By an application of Lemma 8.1 of [H12], there is a tree S_2 in the simplex defined by S_1, and there is a rescaling T'' of T and a train track map $f : S_2 \to T''$ which guides the fast folding path ξ. The map f is an F_n-equivariant isometry on edges.

If α fixes a point in S_2 then the folding path ξ passes through S_2 which is impossible as ξ is contained in $\mathrm{Thick}_\nu(F_n)$. Thus S_2/F_n is a rose with n petals. Let $\sigma \subset S_2$ be an axis for α and let $x \in \sigma$ be a vertex of S_2 on σ. As the translation length of α on S_2 is positive and f is an edge isometry, the point $f(x) \in T''$ is not stabilised by α.

Connect $f(x)$ by a minimal segment s to the fixed point set $\mathrm{Fix}(\alpha)$ of α in T''. Let $y \in \mathrm{Fix}(\alpha)$ be the endpoint of s. As in the proof of Lemma 8.1 of [H12], we observe that the geodesic segment in T'' connecting $f(x)$ to $\alpha f(x)$ passes through y. The turn at x defining the two directions of the axis of α is illegal.

Now recall that the quotient graph S_2/F_n is a rose. By the above discussion, the folding procedure identifies the initial and terminal subsegment of the petal defined by α with unit speed. As all illegal turns are folded at once, with unit speed, we conclude that each of the quotient graphs of the fast folding path ξ contains an embedded loop consisting of a single edge which corresponds to α. Lemma 4.4 now shows that ξ can not be stable. The proposition is proven. □

Remark 4.8. Similar to the case of Teichmüller space with the Weil-Petersson metric (see [CS15] for a discussion), we believe that there are strongly Morse coarse geodesics in the metric completion of Outer space. Note that by [A12], this metric completion is the space of simplicial F_n-trees with quotient of volume one and with all edge stabilisers trivial.

5. Convex cocompact subgroups of $\mathrm{Out}(F_n)$

In Section 3 we showed that a contracting line of minima projects to a parametrized quasi-geodesic in the free factor graph. This quasi-geodesic admits a pair of endpoints in the boundary $\partial\mathcal{FF}$ of the free factor graph, and by Corollary 4.6, these endpoints are represented by arational uniquely ergometric projective trees.

In fact, more is true. We use the following definition which is taken from Section 2 of [H14].

Definition 5.1. A projective tree $[T] \in \partial\mathrm{CV}(F_n)$ is *doubly uniquely ergodic* if the following two conditions are satisfied.

(1) There exists a unique projective measured lamination $[\mu] \in \mathcal{PML}$ which is dual to $[T]$.
(2) If $[\mu]$ is dual to $[T]$ and if $[S]$ is dual to $[\mu]$ then $[S] = [T]$.

Denote by $\mathcal{UT} \subset \partial\mathrm{CV}(F_n)$ the $\mathrm{Out}(F_n)$-invariant set of doubly uniquely ergodic trees. Lemma 2.9 of [H14] shows that a fixed point in $\partial\mathrm{CV}(F_n)$ of any iwip element of $\mathrm{Out}(F_n)$ is contained in $\mathcal{UT}$. Moreover, the action of $\mathrm{Out}(F_n)$ on the closure of $\mathcal{UT}$ in $\partial\mathrm{CV}(F_n)$ is minimal.

The following corollary is immediate from Corollary 4.6 and Theorem 1.1 of [NPR14].

Proposition 5.2. *Any contracting line of minima admits a pair of endpoints in $\partial\mathrm{CV}(F_n)$, and such an endpoint is doubly uniquely ergodic and dual to a defining measured lamination for the line of minima.*

Proof. Let γ be a B-contracting line of minima. By Corollary 4.1, there exists a fast folding path β whose Hausdorff distance to γ is at most $M(B)$. This fast folding path is contained in $\mathrm{Thick}_\epsilon(F_n)$ for some $\epsilon > 0$, and it is stable.

Corollary 4.6 shows that there is a uniquely ergometric arational tree $[T] \in \partial\mathrm{CV}(F_n)$ such that $[\beta(t)] \to [T]$ in $\overline{\mathrm{CV}(F_n)}$. In fact, $[T]$ is the unique projective tree in $\partial\mathrm{CV}(F_n)$ defining the forward endpoint in $\partial\mathcal{FF}$ of the quasi-geodesic ray $\Upsilon(\gamma[0,\infty))$.

By Proposition 3.2, γ projects to a uniform quasi-geodesic in $\mathcal{FF}$ and hence the same holds true for β. Thus the assumptions of Theorem 1.1 of [NPR14] are fulfilled. Theorem 1.1 of [NPR14] now implies that $[T]$ is doubly uniquely ergodic.

Now recall that γ is defined by the B-contracting pair $([\mu],[\nu]) \in \mathcal{PML}\times\mathcal{PML}$. It follows from the definition of a B-contracting line of minima that $\langle\gamma(t),\mu\rangle \to 0$ $(t\to\infty)$. By continuity of the length function, this implies that μ is dual to $[T]$. $\square$

The following result is the main characterization of convex cocompact subgroups of $\mathrm{Out}(F_n)$. It gives a sufficient criterion for a subgroup to be convex cocompact. That this characterization is also necessary follows from the results in Section 6.

For the purpose of its proof and for later use, following Definition 3.1 of [H14] we call a family F of non-negative functions ρ on $cv_0(F_n)$ *uniformly proper* if for every $a>0$ there is a compact subset $C(a)$ of $\mathrm{Thick}_\epsilon(F_n)$ such that $\rho^{-1}[0,a]\cap \mathrm{Thick}_\epsilon(F_n)\subset C(a)$ for every $\rho\in F$. Let moreover $\mathcal{UE}\supset\mathcal{UT}$ be the $\mathrm{Out}(F_n)$-invariant subset of $\partial\mathrm{CV}(F_n)$ of arational uniquely ergometric projective trees.

Proposition 5.3. *Let $\Gamma<\mathrm{Out}(F_n)$ be a word hyperbolic subgroup with the following properties.*

(1) *There is a Γ-equivariant homeomorphism of the Gromov boundary $\partial\Gamma$ of Γ onto a compact subset Λ of $\mathcal{UE}$.*
(2) *There is some $B>0$ so that for any two points $[S]\neq[T]\in\Lambda$, there is a pair of dual projective measured laminations $([\mu],[\nu])$ for $[S],[T]$ which is a B-contracting pair.*

Then Γ is convex cocompact.

Proof. Let γ be a contracting line of minima defined by the B-contracting pair (μ,ν), with balancing projection Π_γ. Recall that this projection associates to a tree S the unique point $\Pi_\gamma(S)=\gamma(t)$, determined by the requirement that $f_{\mu,\nu}([S])=0$ (notations as in Section 2).

By Corollary 4.6 (see also Proposition 5.2), there are unique projective trees $[T_\mu],[T_\nu]\in\mathcal{UE}\subset\partial\mathrm{CV}(F_n)$ which are dual to μ,ν. Thus the balacing projection Π_γ extends to a map

$$\Pi_\gamma:\overline{\mathrm{CV}(F_n)}-\{[T_\mu],[T_\nu]\}\to\gamma(\mathbb{R}).$$

Let $\mathcal{T}_B$ be the space of triples of points in $\mathcal{PML}$ with the property that any pair from this triple is B-contracting. Let $([\mu],[\nu],[\xi])\in\mathcal{T}_B$ and let $[U]\in\mathcal{UE}$ be the tree which is dual to $[\xi]$. Choose representatives μ,ν of $[\mu],[\nu]$ such $f_{\mu,\nu}[U]=0$. Let γ be a contracting line of minima, defined by the B-contracting pair (μ,ν); then $\Pi_\gamma([U])=\gamma(0)$.

Let ζ be a contracting line of minima defined by the contracting pair $([\mu],[\xi])$. By Corollary 4.6, $[\zeta(t)]\to[U]$ $(t\to-\infty)$ in $\overline{\mathrm{CV}(F_n)}$. Now $f_{\mu,\nu}([U])=0$ and hence by continuity of the function $f_{\mu,\nu}$ established in Lemma 2.12 and by the definitions, we have $\Pi_\gamma([\zeta(s)])\in\gamma[-1,1]$ for sufficiently small s.

By the definition of a B-contracting pair, applied to the pair $([\mu],[\nu])$, we obtain that for sufficiently small s, any fast folding path connecting $\zeta(s)$ to a tree $T\in\partial cv(F_n)$ which is dual to μ (recall that such a tree is unique up to scale) passes through a uniformly bounded neighborhood of $\gamma(0)$. An application of Corollary 4.1 then yields that the line of minima ζ passes through a uniformly bounded neighborhood of $\gamma(0)$. Furthermore, the Hausdorff distance between $\gamma[0,\infty)$ and a half-ray of ζ is uniformly bounded, and $\gamma[0,\infty)$ is up to a uniformly bounded error the largest subray of γ with this property.

An application of this reasoning to a line of minima ρ which is defined by the B-contracting pair $([\nu],[\xi])$ now shows that ρ passes through a uniformly bounded neighborhood of $\gamma(0)$ is as well. Furthermore, the set of points in $cv_0(F_n)$ which are uniformly close to all three lines of minimal γ,ζ,ρ has uniformly bounded diameter.

As a consequence, there is a coarsely well defined map

$$\Theta : \mathcal{T}_B \to \mathrm{Thick}_\epsilon(F_n)$$

which maps an ordered triple $([\mu],[\nu],[\xi]) \in \mathcal{T}_B$ to some point $\Theta([\mu],[\nu],[\xi]) \in \mathrm{Thick}_\epsilon(F_n)$ which is uniformly close to all three contracting lines of minimal defined by the three different pairs of points in the triple. The map Θ depends on choices, but there is a number $D > 0$ such that for any other choice Θ', we have

$$d(\Theta([\mu],[\nu],[\xi]),\Theta'([\mu],[\nu],[\xi])) \le D \text{ for all } ([\mu],[\nu],[\xi]) \in \mathcal{T}_B.$$

Moreover, for $\varphi \in \mathrm{Out}(F_n)$ we have $\Theta(\varphi[\mu],\varphi[\nu],\varphi[\xi]) = \varphi(\Theta([\mu],[\nu],[\xi]))$ coarsely (ie up to replacing points by points of uniformly bounded distance). The map Θ also is coarsely invariant under permutations of the three variables.

Let now $\Gamma < \mathrm{Out}(F_n)$ be as in the proposition. Let $F : \partial\Gamma \to \Lambda \subset \mathcal{UE}$ be the equivariant homeomorphism whose existence is assumption (1) of the proposition. Let $\mathcal{H}_\Gamma$ be the closure of the collections of all lines of minima defined by any two distinct points in Λ. The set $\mathcal{H}_\Gamma$ is a closed Γ-invariant subset of $\mathrm{Thick}_\epsilon(F_n)$.

The group Γ is word hyperbolic and hence it acts properly and cocompactly on the space of triples of distinct points in $\partial\Gamma$. Let A be a compact fundamental domain for this action. The subset $F^3(A) \subset \mathcal{T}_B$ is mapped by Θ to a subset of $\mathcal{H}_\Gamma \subset \mathrm{Thick}_\epsilon(F_n)$ of uniformly bounded diameter.

Namely, to a triple $([\mu],[\nu],[\xi]) \in \mathcal{T}_B$ associate the pair

$$G([\mu],[\nu],[\xi]) = (\mu,\nu)$$

of representatives of $[\mu],[\nu]$ with the following two properties.

(1) Let $[U] \in \mathcal{UE}$ be the tree which is dual to $[\xi]$; then $\langle U,\mu\rangle = \langle U,\nu\rangle$.
(2) $\mathrm{Min}\{\langle S,\mu+\nu\rangle \mid S \in \mathrm{Thick}_\epsilon(F_n)\} = 1$.

It is immediate from the earlier discussion that the map G is continuous and therefore the family of functions

$$\mathcal{F} = \{\langle\cdot,\mu+\nu\rangle \mid (\mu,\nu) \in GF^3(A)\}$$

is compact. Furthermore, $\mathcal{F}$ is uniformly positive (compare [H14]) and hence uniformly proper.

Now the map Θ can be chosen to associate to a point $z \in \mathcal{T}_B$ a point in $\mathrm{Min}_\epsilon(\mu+\nu) \subset \mathrm{Thick}_\epsilon(F_n)$ where $G(z) = (\mu,\nu)$. From this the diameter bound of the image of F^3A under Θ is immediate.

In particular, its closure K is compact. Thus by coarse equivarianc of the map Θ, Γ acts on $\mathcal{H}_\Gamma$ cocompactly. The action is proper as well since Γ acts properly on $\mathrm{Thick}_\epsilon(F_n)$.

Choose a path connected closed neighborhood U of the compact set K so that the union of the Γ-translates of this set is a path connected closed neighborhood Ω of $\mathcal{H}_\Gamma$ on which Γ acts properly and cocompactly.

Equip Ω with a Γ-invariant length metric. As Γ acts on Ω cocompactly, for $x \in \Omega$ the orbit map $g \in \Gamma \to gx \in \Omega$ is a quasi-isometry.

Let γ be a geodesic in Γ with endpoints $\gamma(-\infty) \in \partial\Gamma, \gamma(\infty) \in \partial\Gamma$. There is a corresponding B-contracting line of minima ζ in $\mathcal{H}_\Gamma$ connecting $F(\gamma(-\infty))$ to $F(\gamma(\infty))$, and this line of minima is a c-coarse geodesic in $cv_0(F_n)$ for the symmetrized Lipschitz metric for a number $c > 0$ not depending on γ. In particular, it is a c-coarse geodesic in $\Omega \supset \mathcal{H}_\Gamma$ equipped with the intrinsic path metric.

As an orbit map $\Gamma \to \Omega$ is a quasi-isometry, the contracting line of minima ζ determines an equivalence class of uniform quasi-geodesics in Γ, where two quasi-geodesics are equivalent if and only if their Hausdorff distance is uniformly bounded. We claim that the geodesic γ is contained in this class.

To this end note that by Corollary 4.6, as $t \to \pm\infty$ the projective trees $[\zeta(t)]$ converge in $\overline{\mathrm{CV}(F_n)}$ to $F(\gamma(\pm\infty))$. By hyperbolicity of Γ, the equivalence class of γ consists precisely of quasi-geodesics in Γ with the same endpoints in $\partial\mathcal{FF}$ as γ and hence the geodesic γ is contained in this class. As a consequence, for some fixed $x \in \zeta$, the orbit γx is contained in a uniformly bounded neighborhood of ζ.

Recall that the map $\Upsilon : cv_0(F_n) \to \mathcal{FF}$ is coarsely Lipschitz and coarsely $\mathrm{Out}(F_n)$-equivariant, and it maps ζ to a parametrized uniform quasi-geodesic in $\mathcal{FF}$. Together with Lemma 3.4 and Remark 3.5, this yields that an orbit map $g \in \Gamma \to gA \in \mathcal{FF}$ $(A \in \mathcal{FF})$ is a quasi-isometric embedding. □

6. Free factor graph and Outer space

In this section we consider Lipschitz paths in Outer space which project to parametrized quasi-geodesics in the free factor graph. The endpoints of such a path in $\partial\mathcal{FF}$ is a pair of equivalence classes of arational trees. We show that a pair of dual laminations for any two representatives of such trees is a B-contracting pair.

The main result is the following

Proposition 6.1. *For every $L > 1$ there exists a number $B = B(L) > 0$ with the following property. Suppose that $\gamma : \mathbb{R} \to cv_0(F_n)$ is a one-Lipschitz path which projects to an L–quasi-geodesic in $\mathcal{FF}$. Then there are unique projective arational trees which define the endpoints of γ in $\partial\mathcal{FF}$. If (μ, ν) is a pair of measured laminations which is dual to this pair of trees, then (μ, ν) is B–contracting. Furthermore, the Hausdorff distance between γ and the line of minima ζ defined by (μ, ν) is at most $D(L)$ where $D = D(L)$ only depends on L.*

We break the proof of this proposition into several lemmas which will also be useful for the extension of Proposition 6.1 formulated in Corollary 6.10.

Consider the space A of finite, one-sided infinite or biinfinite one-Lipschitz paths $\gamma : J \to cv_0(F_n)$ (for the symmetrized Lipschitz metric), parametrized on a connected closed interval $J \subset \mathbb{R}$ containing 0, whose image under the projection $\Upsilon : cv_0(F_n) \to \mathcal{FF}$ is an L-quasigeodesic.

Lemma 6.2. *There are numbers $\epsilon > 0, M_0 > 0$ with the following property. If $\gamma \in A$ is parametrized on an interval J of length $|J| \geq M_0$ then $\gamma(J) \subset \mathrm{Thick}_\epsilon(F_n)$.*

Proof. Let $T \in cv_0(F_n)$ and let $\alpha = \Upsilon(T)$; then α is a primitive conjugacy class whose length on T is at most two.

The diameter in $\mathcal{FF}$ of the set of all primitive conjugacy classes whose length on T is at most two is uniformly bounded, independent of $T \in cv_0(F_n)$ (see [H12] for details). Let $k > 0$ be such an upper bound. We may assume that $k \geq L$.

Let $M = 4kL$; if $S \in cv_0(F_n) - \mathrm{Thick}_{e^{-M}}(F_n)$ then there exists a primitive conjugacy class α of length at most e^{-M} on S. Furthermore, by the definition (3) of the one-sided Lipschitz metric, the length of α on a tree T with $d_L(S,T) \leq M$ is at most one. This implies that the diameter in $\mathcal{FF}$ of the set of all primitive conjugacy classes whose length on either S or T is at most two does not exceed $2k$.

However, if $\gamma : J \to cv_0(F_n)$ is a path in A with $|J| \geq 2M$ and if $s \in J$ is such that $\gamma(s) = S$ then there is a point $t \in J$ with $|s-t| = M$. The above discussion shows that $d(\Upsilon(\gamma(s)), \Upsilon(\gamma(t))) \leq 2k$. On the other hand, as $\gamma \in A$, we have $d(\Upsilon(\gamma(s)), \Upsilon(\gamma(t))) \geq M/L - L \geq 3k$. This contradiction completes the proof of the lemma for $M_0 = 2M$ and $\epsilon = e^{-M}$. □

From now on we assume by abuse of notation that the set A only contains paths $\gamma : J \to cv_0(F_n)$ with $J \supset [-M_0, M_0]$ where $M_0 > 0$ is as in Lemma 6.2. Then $\gamma(J) \subset \mathrm{Thick}_\epsilon(F_n)$ for all $\gamma \in A$ where $\epsilon > 0$ is as in Lemma 6.2. We equip A with the topology of uniform convergence on compact sets.

The group $\mathrm{Out}(F_n)$ acts on $\mathrm{Thick}_\epsilon(F_n)$ properly and cocompactly. Let

$$K_0 \subset \mathrm{Thick}_\epsilon(F_n)$$

be a compact fundamental domain for this action. The action of $\mathrm{Out}(F_n)$ on A is cocompact as well: given any sequence γ_i in A, choose elements $\varphi_i \in \mathrm{Out}(F_n)$ so that $\varphi_i\gamma_i(0) \in K_0$. The desired compactness follows from the Arzela-Ascoli theorem. Thus if we denote by A_0 the subset of A consisting of paths γ with $\gamma(0) \in K_0$, then for the purpose of Proposition 6.1, by invariance under the action of $\mathrm{Out}(F_n)$ it suffices to investigate the set A_0.

Let $\mathcal{FT} \subset \partial CV(F_n)$ be the set of all arational trees and let

$$\Pi : \mathcal{FT} \to \partial\mathcal{FF}$$

be the natural $\mathrm{Out}(F_n)$-equivariant projection.

Lemma 6.3. *Let $\mathcal{Q} \subset \partial\mathcal{FF}$ be the set of all endpoints of biinfinite paths in A_0, viewed as quasi-geodesics in $\mathcal{FF}$. The set $\Xi = \Pi^{-1}\mathcal{Q} \subset \partial\mathrm{CV}(F_n)$ is compact.*

Proof. Since $\partial\mathcal{FF}$ is metrisable and A_0 is sequentially compact, it follows from the definition of the Gromov topology on $\partial\mathcal{FF}$ that $\mathcal{Q}$ is sequentially compact and hence compact.

Our goal is to show that $\Xi = \Pi^{-1}(\mathcal{Q})$ is a compact subset of $\mathcal{FT} \subset \partial\mathrm{CV}(F_n)$, and for this it suffices to show that Ξ is sequentially compact. To this end take a sequence $[T_i] \subset \Xi$ which limits to a tree $[T] \in \partial\mathrm{CV}(F_n)$. As $\mathcal{Q}$ is compact, up to a passing to a subsequence, there is an element $\xi \in \mathcal{Q}$ so that $\Pi([T_i])$ converges to ξ. Now for each i choose a measured lamination ν_i dual to $[T_i]$. Since $\mathcal{PML}$ is compact, up to passing to a subsequence and normalization we may assume that $\nu_i \to \nu \in \mathcal{ML}$. By continuity of the length pairing, $[T]$ is dual to ν. On the other hand, as $\Pi([T_i]) \to \xi$, we have $\langle S, \nu\rangle = 0$ if and only if $[S] \in \mathcal{FT}$ and $\Pi(S) = \xi$. Thus indeed, Ξ is compact. ☐

For a path $\gamma : J \to cv_0(F_n)$ in the set A_0 define a *pair of projective ending laminations* to be an ordered pair $([\mu], [\nu]) \in \mathcal{PML}^2$ so that the following holds true. Assume first that $J = [-b, a]$ for some $a < \infty$; we then require that $[\mu]$ is the projective class of a lamination which is dual to a primitive basic conjugacy class for $\gamma(a)$. If $[0, \infty) \subset J$ then we require that $[\mu]$ is dual to a tree in $\Pi^{-1}(\Upsilon\gamma(\infty))$ (note that $\Upsilon\gamma$ has well defined endpoints in $\partial\mathcal{FF}$). Define similarly $[\nu]$ for the backward ray $\gamma(J \cap (-\infty, 0])$.

Let $P_0 \subset \mathcal{PML}^2$ be the set of pairs $([\mu], [\nu])$ of projective measured laminations which are pairs of projective ending laminations for all paths in A_0.

Lemma 6.4. *Up to perhaps increasing the number M_0 in the definition of A_0, the set $P_0 \subset \mathcal{PML}^2$ is compact and consists of positive pairs.*

Proof. For compactness of P_0, it suffices to establish sequential compactness. To this end let $([\mu_i], [\nu_i])$ be a sequence of points in P_0 defined by maps $\gamma_i : J_i \to cv_0(F_n)$ in A_0. Assume first that $J_i = [b_i, a_i)$ for some $b_i \leq -M$ and that $a_i \to a$ for some $a < \infty$. Up to passing to a subsequence, we may assume that $\gamma_i \to \gamma$ in A_0. Then $\gamma_i(a_i) \to \gamma(a)$, furthermore for large i the forward endpoint laminations $[\mu_i]$ of γ_i are dual to a primitive basic conjugacy class α_i for $\gamma_i(a_i)$.

Now $\gamma \subset \mathrm{Thick}_\epsilon(F_n)$ and hence the number of primitive conjugacy classes which are basis for a tree in a small neighborhood of $\gamma(0)$ is finite. Thus for infinitely many i, the conjugacy class α_i coincides with a primitive basic conjugacy class for $\gamma(0)$ and hence the same is true for the dual lamination. This is what we wanted to show.

The reasoning for an infinite endpoint is the same, using again continuity of the length function. This shows compactness of P_0.

We are left with showing that up to perhaps increasing the number M_0 in the definition of the set A_0, the set P_0 consists of positive pairs. To this end note that the subset Q of P_0 of pairs $([\mu], [\nu])$ consisting of projective laminations which are dual to a pair of endpoints of a biinfinite path in A_0 is compact and consists of positive pairs by Lemma 2.7. Furthermore, as $N \to \infty$, the sets Q_N of pairs of projective laminations defined by paths in A_0 for intervals $J \supset [-N, N]$ define a neighborhood basis of Q in P_0. Thus by Lemma 2.6, there exists $M > 0$ so that Q_M consists of positive pairs as claimed. □

Define

$$R = \{(\mu, \nu) \in \mathcal{ML}^2 \mid ([\mu], [\nu]) \in P_0, \langle T, \mu \rangle = \langle T, \nu \rangle = 1 \text{ for some } T \in K_0\}.$$

Continuity of the length pairing and compactness of P_0 and K_0 show that R is compact.

By Lemma 6.4 and Lemma 3.2 of [H14], the family of functions $\{\langle \cdot, \mu + \nu \rangle \mid (\mu, \nu) \in R\}$ on $cv_0(F_n)$ is uniformly proper. Thus the closure

$$W = \overline{\bigcup_{(\mu,\nu) \in R} \mathrm{Min}_\epsilon(\mu + \nu)}$$

is compact.

The following observation is immediate from continuity.

Lemma 6.5. *There is a number $B_1 > 0$ such that*

$$\frac{\langle T, \mu \rangle}{\langle T, \nu \rangle} \in [B_1^{-1}, B_1]$$

for all $T \in W$ and all $(\mu, \nu) \in R$.

Proof of Proposition 6.1. Using the above notations, let $P \subset R$ be the subset of R of all pairs (μ, ν) which are dual to a pair of arational trees (ie which correspond to pairs of ending laminations of biinfinite paths in A_0). Our goal is to show that each $(\mu, \nu) \in P$ is a B-contracting pair for some fixed number $B > 0$.

To this end we now use an argument from the proof of Proposition 3.8 of [H14]. Namely, using the above notation, the first requirement in the definition of a B-contracting pair is immediate from Lemma 6.5 and equivariance for $B = B_1$.

For $S \in \mathrm{Thick}_\epsilon(F_n)$ let

$$\Lambda(S) = \{\nu \in \mathcal{ML} \mid \langle S, \nu \rangle = 1\}.$$

If $S \in W$ and if $\tilde{\mu}, \tilde{\nu} \in \Lambda(S)$ are rescalings of $(\mu, \nu) \in P$ then using once more positivity, continuity and compactness, we have $\langle U, \tilde{\mu} + \tilde{\nu} \rangle \geq 1/B_2$ for all

$$U \in \Sigma(S) = \{V \mid \max\{\langle V, \nu \rangle \mid \nu \in \Lambda(S)\} = 1\}$$

where $B_2 > 0$ does not depend on $S \in W$ and $(\mu, \nu) \in P$. Thus the second requirement in the definition of a B-contracting pair holds true for P.

For measured laminations $\mu, \nu \in \mathcal{ML}$ let as before

$$\mathrm{Bal}(\mu, \nu) = \{S \in cv(F_n) \mid \langle S, \mu \rangle = \langle S, \nu \rangle\}.$$

We claim that if $[T],[T']$ is a pair of projective arational trees defining two distinct boundary points of $\mathcal{FF}$ and if μ,ν are two measured laminations supported in the zero lamination of T,T' then the sets

$$U(p)=\{[S]\in\overline{[\text{Thick}_\epsilon(F_n)]}\mid S\in\text{Bal}(e^t\mu,e^{-t}\nu)\text{ for some }t>p\}$$

$(p>0)$ form a neighborhood basis in $\overline{[\text{Thick}_\epsilon(F_n)]}$ for the set of all projective trees which are equivalent to $[T]$. By this we mean that for any open set $\mathcal{U}\subset\overline{[\text{Thick}_\epsilon(F_n)]}$ which contains the set of all projective trees equivalent to $[T]$, we have $U(p)\subset\mathcal{U}$ for all sufficiently large p.

Namely, fix a tree $V\in\text{Thick}_\epsilon(F_n)$. For $t\geq 0$ let

$$\beta(t)=e^t\mu+e^{-t}\nu/\langle V,e^t\mu+e^{-t}\nu\rangle.$$

Then $\{\beta(t)\mid t\geq 0\}$ is a compact subset of the set of all *currents* for F_n, i.e. F_n-invariant locally finite Borel measures on $\partial F_n\times\partial F_n-\Delta$. As $t\to\infty$, we have

$$\beta(t)\to\hat{\mu}=\mu/\langle V,\mu\rangle$$

in the space of currents equipped with the weak*-topology [H12]. As $\hat{\mu}$ is dual to an arational tree, we have $\langle S,\hat{\mu}\rangle=0$ if and only if $[S]$ is equivalent to $[T]$. The above claim now follows once more from continuity of the length pairing (as a pairing between F_n-trees and currents, see [KL09]).

Let $T\in\text{Min}_\epsilon(\mu+\nu)$ and assume that the first and the second property in the definition of a B-contracting pair hold true for T. Let

$$\mathcal{B}(T)\subset\Lambda(T)$$

be the closure of the set of all normalized measured laminations which are up to scaling induced by a basic primitive conjugacy class for a tree $U\in\text{Bal}(\mu,\nu)$. Then $\mathcal{B}(T)$ is a compact subset of $\Lambda(T)$ which does not contain the representatives $\hat{\mu},\hat{\nu}\in\Lambda(T)$ of the measured laminations μ,ν.

Let $D(\mu),D(\nu)\subset\Sigma(T)$ be the set of all normalized arational trees which are dual to μ,ν. By continuity of the length pairing, the set of functions

$$\mathcal{F}=\{U\to\langle U,\zeta\rangle\mid\zeta\in\mathcal{B}(T)\}$$

is compact for the compact open topology on the space of continuous functions on $\Sigma(T)$. Thus by the above discussion, their values on the set $D=D(\mu)\cup D(\nu)$ are bounded from below by a positive number $c>0$.

By continuity, there is some $p>0$ so that these functions are bounded from below by $c/2$ on $\tilde{U}(p)=\{S\in\Sigma(T)\mid[S]\in U(p)\}$. Note that $\tilde{U}(p)$ is a neighborhood of $D(\mu)$ in $\Sigma(T)$. Similarly, we find a neighborhood $\tilde{V}(q)\subset\Sigma(T)$ of $D(\nu)$ so that these functions are bounded from below by $c/2$ on $\tilde{V}(q)$. As a consequence, property (3) in Definition 2.9 holds true for T and for $B=\max\{p,q,2/c\}$.

Now by compactness and continuity of the length pairing, the same property holds true for pairs (μ',ν') in a small neighborhood of (μ,ν) in P and for trees S in a small neighborhood of T, perhaps after replacing the constant B by $2B$. As the set P is compact and hence the same holds true for

$$Z=\{((\mu,\nu),S)\in P\times W\mid S\in\text{Min}_\epsilon(\mu+\nu)\}$$

it can be covered by finitely many open sets which are controlled in this way.

Together with Remark 3.5 and Proposition 5.2, this shows that biinfinite paths in the set A_0 determine B-contracting line of minima where $B = B(L) > 0$ only depends on L. By invariance under the action of $\mathrm{Out}(F_n)$ and cocompactness, this then holds true for every biinfinite path from the collection A.

To summarize, each biinfinite path $\gamma \in A$ determines a (family of) B-contracting lines of minima $\Psi(\gamma)$, and the map $\Psi : \gamma \to \Psi(\gamma)$ is equivariant with respect to the action of $\mathrm{Out}(F_n)$. Now let $\Pi_{\Psi(\gamma)}$ be a balancing projection. The map $(\gamma, X) \to \Pi_{\Psi(\gamma)}(X)$ is equivariant as well. Thus by cocompactness of the action of $\mathrm{Out}(F_n)$, the distance between $X \in \gamma$ and $\Pi_{\Psi(\gamma)}(X) \in \Psi(\gamma)$ is bounded from above by a universal constant $D(L)$ only depending on L. This is what we wanted to show. □

Using Lemma 3.4, we obtain

Corollary 6.6. *Let $\gamma : \mathbb{R} \to \mathrm{Thick}_\epsilon(F_n)$ be a c-coarse geodesic for the symmetrized Lipschitz metric. If the path $t \to \Upsilon(\gamma(t))$ is a uniform quasi-geodesic in $\mathcal{FF}$ then γ is strongly Morse.*

Proposition 4.7 and Proposition 6.1 imply

Corollary 6.7. *If $\gamma : \mathbb{R} \to \mathrm{Thick}_\epsilon(F_n)$ is uniformly strongly Morse then there exists a uniformly contracting line of minima ζ whose Hausdorff distance to γ is uniformly bounded.*

Let $\Gamma < \mathrm{Out}(F_n)$ be convex cocompact. Then Γ is finitely generated, and for one (and hence any) $\alpha \in \mathcal{FF}$ the orbit map $g \in \Gamma \to g\alpha \in \mathcal{FF}$ is a quasi-isometric embedding. As $\mathcal{FF}$ is hyperbolic, this implies that Γ is word hyperbolic. Moreover, the Gromov boundary $\partial\Gamma$ of Γ admits a Γ-equivariant embedding into $\partial\mathcal{FF}$. We denote by

$$Q_\Gamma \subset \partial\mathcal{FF}$$

its image. Since $\partial\Gamma$ is compact, the set Q_Γ is closed, Γ-invariant and minimal for the Γ-action. Proposition 6.1 now immediately implies

Corollary 6.8. *Let $\Gamma < \mathrm{Out}(F_n)$ be convex cocompact. There is a number $B > 0$ with the following property. Let $([\mu], [\nu]) \in \mathcal{PML}^2$ be a pair of measured laminations which are dual to projective trees defining distinct points in Q_Γ. Then $([\mu], [\nu])$ is a B-contracting pair. For $R > 0$ the closed R-neighborhood of the union of all lines of minima obtained from all such pairs is Γ-invariant and Γ-cocompact.*

Corollary 6.8 and Proposition 5.2 together show

Corollary 6.9. *A convex cocompact group has the properties stated in Corollary 5.3.*

Corollary 6.9, Proposition 5.3 and Lemma 3.4 complete the proof of Theorem 3 from the introduction.

To show the implication 1) $\Rightarrow$ 2) in Theorem 1, it now suffices to establish a local version of Proposition 6.1 which holds true for all paths in the set A_0 (recall that by our convention, these paths γ are defined on a closed inverval $J \supset [-M_0, M_0]$, and any pair of ending laminations for γ is a positive pair).

For numbers $C > 0, N > 0$ call a path $\gamma : [-b, a] \to \mathrm{Thick}_\epsilon(F_n)$ *C-contracting N-relative to the endpoints* if the following holds true. Let (μ, ν) be a pair of ending laminations for γ. Then (μ, ν) is a positive pair, and for all $t \in [-b + N, a - N]$ the following holds true.

(i) There exists a number $\sigma(t) \in \mathbb{R}$ such that $d(\gamma(t), \mathrm{Min}_\epsilon(e^{\sigma(t)}\mu + e^{-\sigma(t)}\nu)) \leq C$.
(ii) Let $\mathcal{B}(T) \subset \Lambda(\gamma(t))$ be the set of all normalized measured laminations which are up to scaling induced by a basic primitive conjugacy class for a tree $U \in \mathrm{Bal}(e^{\sigma(t)}\mu, e^{-\sigma(t)}\nu)$. Then $\langle S, \xi\rangle \geq 1/C$ for every $\xi \in \mathcal{B}(T)$ and for every tree
$$S \in \Sigma(\gamma(t)) \cap \Big(\cup_{s\in(-\infty,\sigma(t)-C)\cup(\sigma(t)+C,\infty)}\mathrm{Bal}(e^s\mu, e^{-s}\nu)\Big).$$

In other words, the restriction of the path γ to the subinterval $[-b + N, a - N]$ has all the contraction properties of a C-contracting line of minima. Note that by the above definition and Lemma 6.5, a contracting line of minima is precisely an endpoint-relative contracting biinfinite path.

For the formulation of the next corollary, let $B = B(L) > 0$ be as in Proposition 6.1.

Corollary 6.10. *For every $L > 1$ there are numbers $M = M(L) > 0, N = N(L) > 0$ with the following property. Let $J \supset [-M, M]$ and let $\gamma : J \to cv_0(F_n)$ be any one-Lipschitz path whose image under Υ is an L-quasi-geodesic in $\mathcal{FF}$. Then γ is $2B$-contracting $N(L)$-relative to its endpoints.*

Proof. We argue by contradiction and assume that the corollary does not hold for $2B(L)$. This means that no number $N(L)$ can be found which fulfills the above requirements. Assume without loss of generality that $2B(L) \leq M_0$ where $M_0 > 0$ is as in the definition of the set A_0.

By invariance under the action of the mapping class group, there is then a sequence of paths $\gamma_i \in A_0$, defined on intervals $[-M_0 - i, M_0 + i]$, such that property (ii) above is violated at $\gamma(0)$ with $C = 2B(L)$ and a pair $([\mu_i], [\nu_i])$ of ending laminations for γ_i. Choose representatives μ_i, ν_i which are normalized in such a way that $\langle\gamma_i(0), \mu_i\rangle = \langle\gamma_i(0), \nu_i\rangle = 1$.

As A_0 is compact, we may extract a converging subsequence whose limit is a biinfinite path $\gamma \in A_0$. By Lemma 6.4 and its proof, by passing to a subsequence we may assume that $\mu_i \to \mu$ and $\nu_i \to \nu$ where μ, ν are dual to the endpoints of γ. Proposition 6.1 shows that μ, ν are unique (as they are normalized at $\gamma(0)$).

Denote by $\mathcal{B}(\mu,\nu) \subset \Lambda(\gamma(0))$ the set of all normalized measured laminations which are up to scaling induced by a basic primitive conjugacy class for a tree $U \in \mathrm{Bal}(\mu,\nu)$, and for large i define similarly $\mathcal{B}(\mu_i,\nu_i)$. By continuity, we have $\Lambda(\gamma_i(0)) \to \Lambda(\gamma(0))$, $\mathrm{Bal}(\mu_i,\nu_i) \to \mathrm{Bal}(\mu,\nu)$ in the Hausdorff topology for compact subsets of $\mathcal{ML}$.

For large i let now $\mathcal{S}_i$ be the closure in $\overline{cv(F_n)}$ of the set

$$\Sigma(\gamma_i(0)) \cap \Big(\cup_{s\in(-\infty,-2B]\cup[2B,\infty)}\mathrm{Bal}(e^s\mu_i, e^{-s}\nu_i)\Big)$$

and denote by $\mathcal{S}$ the closure of $\Sigma(\gamma(0)) \cap \Big(\cup_{s\in(-\infty,2B]\cup[2B,\infty)}\mathrm{Bal}(e^s\mu, e^{-s}\nu)\Big)$. As before, $\mathcal{S}_i$ is a compact subset of $\overline{cv(F_n)}$ and the same holds true for $\mathcal{S}$. Furthermore, using again continuity of the length function, we have $\mathcal{S}_i \to \mathcal{S}$ in the Hausdorff topology for compact subsets of $cv(F_n)$.

Now $\langle S,\xi\rangle \geq 1/B$ for every $\xi \in \mathcal{B}(\mu,\nu)$ and every $S \in \mathcal{S}$. Thus by continuity of the length function and compactness, for sufficiently large i we have $\langle S,\xi\rangle \geq 1/2B$ for all $S \in \mathcal{S}_i$ and all $\xi \in \mathcal{B}(\mu_i,\nu_i)$. However, this is a contradiction to the assumption on the sequence γ_i. The corollary follows. ☐

The discussion in Section 3, in particular Lemma 3.4, now shows the following: Let $\gamma : (a,b) \to cv_0(F_n)$ be a one-Lipschitz path which projects to an L-quasi-geodesic in $\mathcal{FF}$ and whose length $b-a$ is at least $2M(L)$; then $\gamma(-b+M(L), a-M(L))$ is a strongly M-Morse quasi-geodesic. Thus the implication (1) $\Rightarrow$ (2) in Theorem 1 is established.

7. Examples

7.1. **Schottky groups.** In analogy to the theory of Kleinian groups, we call a finitely generated free convex cocompact subgroup of $\mathrm{Out}(F_n)$ a *Schottky group.* Such groups can be generated by a standard ping-pong construction [KL10, H14]. Namely, an iwip element acts with north-south dynamics on $\partial\mathrm{CV}(F_n)$. There is a unique attracting and a unique repelling fixed point. Each of these fixed points is a projective arational tree.

Call iwip elements α,β of $\mathrm{Out}(F_n)$ *independent* if the fixed point sets for the action of α,β on $\partial\mathrm{CV}(F_n)$ are disjoint. If α,β are independent then there are $k>0, \ell>0$ such that α^k,β^ℓ generate a free convex cocompact subgroup of $\mathrm{Out}(F_n)$ ([KL10] and Section 6 of [H14]). As in [FM02], this construction can be extended to groups generated by an arbitrarily large finite number of independent iwips.

7.2. **Convex cocompact subgroups of mapping class groups.** Let S be a compact surface of genus $g \geq 2$ with one puncture. Let $\mathrm{Mod}(S)$ be the mapping class group of S; then $\mathrm{Mod}(S)$ is the subgroup of $\mathrm{Out}(F_{2g})$ of all outer automorphisms which preserve the conjugacy class of the puncture of S.

Proposition 7.1. *If* $\Gamma < \mathrm{Mod}(S)$ *is convex cocompact in the sense of Farb-Mosher, then its image in* $\mathrm{Out}(F_{2g})$ *is convex cocompact in the sense of this article.*

To prove this proposition, we use several combinatorial complexes. For the surface, we require the arc graph $\mathcal{A}(S)$ and the arc-and-curve graph $\mathcal{AC}(S)$ (which is quasi-isometric to the curve graph). By (1) of Theorem 2, a subgroup Γ of $\mathrm{Mod}(S)$ is convex cocompact if and only if the orbit map on the arc-and-curve graph is a quasi-isometric embedding.

On the free group side we use the free factor graph $\mathcal{FF}$ and the free splitting graph $\mathcal{FS}$. These four graphs naturally admit maps as follows

$$\begin{array}{ccc} \mathcal{A}(S) & \longrightarrow & \mathcal{FS} \\ \downarrow & & \downarrow \\ \mathcal{AC}(S) & \longrightarrow & \mathcal{FF} \end{array}$$

The map $\mathcal{A}(S) \to \mathcal{FS}$ associates to an arc $a \subset S$ the corank one free factor $\pi_1(S-a)$ which defines a free splitting of F_{2g} (see [HH15]). Similarly, the map $\mathcal{AC}(S) \to \mathcal{FF}$ associates to an arc or curve on S some primitive element of F_n which can be realized by a simple closed curve in the complement.

In [HH15] it is shown that the map $\mathcal{A}(S) \to \mathcal{FS}$ is a quasi-isometric embedding. We expect that the natural map $\mathcal{AC}(S) \to \mathcal{FF}$ is a quasi-isometric embedding as well (compare with the related statement in [F15]). If this were known then then Proposition 7.1 would be immediate. In the proof below, we circumvent this difficulty by working directly with the groups.

Proof. By [HH15], the natural inclusion $\mathrm{Mod}(S) \to \mathrm{Out}(F_n)$ is a quasi-isometric embedding. If $\Gamma \subset \mathrm{Mod}(S)$ is convex cocompact, then there is an equivariant embedding $\partial\Gamma \to \partial\mathrm{CV}(F_n)$. Its image consists of trees which are dual to uniquely ergodic measured geodesic laminations on S, and such trees are arational.

Now a geodesic in Γ defines a uniform quasi-geodesic in $\mathrm{Mod}(S)$ which projects to a uniform quasi-geodesic in $\mathcal{A}(S)$ and hence $\mathcal{FS}$. But uniform quasi-geodesics in $\mathcal{FS}$ map to uniform unparametrized quasi-geodesics in $\mathcal{FF}$ [KR14] and hence following the reasoning in the proof of Proposition 4.7, by cocompactness all we need to establish is that the endpoint of such a parametrized quasi-geodesic in $\mathcal{FS}$ in the boundary of the free splitting graph (which contains the boundary of the free factor graph) is in fact an arational tree. However, we observed above that this is indeed the case. $\square$

Appendix A. Uniquely ergometric trees: The work of Namazi, Pettet and Reynolds

In this appendix we summarize some technical results of [NPR14] used in the proof of Lemma 4.4.

We are looking at compact graphs without univalent or bivalent vertices and fundamental group F_n. A *combinatorial morphism* between such graphs G, H is a homotopy equivalence $f : G \to H$ which takes edges of G to non-degenerate reduced edge paths in H. In particular, f maps the vertices of G to vertices in H.

Following Section 3 of [NPR14], define a *combinatorial folding sequence* to be a sequence of graphs $(G_m)_{a\leq m\leq b}$ with combinatorial morphism $f_{k,\ell} : G_k \to G_\ell$ for $k < \ell$ such that $f_{k,m} = f_{\ell,m} \circ f_{k,\ell}$. An *invariant sequence of subgraphs* is a sequence of non-degenerate proper subgraphs $E_m \subset G_m$ with the property that $f_{m,m+1}$ restricts to a change of marking morphism $E_m \to E_{m+1}$ for all $m < b$. Such a sequence is *stabilized* if the restriction of $f_{m,m+1}$ to E_m is a permutation for sufficiently large m. A sequence is *reduced* if it does not admit any stabilized sequence of subgraphs.

Let EG be the set of edges of the graph G. The *incidence matrix* M_f of a combinatorial morphism $f : G \to H$ is the matrix whose columns are indexed by the edges of G and whose rows are indexed by the edges of H. For $e \in EG, e' \in EH$, the (e, e')-entry $M_f(e', e)$ of M_f is the number of occurrences of e' in the reduced edge path $f(e)$. A combinatorial folding sequence

$$G_0 \xrightarrow{f_{0,1}} G_1 \xrightarrow{f_{1,2}} \dots$$

determines a sequence μ_n of *frequency vectors* defined by $\mu_0 = 1$ and

$$\mu_{n+1} = M_{f_{n,n+1}}\mu_n.$$

If we write $\varphi_n = f_{n-1,n} \circ \cdots \circ f_{0,1}$ then $\mu_n(e)$ is the number of times that φ_n-images of edges of G_0 travers e.

A *length vector* for a graph H associates to each edge of H a non-negative length. If $f : G \to H$ is a combinatorial morphism then any such length vector λ_H for H induces a length vector λ_G for G by the formula

$$\lambda_G = M_f^T \lambda_H,$$

ie λ_G is the pullback by f of the metric on H defined by λ_H. A length vector for G defines a *length measure* on G.

Now let $\beta(t) \subset \mathrm{Thick}_\epsilon(F_n)$ be a fast folding path which converges to a non-uniquely ergometric tree $[T] \in \partial\mathrm{CV}(F_n)$ with dense action of F_n. Assume that $\beta(t)$ is guided by an F_n-equivariant train track map $F_0 : \beta(0) \to T$. The map F_0 maps an edge of $\beta(0)$ to an embedded path in T, and it is a homothety on edges with scaling factor not depending on the edge. For each t, there exists a train track map $F_t : \beta(t) \to T$, and for $t > s$ there is a map $F_{s,t} : \beta(s) \to \beta(t)$ with $F_s = F_t \circ F_{s,t}$. As $[T]$ is not uniquely ergometric, there is a number $k \geq 2$, and there are pairwise non-homothetic length measures $\lambda^1, \dots, \lambda^k$ on $[T]$ which are ergodic under the action of F_n. For each t, these measures pull back to a length measure λ^i_t on $\beta(t)$ and hence on the $G_t = \beta(t)/F_n$, viewed as an abstract finite graph. Note that these length measures are not normalized, ie the volume of the quotient graphs may be different from one.

As there are only finitely many isomorphism classes of graphs without vertices of valence one or two and fundamental group F_n, by passing to a subsequence we may assume that the graphs G_m are all isomorphic to a fixed graph G. We may moreover assume that the maps $F_{\ell,m}$ descend to a combinatorial folding sequence

$$G_0 \xrightarrow{f_0} G_1 \xrightarrow{f_1} \dots$$

Write $\lambda_m = \sum_i \lambda_m^i$. After passing to a subsequence, we may assume that for each edge e in G, either $\liminf_m \mu_m(e)\lambda_m(e) > 0$ or $\lim_m \mu_m(e)\lambda_m(e) = 0$. Define H^i to be the union of all edges of G so that $\liminf_m \mu_m(e)\lambda_m^i(e) > 0$ and let H^0 be the union of all edges with $\lim_m \mu_m(e)\lambda_m(e) = 0$.

A *transverse decomposition* of G consists of a finite collection of subgraphs H^i of G so that every edge of G is contained in precisely one of the graphs H^i. The following is Theorem 5.6 of [NPR14].

Theorem A.1. *The sets H^i $(i = 0, \dots, k)$ define a transverse decomposition of G. Furthermore, after passing to a another subsequence, for every $e \in EH^i$ and for $j \neq i$, we have*

(1) $\liminf_m \mu_m(e)\lambda_m^i(e) > 0$,
(2) $\sum_m \mu_m(e)\lambda_m^j(e) < \infty$, *and*
(3) $\lim_m \frac{\lambda_m^j(e)}{\lambda_m^i(e)} = 0$.

The proof of this result is based on the following two observations. Lemma 5.7 of [NPR14] investigates edges e of G with $\limsup_m \mu_m(e)\lambda_m^i(e) > 0$.

For each m let $e_m \in G_m$ be the edge corresponding to e. By definition, $\mu_m(e)$ is the number of times that φ_m-images of edges in G_0 traverse e_m. Let $B_m(e) = \varphi_m^{-1}(e_m)$ be the preimage of e_m in G_0. Then

$$\lambda_0^i(B_m(e)) = \mu_m(e)\lambda_m^i(e).$$

Hence if $\limsup_m \mu_m(e)\lambda_m^i(e) = 2\epsilon > 0$ then $\lambda_0^i(B_m(e)) \geq \epsilon$ for infinitely many m. This implies that $\lambda_0^i(\cap_m \cup_{\ell \geq m} B_\ell(e)) \geq \epsilon$. Moreover, if x is in this set then $\varphi_m(x) \in e_m$ for infinitely many m.

Lemma 5.8 of [NPR14] exploits the fact that the length measures λ^i are singular. Namely, let $x \in G_0$ be a λ_0^i-generic point in G_0. For $m > 0$ denote by $e(x,m)$ the edge of G_m containing $\varphi_m(x)$. Let us denote by I_m the connected component of $\varphi_m^{-1}(e(x,m))$ containing x. Since the action of F_n on T has dense orbits for the metric λ, the λ_0-length of I_m tends to zero with m. Hence we can evaluate the Radon Nikodym derivative of λ_0^i with respect to λ_0 at x by taking the limit $\lim_m \frac{\lambda_0^i(I_m)}{\lambda_0(I_m)} = 1$. Since φ_m is a path isometry for the lengths measures on G_0, G_m induced by λ^i, λ, this shows that

$$\lim_m \frac{\lambda_m^i(e(x,m))}{\lambda_m(e(x,m))} = 1. \tag{8}$$

This observation is used as follows. Let us assume that the edge e of G satisfies $\liminf_m \mu_m(e)\lambda_m^i(e) > 0$ for some $i \leq k$. By Lemma 5.7 of [NPR14] as recorded above, there is a subset A of G_0 with $\lambda_0^i(A) > 0$ such that for every $x \in A$, we have $\varphi_m(x) \in e_m$ for infinitely many m. By passing to a subsequence, we may assume that $\varphi_m(x) \in e_m$ for all m where $x \in G_0$ is a density point for λ_0^i.

Assume for contradiction that $\limsup_m \mu_m(e)\lambda_m^j(e) > 0$ for some $j \neq i$. Using again Lemma 5.7 of [NPR14], by passing to another subsequence we may assume that there exists a generic point $y \in G_0$ for λ_0^j such that $\varphi_m(y) \in e_m$ for all m. However, this violates the estimate (8) on Radon Nikodym derivatives together with the fact that $\lambda^i + \lambda^j \leq \lambda$.

To summarize, by passing to suitable subsequences of the sequence G_m we can arrange that (1) -(3) in Theorem A.1 hold true.

The above argument shows more. By inequality (8) and the construction of the transverse decomposition of G into the subgraphs H^i, for all $\delta > 0$ there exists an infinite sequence $m_\delta(u)$ tending to infinity as $u \to \infty$ such that for all u, for every $1 \leq i \leq k$ and every edge $e \in H^i$, we have

$$\lambda^i_{m_\delta(u)}(e) \geq \lambda_{m_\delta(u)}(e)(1-\delta).$$

Let now $\ell > 0$. For $p > m$ write $f_{m,p} = f_{p-1} \circ \cdots \circ f_m : G_m \to G_p$. We claim that for sufficiently small δ and $m = m_\delta(u)$, $p \leq m + \ell$, any edge contained in the intersection of the subgraphs $f_{m,p}(H^i), f_{m,p}(H^j)$ is contained in $f_{m,p}(H^0)$. To this end note that for fixed ℓ, the entries of the incident matrix describing the morphism $f_{m,p}$ for $p \leq m + \ell$ are uniformly bounded.

Now let $p \leq m + \ell$ and let e' be an edge in G_p. If $e^i \in G_m$ is an edge so that the image of e^i under the morphism $f_{m,p} : G_m \to G_p$ contains e' and if $c^i \subset e^i$ is a connected component of $f^{-1}_{m,p}(e') \cap e^i$, then as the map $f_{m,p}$ is a path isometry, we conclude that $\lambda_m(c^i) \geq \theta \lambda_m(e^i)$ where $\theta > 0$ only depends on ℓ and the length of e' in G_p, equipped with the rescaling of λ_p whose volume equals one.

By inequality (8) and the choice of m, if $e \in H^i$ then $|\lambda^i_m(c^i)/\lambda_m(c^i) - 1| < \delta$. As this is true for all i and as $\lambda_m(c^i)$ equals $\lambda_p(e')$, independent of i, we conclude the following.

Let us assume that there is $j \neq i$, and there is another edge $e^j \in H^j$ whose image under $f_{m,p}$ contains e'. Assuming that $\delta << \theta$, on the one hand, the λ^j_p-length of e' is very close to its λ-length and hence $\lambda^j_m(e^i) \geq \theta \lambda_m(e^i)/2$. On the other hand, by the definition of the decomposition of G_m and since $\lambda \geq \lambda_i + \lambda_j$, we have $\lambda^j_m(e^i) < \delta \lambda_m(e^i)$. This is a contradiction leading to Lemma 6.7 of [NPR14] which is summarized as follows.

Proposition A.2. *Let $\beta(t) \subset \mathrm{Thick}_\epsilon(F_n)$ be a fast folding path converging to a non-uniquely ergometric tree $[T] \in \partial \mathrm{CV}(F_n)$. For $s < t$ let $f_{s,t} : G_s = \beta(s)/F_n \to G_t = \beta(t)/F_n$ be the induced morphism. Then for every $\ell > 0$ there exists a number $t > 0$ with the following properties. Let $H^0, \dots, H^k$ be the transverse decomposition of G_t as constructed in Theorem A.1. For $t \leq s \leq t + \ell$ there is a collection E_s of edges in G_s of total length smaller than an arbitrarily prescribed number with the following properties. The quotient graph G_s/E_s has fundamental group F_n, and if $\Pi : G_s \to G_s/E_s$ is the canonical projection, then for $1 \leq i < j \leq k$, the subgraphs $\Pi(f_{t,s}(H^i)), \Pi(f_{t,s}(H^j))$ of G_s/E_s do not share any edge.*

References

[A12] Y. Algom-Kfir, *The metric completion of Outer space*, arXiv:1202.6392.

[AB12] Y. Algom-Kfir, M. Bestvina, *Asymmetry of Outer space*, Geom. Dedicata 156 (2012), 81–92.

[BF92] M. Bestvina, M. Feighn, *Outer limits*, unpublished manuscript 1992.

[BF14] M. Bestvina, M. Feighn, *Hyperbolicity of the free factor complex*, Advances in Math. 256 (2014), 104–155.

[BR15] M. Bestvina, P. Reynolds, *The Gromov boundary of the complex of free factors*, Duke Math. J. 164 (2015), no. 11, 2213–2251.

[CS15] R. Charney, H. Sultan, *Contracting boundaries of* CAT(0)*-spaces*, J. Topol. 8 (2015), no. 1, 93–117.

[CL95] M. Cohen, M. Lustig, *Very small group actions on* $\mathbb{R}$*-trees and Dehn twist automorphisms*, Topology 34 (1995), 575–617.

[DT14] S. Dowdall, S. Taylor, *Hyperbolic extensions of free groups*, to appear in Geom. Top.

[DT15] S. Dowdall, S. Taylor, *Contracting orbits in Outer space*, arXiv:1502.04053.

[DT16] S. Dowdall, S. Taylor, *The co-surface graph and the geometry of hyperbolic free group extensions*, J. Topol. 10 (2017), 447–482.

[DuT15] M. Durham, S. Taylor, *Convex cocompactness and stability in mapping class groups*, Alg. & Geom. Top. 15 (2015), 2839–2859.

[FM02] B. Farb, L. Mosher, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. 6 (2002), 91–152.

[F15] M. Forlini, *Splittings of free groups from arcs and curves*, arXiv:1511.09446

[FM11] S. Francaviglia, A. Martino, *Metric properties of outer space*, Publ. Mat. 55 (2011), 433–473.

[H05] U. Hamenstädt, *Word hyperbolic extensions of surface groups*, unpublished manuscript, 2005.

[H10] U. Hamenstädt, *Stability of quasi-geodesics in Teichmüller space*, Geom. Dedicata 146 (2010), 101–116.

[H12] U. Hamenstädt, *The boundary of the free factor graph and the free splitting graph*, arXiv:1211.1630.

[H14] U. Hamenstädt, *Lines of minima in Outer space*, Duke Math. J. 163 (2014), 733–776.

[HH15] U. Hamenstädt, S. Hensel, *Spheres and projections in* $\mathrm{Out}(F_n)$, J. Topol. 8 (2015), no. 1, 65–92.

[HOP15] S. Hensel, P. Przytycki, R. Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, J. Eur. Math. Soc. (JEMS) 17 (2015), no. 4, 755–762.

[KL09] I. Kapovich, M. Lustig, *Geometric intersection numbers and analogues of the curve complex for free groups*, Geom. & Top. 13 (2009), 1805–1833.

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

[KL10] I. Kapovich, M. Lustig, *Ping-pong and outer space*, J. Topol. Anal. 2 (2010), no. 2, 173–201.

[KeL08] R. Kent, C. Leininger, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. 18 (2008), 1270–1325.

[Ma95] R. Martin, *Non-uniquely ergodic foliations of thin type, measured currents and automorphisms of free groups*, Dissertation, Los Angeles 1995.

[NPR14] H. Namazi, A. Pettet, P. Reynolds, *Ergodic decomposition for folding and unfolding paths in outer space*, arXiv:1410.8870.

[R12] P. Reynolds, *Reducing systems for very small trees*, arXiv:1211.3378.

[S14] H. Sultan, *Hyperbolic quasi-geodesics in* CAT(0)*-spaces*, Geom. Dedicata 169 (2014), 209–224.

Math. Institut der Universität Bonn
Endenicher Allee 60, 53115 Bonn, Germany
e-mail: Ursula Hamenstädt: ursula@math.uni-bonn.de,
Sebastian Hensel: hensel@math.uni-bonn.de