

Convexity of balls in outer space

Yulan Qing and Kasra Rafi

Abstract. In this paper we study the convexity properties of geodesics and balls in Outer space equipped with the Lipschitz metric. We introduce a class of geodesics called balanced folding paths and show that, for every loop α , the length of α along a balanced folding path is not larger than the maximum of its lengths at the endpoints. This implies that out-going balls are weakly convex. We then show that these results are sharp by providing several counterexamples.

Mathematics Subject Classification (2020). 20F65.

Keywords. Outer space, $\text{Out}(\mathbb{F}_n)$, folding path.

Contents

1	Introduction	893
2	Preliminaries	897
3	Weak convexity	906
4	Decorated difference of markings map	919
5	Construction of the balanced folding paths	924
6	Non-convexity	926
7	In-coming balls	931
	References	932

1. Introduction

Let $\text{Out}(\mathbb{F}_n)$ be the group of outer automorphisms of a free group of rank n , and let Outer space CV_n be the space of marked metric graphs of rank n . The Outer space, which is a simplicial complex with an $\text{Out}(\mathbb{F}_n)$ action, was introduced by Culler and Vogtmann [6] to study $\text{Out}(\mathbb{F}_n)$ as an analogue of the action of mapping class group on Teichmüller space or the action of a lattice on a symmetric space. The Outer space can be equipped with a natural asymmetric metric, namely the

Lipschitz metric. For points $x, y \in \text{CV}_n$,

$$d(x, y) = \inf_{\phi} \log(L_{\phi})$$

where $\phi: x \rightarrow y$ is a difference of markings map from x to y and L_{ϕ} is the Lipschitz constant of the map ϕ . The geometry of CV_n equipped with the Lipschitz metric is closely related to the large scale geometry of $\text{Out}(\mathbb{F}_n)$ and has been the subject of extensive study (see for example, [16, 1, 3]).

In this paper, we examine the convexity properties of geodesics and balls in CV_n . However, we need to be careful with our definitions since the metric d is asymmetric (the ratio of $d(x, y)$ and $d(y, x)$ can be arbitrarily large [2]) and there may be many geodesics connecting two points in CV_n . We introduce the notion of a *balanced folding path* along which we have more control over the lengths of loops. Recall that, a geodesic in CV_n (not necessarily parametrized with unit speed) is an injective continuous map $\gamma: [a, b] \rightarrow \text{CV}_n$. For $a \leq t \leq b$, let $x = \gamma(a)$ and $y = \gamma(b)$, we have

$$d(x, \gamma(t)) + d(\gamma(t), y) = d(x, y).$$

We often denote the image of γ by $[x, y]$. The length of a loop α in a graph x is denoted by $|\alpha|_x$ and we use $|\alpha|_t$ to denote the length of α at $\gamma(t)$. The balanced folding paths from x to y is denoted by $[x, y]_{\text{bf}}$. We show that, lengths of loops along a balanced folding path satisfy a weak notion of convexity.

Theorem 1.1. *Given points $x, y \in \text{CV}_n$, there exists a geodesic $[x, y]_{\text{bf}}$ from x to y so that, for every loop α , and every time t ,*

$$|\alpha|_t \leq \max(|\alpha|_x, |\alpha|_y).$$

The proof of this theorem is by construction. We then apply Theorem 1.1 to the convexity of balls. There are two different notions of a round ball in CV_n . For $x \in \text{CV}_n$ and $R > 0$, we define the out-going ball of radius R centered at x to be

$$B_{\text{out}}(x, R) = \{y \in \text{CV}_n \mid d(x, y) \leq R\}.$$

As an immediate corollary of Theorem 1.1 we have

Theorem 1.2. *Given a point $x \in \text{CV}_n$, a radius $R > 0$ and points $y, z \in B_{\text{out}}(x, R)$,*

$$[y, z]_{\text{bf}} \subset B_{\text{out}}(x, R).$$

That is, the ball $B_{\text{out}}(x, R)$ is weakly convex.

Corollary 1.3. *The intersection of any number of out-going balls is connected.*

Proof. The same balanced folding path is contained in all balls. □

We will also show that these theorems are sharp in various ways by providing examples of how possible stronger statements fail. There are other ways to choose a geodesic connecting y to z , for example, a *standard geodesic* which is a concatenation of a rescaling of the edges and a *greedy folding path* (see Section 2 for definitions). In fact, when there is a greedy folding path connecting y to z , Dowdall and Taylor [7, Corollary 3.3] following Bestvina and Feighn [4, Lemma 4.4] have shown that the lengths of loops are quasi-convex. However, we will show that these paths do not satisfy the conclusion of Theorem 1.1 and a standard path or a greedy folding path with endpoints in $B_{\text{out}}(x, R)$ may leave the ball. Here is a summary of our counter-examples.

Theorem 1.4. *Theorems 1.1 and 1.2 are sharp.*

- (1) **Lengths cannot be made convex.** *There are points $x, y \in \text{CV}_n$ and a loop α so that along any geodesic connecting x to y , the length of α is not a convex function of distance in CV_n .*
- (2) **The ball B_{out} is not quasi-convex.** *This is true even when one restricts attention to (non-greedy)-folding paths. Namely, for any $R > 0$, there are points $x, y, z \in \text{CV}_n$ and there is a folding path $[y, z]_{\text{ng}}$ connecting y to z so that*

$$y, z \in B_{\text{out}}(x, 2) \quad \text{and} \quad [y, z]_{\text{ng}} \not\subset B_{\text{out}}(x, R).$$

That is, a folding path with endpoints in $B_{\text{out}}(x, 2)$ can travel arbitrarily far away from x .

- (3) **Standard geodesics could behave very badly.** *There exists a constant $c > 0$ such that, for every $R > 0$, there are points $x, y, z \in \text{CV}_n$ and a standard geodesic $[y, z]_{\text{std}}$ connecting y to z such that*

$$y, z \in B_{\text{out}}(x, R) \quad \text{and} \quad [y, z]_{\text{std}} \not\subset B_{\text{out}}(x, 2R - c).$$

That is, the standard geodesic path can travel nearly twice as far from x as y and z are from x .

- (4) **Greedy folding paths may not stay in the ball.** *For every $R > 0$, there are points $x, y, z \in \text{CV}_n$, $n \geq 6$, where y and z are connected by a greedy folding path $[y, z]_{\text{gf}}$ such that*

$$y, z \in B_{\text{out}}(x, R) \quad \text{but} \quad [y, z]_{\text{gf}} \not\subset B_{\text{out}}(x, R).$$

Construction of a balanced folding path. Given an optimal difference of markings map $\phi: x \rightarrow y$ where the tension graph is the whole x , there are many folding paths connecting x to y . We need a controlled and flexible way to construct a folding path between x to y . To this end, we introduce a notion of a speed assignment (see Section 2.4) which describes how fast every illegal turn in x folds. Given a speed assignment, one can write a concrete formula for the rate of change of

the length of a loop α (Lemma 3.2). To prove Theorem 1.1, we need to find a speed assignment so that, whenever $|\alpha|_y < |\alpha|_x$, the derivative of the length of α is negative and if $|\alpha|_y > |\alpha|_x$ the length α does not grow too fast.

A difference of markings map $\phi: x \rightarrow y$ (again assuming the tension graph is the whole x) can be decomposed to a quotient map $\bar{\phi}: x \rightarrow \bar{y}$ which is a local isometry and a scaling map $\bar{y} \rightarrow y$. Our approach is to determine the contribution ℓ_τ of every sub-gate τ in x to the length loss from x to \bar{y} . For an easy example, consider $\mathbb{F}_3 = \langle a, b, c \rangle$, let x be a rose with 3 pedals where the edges are labeled ac^2, bc and c and the edge lengths are $\frac{1}{2}, \frac{1}{3}$ and $\frac{1}{6}$ and let \bar{y} be the rose with labels a, b and c and edges lengths all $\frac{1}{6}$. Then \bar{y} is obtained from x by wrapping ac^2 around c twice and bc around c once. The length loss going from x to \bar{y} is

$$|x| - |\bar{y}| = 1 - \frac{1}{2} = \frac{1}{2}.$$

Here, the sub-gate $\langle bc, c \rangle$ is contributing $\frac{1}{6}$ to the length loss and the sub-gate $\langle ac^2, c \rangle$ is contributing $2 \times \frac{1}{6}$ to the length loss. Of course, in general, the definition of length loss contribution needs to be much more subtle.

These length loss contributions are then used to determine the appropriate speed assignment. That is, we fold each sub-gates proportional to the length loss they eventually induce (see Section 3). For instance, in the above example, the sub-gate $\langle ac^2, c \rangle$ should be folded with twice the speed of the sub-gate $\langle bc, c \rangle$.

Decorated difference of markings map. If the tension graph of $\phi: x \rightarrow y$ is a proper subset of x , there is no folding path between x to y . As we see in part (3) of Theorem 1.4, standard paths are not suitable for our purposes. Instead, we emulate a folding path even in this case. Namely, we introduce a decorated difference of markings map. That is, by adding decoration to x and y (marked points are added to x and some *hair* are added to y), we can ensure that the difference of markings map is tight. Then, we show, a folding path can be defined as before and the discussion above carries through (see Section 4).

A criterion for the uniqueness of geodesics. To prove part one of Theorem 1.4 and Theorem 1.6 below, we need to know what all the geodesics connecting x to y are accounted for. In general this is hard to characterize. Instead, we focus on a case where there is a unique geodesic connecting x to y . It is not hard to prove the uniqueness of geodesic in special cases, however, we prove a general statement giving a criterion for uniqueness. A *yo-yo* is an illegal turn formed by a one-edge loop and a second edge, with no other edges incident to the vertex of this illegal turn. We say a folding path from x to y is *rigid* if at every point along the path there is exactly one illegal turn and it is not a yo-yo.

Theorem 1.5. *For points $x, y \in \text{CV}_n$, the geodesic from x and y is unique (up to reparametrization) if and only if there exists a rigid folding path connecting x to y .*

The in-coming balls. In the last section we examine the convexity of in-coming balls. For $x \in \text{CV}_n$ and $R > 0$, we define the in-coming ball of radius R centered at x to be

$$B_{\text{in}}(x, R) = \{y \in \text{CV}_n \mid d(y, x) \leq R\}.$$

We show that a ball $B_{\text{in}}(x, R)$, in general, is not even weakly quasi-convex.

Theorem 1.6. *For any constant $R > 0$, there are points $x, y, z \in \text{CV}_n$ such that, $y, z \in B_{\text{in}}(x, 2)$ but, for any geodesic $[y, z]$ connecting y to z ,*

$$[y, z] \not\subset B_{\text{in}}(x, R).$$

Once again we use Theorem 1.5; we construct an example where there is a unique geodesic between y and z and show that it can go arbitrarily far out.

Analogies with Teichmüller space. The problem addressed in this paper has a long history in the setting of Teichmüller space. Let $(\mathcal{T}(\Sigma), d_{\mathcal{T}})$ be the Teichmüller space of a surface Σ equipped with the Teichmüller metric. It was claimed by Kravetz [13] that round balls in Teichmüller space are convex and he used it to give a positive answer to the Nielsen-realization problem. However, his proof turned out to be incorrect. Even-though the Nielsen-realization problem was solved by Kerckhoff [11], it was open for many years whether or not the rounds balls are convex and was only resolved recently; It was shown in [14] that round balls in $(\mathcal{T}(\Sigma), d_{\mathcal{T}})$ are quasi-convex and it was shown in [8] that there are non-convex balls in the Teichmüller space. The problem is still open for the Teichmüller space equipped with the Thurston metric [15], which is an asymmetric metric and is more directly analogous to the Lipschitz metric in CV_n . Hence, the weak convexity in the case of Outer space is somewhat surprising.

Acknowledgements. We would like to thank Yael Algom-Kfir and Mladen Bestvina for helpful comments on an earlier version of this paper.

2. Preliminaries

2.1. Outer space. Let \mathbb{F}_n be a free group of rank n and let $\text{Out}(\mathbb{F}_n)$ be the outer automorphism group of \mathbb{F}_n . Let cv_n be the space of free, minimal actions of \mathbb{F}_n by isometries on metric simplicial trees [6]. Two such actions are considered isomorphic if there is an equivariant isometry between the corresponding trees. Equivalently we think of a point in cv_n as the quotient metric graph of the tree by

the corresponding action. The quotient graph is marked, that is, its fundamental group is identified (up to conjugation) with \mathbb{F}_n .

The *Culler–Vogtmann Outer space*, CV_n , (or simply the *Outer space*) is the subspace of cv_n consisting of all *marked metric graphs* of total length 1. Let x be a metric graph of total length 1, in which every vertex has degree at least 3. Let R_n be the graph of n edges that are all incident to one vertex. A *marking* is a homotopy equivalence $f: R_n \rightarrow x$. Two marked graphs $f: R_n \rightarrow x$ and $f': R_n \rightarrow x'$ are *equivalent* if there is an isometry $\phi: x \rightarrow x'$ such that $\phi \circ f \simeq f'$ (homotopic). When the context is clear, we often drop the marking out of the notation and simply write $x \in CV_n$. In this paper, we refer to metric graphs as x, y , etc. and the corresponding trees as T_x, T_y , etc. We also use $\tilde{\phi}: T_x \rightarrow T_y$ for the lift of ϕ .

The set of marked metric graphs that are isomorphic as marked graphs to a given point $x \in CV_n$ makes up an open simplex in CV_n which we denote Δ_x . Outer space CV_n consists of simplices with missing faces. The group $\text{Out}(\mathbb{F}_n)$ acts on CV_n by precomposing the marking: for an element $g \in \text{Out}(\mathbb{F}_n)$, $(x, f)g = (x, f \circ g)$. This is a simplicial action.

2.2. Lipschitz metric. A map $\phi: x \rightarrow y$ is a *difference of markings* map if $\phi \circ f_x \simeq f_y$. We will only consider Lipschitz maps and we denote by L_ϕ the Lipschitz constant of ϕ . The Lipschitz metric on CV_n is defined to be:

$$d(x, y) := \inf_{\phi} \log L_{\phi}$$

where the infimum is taken over all difference of markings maps. There exists a non-unique difference of markings map that realizes the infimum [9]. Since such a difference of markings map is homotopic rel vertices to a map that is linear on edges, we also use ϕ to denote the representative that realizes the infimum and is linear on edges and refer to such a map as an *optimal map* from x to y . Given an optimal map $\phi: x \rightarrow y$, the *slope* of an edge $e \in x$ associated to ϕ is the ratio of lengths $|\phi(e)|$ to $|e|$. For the remainder of the paper, we always assume the difference of markings maps are optimal.

By a *loop* or an *immersed loop* in x , we mean a free homotopy class of a map from the circle into x , or equivalently, a conjugacy class in \mathbb{F}_n . Meanwhile, we use a *simple loop* to mean a union of edges in a graph that forms a circle with no repeated vertices. Both kinds of loops can be identified with a conjugacy class in the free group, call it α . We use $|\alpha|_x$ to denote the metric length of the shortest representative of α in x , where x can be a point in CV_n or cv_n , depending on the context. It is shown [9] that if $x, y \in CV_n$ then the distance $d(x, y)$ can be computed as:

$$d(x, y) = \sup_{\alpha} \log \frac{|\alpha|_y}{|\alpha|_x}, \tag{1}$$

where the sup is over all conjugacy classes. In fact, it is shown in [9] that:

Theorem 2.1. *Given two points in Outer space x and y , the immersed loop that represents α which realizes the supremum can be taken from a finite set of sub-graphs of x of the following forms*

- *simple loops;*
- *figure-eight: an immersed loop where there is exactly one vertex with two pre-images in the circle;*
- *dumbbell: an immersed loop that crosses edges in two disjoint loops once and edges in the connecting arc twice.*

This result implies we can compute distances between two points by calculating the ratio $\frac{|\alpha|_y}{|\alpha|_x}$ for a finite set of immersed loops.

2.3. Train track structure. It is often convenient to use a difference of markings map that has some additional structure. We define

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

to be the stretch factor of a shortest immersed loop that represents α . For an optimal map $\phi: x \rightarrow y$, since it is linear on edges, one can define

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

to be the stretch factor of an edge e and define the *tension sub-graph*, x_ϕ , to be the sub-graph of x consisting of maximally stretched edges.

Let $\phi: x \rightarrow y$ be an optimal map. A *direction* at a vertex $v \in x_\phi$ is a germ of geodesic path $[0, \epsilon] \rightarrow x_\phi$ sending 0 to v . Let $D(v)$ be the set of all directions at v . Now ϕ induces a map

$$\phi_*: D(v) \longrightarrow D(\phi(v)).$$

Since it sends a geodesic $\gamma: [0, \epsilon] \rightarrow x_\phi$ to a geodesic $\phi \circ \gamma: [0, \epsilon] \rightarrow y$. Thus we have an equivalence relation on $D(v)$:

$$d \sim d' \iff \phi_*(d) = \phi_*(d').$$

A *gate* at v is an equivalence class. The size of a gate τ is denoted $|\tau|$ and is defined to be the number of directions in the equivalence class. An unordered pair $\{d, d'\}$ of distinct directions at a vertex v of x_ϕ is called a *turn*. The turn $\{d, d'\}$ is *ϕ -illegal* if d and d' belong to a same gate and is *legal* otherwise. The set of gates at x_ϕ is also called the *illegal turn structure* on x_ϕ induced by ϕ . We can also use a pair of edges incident to a given vertex to indicate a turn. That is, if \vec{e}_1 is an oriented edge ending on v and \vec{e}_2 is an oriented edge starting from v then we use $\langle \vec{e}_1, \vec{e}_2 \rangle$ to denote the turn that traverses \vec{e}_1 first and \vec{e}_2 second. If e_1 and e_2

share only one vertex then we do not even need to specify the orientation and we use the notation $\langle e_1, e_2 \rangle$. But if one or both of e_1 and e_2 start and end at the same vertex v we need to be more careful with orientation.

Definition 2.2. A *sub-gate* is a subset of directions in a gate (including the gate itself). The set of all sub-gates of $x \in CV_n$ under the difference of markings map ϕ is denoted \mathcal{T}_ϕ , or simply \mathcal{T} if the associated map is clear from the context. With this terminology, we can view an illegal turn as a sub-gate of size 2. A *speed assignment* is an assignment of non-negative real numbers s_τ to all elements of \mathcal{T}_ϕ of size 2 and denote the assignment

$$\mathcal{S} = \{s_\tau : \tau \in \mathcal{T}, |\tau| = 2\}.$$

For a gate $\tau = d_1, d_2, d_3$, we might have $s_{\{d_1, d_2\}}$ and $s_{\{d_2, d_3\}}$ are larger than $s_{\{d_1, d_3\}}$. But if we identify the edges associated to d_1 and d_2 as well as d_2 and d_3 along some segment of size ϵ , then the edges associated to d_1 and d_3 are also identified. To address this issue, we modify the assignments of \mathcal{S} .

For $s > 0$, we say directions d and d' at a point v are s -equivalent if there is a sequence of direction $d = d_1, \dots, d_k$ and gate $\tau_1, \dots, \tau_{k-1}$ so that $\{d_i, d_{i+1}\} = \tau_i$ and $s_{\tau_i} \geq s$. We say a speed assignment is *coherent* if for every gate $\tau = d_1, \dots, d_\ell$ and every $s > 0$

$$d d_i, \text{ and } d_j \text{ are } s\text{-equivalent} \implies s_{\{d_i, d_j\}} \geq s \quad \text{for all } 1 \leq i, j \leq k.$$

For every speed assignment \mathcal{S} there is a coherent speed assignment \mathcal{S}^c where the s equivalent classes are identical. In fact, we can define

$$s_\tau^c = \sup\{s > 0 \mid \text{every } d, d' \in \tau \text{ are } s\text{-equivalent}\}.$$

Note that replacing s_τ with s_τ^c does not change the s -equivalent classes since we are increasing the speed only for directions that are already s -equivalent. We call \mathcal{S}^c the *associated coherent speed assignment* to \mathcal{S} .

Remark 2.3. The process of making a speed assignment coherent resembles Gromov’s process of approximation of a hyperbolic space by trees. The condition of coherence of a speed assignment is really a 0-hyperbolicity condition (see Section 6, [10]).

For an immersed loop $\alpha_x \in x$, the set of illegal turns of α is denoted $\mathcal{T}_\alpha(\phi)$, or \mathcal{T}_α when the associated map is clear from the context. Since α is immersed, \mathcal{T}_α is a priori a multi-set because an illegal turn $\tau = \{d, d'\}$ can appear in \mathcal{T}_α more than once.

An illegal turn structure is moreover a *train track* structure if there are at least two gates at each vertex. For any two points $x, y \in CV_n$, there exists an optimal map $\phi : x \rightarrow y$ such that x_ϕ has a train track structure [9].

2.4. Folding paths. In this section we construct a family of paths called folding paths. The general definition can be found in [4], but we introduce it here in the language that is adapted to this paper. Assume there is an optimal difference of markings map $\phi: x \rightarrow y$ with $x_\phi = x$ that induces a train track structure on x [4, Proposition 2.1]. An immersed path in $x_\phi = x$ is *legal* if whenever it passes through a vertex, the entering and the exiting gates are distinct. A *legal loop* in a graph is a legal immersed path whose last vertex coincides with the first vertex on the path. A *legal segment* is a legal immersed path whose last vertex does not coincide with its first vertex.

Given this train track structure and an associated speed assignment $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$, we define a *folding segment*, $\{x_t\}$, for small $t \geq 0$. The optimal difference of markings map $\phi_t: x \rightarrow x_t$ is a composition of a quotient map $\bar{\phi}_t: x \rightarrow \bar{x}_t$ and a scaling map $\bar{x}_t \rightarrow x_t$. For t small enough, the quotient graph \bar{x}_t of x is obtained from x as follows. For every sub-gate τ at vertex v with $|\tau| = 2$ and two points u, w on the two edges e_u, e_w of the sub-gate τ , we identify u and w if $|v, u|_x = |v, w|_x \leq t s_\tau$ (here $|\cdot, \cdot|$ measures the length of the segment in the graph x). The graph \bar{x}_t inherits a natural metric so that this quotient map is a local isometry on each edge of x . Since $\bar{x}_t \in \text{cv}_n$, let x_t be the projective class of \bar{x}_t in CV_n , and let $\phi_t: x \rightarrow x_t$ be the composition $\bar{\phi}_t$ and the appropriate scaling.

Notice that in x_t it is possible for edges in a sub-gate τ to be identified along a segment that is longer than $t s_\tau$, depending on the identification of other edges in the gate containing τ . However we still say x_t constructed this way is the folding path associated with $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$. That is, different speed assignments may result in the same folding path.

We assume t_1 is small enough so that the combinatorial type of x_t does not change on the interval $(0, t_1)$. Also, for t small enough, any u and w as above are also identified under ϕ because e_u and e_w are in the same gate. Hence, $\phi \circ \phi_t^{-1}$ is a well-defined map. We always assume that t is small enough, so that $\phi \circ \phi_t^{-1}$ is well-defined. Given that, we now define the *left-over map* ψ_t at time t to be defined by

$$\psi_t: x_t \longrightarrow y, \quad \psi_t = \phi \circ \phi_t^{-1}$$

Note that the norm of the derivatives of the map $\phi: x \rightarrow y$ is constant along x because it is an optimal map and $x_\phi = x$ hence ϕ stretches x by the same amount everywhere. The same is true for ϕ_t . Therefore, the norm of the derivative of $\psi_t: x_t \rightarrow y$ is also constant along x_t and in fact it is the ratio of the norm of the derivative of ϕ over that of ϕ_t . That is,

$$d(x_t, y) = d(x, y) - d(x, x_t).$$

We call such a path $\{x_t\}$ a *geodesic starting from x towards y* since it does not necessarily reach y . To summarize, we have shown:

Proposition 2.4. *Assume a difference of markings map $\phi: x \rightarrow y$ gives a train track structure on x . For any speed assignment $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$, there is $t_1 > 0$ and a geodesic $\gamma_{\mathcal{S}}: [0, t_1] \rightarrow \text{CV}_n$ starting from x towards y where the graph $x_{t_1} = \gamma(t)$ is obtained by folding every gate τ at speed s_τ .*

Note that when we say the combinatorial type of x_t does not change, it does not mean it is the same as x or even x_{t_1} . Typically, the geodesic segment $\gamma_{\mathcal{S}}$ starts from x which lies on the boundary of simplex in CV_n (not necessarily of maximal dimension) travels in the interior of this simplex and stops when it hits the boundary of the simplex. At this point, if there is a new speed assignment, the folding could continue.

Globally, a *folding path* is a concatenation of such folding segments. Note that at the end point of a folding segment, we have a difference of markings map $\phi_{t_1}: x_{t_1} \rightarrow y$ and still $x_{\phi_{t_1}} = x_{t_1}$. Then ϕ_{t_1} defines a train-track structure on x_{t_1} and, choosing speed assignments, we can continue the folding path further. It may be the case that a folding path is divided into infinitely many folding segments. See Section 5 for more details on how a global folding path can be constructed from folding segments.

In general, a folding path from x to y is denoted $[x, y]_f$. If all values of the speed assignment are equal at every time then the path is called a *greedy folding path* and denoted $[x, y]_{\text{gf}}$. In Section 3 we construct a specific type of folding path whose speed assignment reflects the contribution of each sub-gate to the total length loss along the path. However, one has to be careful to extend the local construction described here to a geodesic connecting x to y . In Section 5 we extend the local construction to a global construction.

2.5. Standard geodesic. Another important class of geodesic paths to consider is the class of standard geodesic paths. For two points $x, y \in \text{CV}_n$, there may not exist a folding path connecting them. There is, however, a non-unique *standard geodesic*, denoted $[x, y]_{\text{std}}$, from x to y [4]. In [4, Proposition 2.5], Bestvina and Feighn give a detailed construction of such a standard geodesic, which we summarize briefly here. First, take an optimal map $\phi: x \rightarrow y$ and consider the tension sub-graph x_ϕ . Let $\Delta_x \subset \text{CV}_n$ denote the smallest simplex containing x . By shortening some of the edges outside of x_ϕ (and rescaling to maintain total length 1), one may then find a point $x' \in \Delta_x$ in the closed simplex Δ_x (and a scaling map $\phi_{\text{scale}}: x \rightarrow x'$) together with an optimal difference of markings $\phi': x' \rightarrow y$ whose tension graph $x'_{\phi'}$ is all of x' and such that

$$d(x, y) = d(x, x') + d(x', y).$$

If γ_1 denotes the linear path in Δ_x from x to x' (which when parameterized by arc length is a directed geodesic) and $\gamma_2 = \gamma^{\phi'}$ denotes the folding path from x' to y induced by ϕ' , it follows from the equation above that the concatenation $\gamma_1\gamma_2$ is

a directed geodesic from x to y which is called a standard geodesic from x to y , which we denote $[x, y]_{\text{std}}$.

2.6. Unique geodesics. We would like to show that, in certain situations, all geodesics connecting a pair of points have some property. Here, we give a criterion for when the geodesic between two points is unique. For a difference of markings map $\phi: x \rightarrow y$, we define a *yo-yo* to be an illegal turn $\langle e, \bar{e} \rangle$, $e \neq \bar{e}$, at a vertex v induced by ϕ satisfying the following (see Figure 1):

- the edge \bar{e} forms a loop at v ;
- there are no other edges attached to v .




Figure 1. A yo-yo illegal turn.

We say a folding path $\gamma_r: [a, b] \rightarrow \text{CV}_n$ is *rigid* if, for every $t \in [a, b]$, there is a difference of markings map $\phi_t: \gamma_r(t) \rightarrow y$ that induces a train-track structure which has exactly one illegal turn, and that illegal turn is not a yo-yo. We show that unique geodesics are exactly rigid folding paths.

Theorem 2.5. *For points $x, y \in \text{CV}_n$, where $n \geq 3$, the geodesic from x to y is unique if and only if there exists a rigid folding path γ_r connecting x to y .*

Proof. Let $\gamma_r: [0, a] \rightarrow \text{CV}_n$ be a rigid folding path connecting x to y . This implies, in particular, that there is a difference of markings map $\phi_r: x \rightarrow y$ where tension sub-graph of ϕ_r is all of x and where the train-track structure associated to ϕ_r has one, non-yo-yo illegal turn. We need the following combinatorial statement.

Claim. *Every edge, and every legal segment $P = \{e_1, e_2\}$ in x is a subpath of an immersed ϕ_r -legal loop α in x .*

Proof of Claim. Let $\tau = \langle e, \bar{e} \rangle$ at v denote the only illegal turn in x . We first address the case when e and \bar{e} are the same edge. Since every vertex of the graph x has degree two or higher, every edge or length-2 legal edge path is in an immersed loop. For this immersed loop to be legal, we need to check that this loop does not go around e twice or more consecutively. If the loop does go around e twice in a row, we modify the loop to go around e only once. The modified loop still contains the edge or the length-2 legal edge path we began with. Furthermore, it is now legal since it does not traverse the only illegal turn in the graph. Thus we have established the claim in the case $e = \bar{e}$.

Now assume $e \neq \bar{e}$. The graph $x \setminus e$ either has a vertex of degree one, or every vertex of $x \setminus e$ has degree at least two. In the first case, e, \bar{e} forms a yo-yo, which contradicts the assumption. In the second case, since every turn in $x \setminus e$ is ϕ_r -legal and every vertex has two or more gates, every edge and every length-2 legal segment is part of an immersed legal loop. Similarly, in $x \setminus \bar{e}$, every edge and every length-2 legal segment is part of an immersed legal loop. Since every edge is contained either in $x \setminus e$ or in $x \setminus \bar{e}$, it is part of an immersed ϕ_r -legal loop in x .

Given a legal segment $P = \{e_1, e_2\}$ of length 2, one of the following holds:

- $P \subset x \setminus e$;
- $P \subset x \setminus \bar{e}$;
- $\{e_1, e_2\} = \{e, \bar{e}\}$, and the segment P starts and ends at vertex v ;
- $\{e_1, e_2\} = \{e, \bar{e}\}$, and the segment P starts at v , and ends at a different vertex w .

For the first two cases, we have already established that P is a subpath of an immersed ϕ_r -legal loop α in x . In the third case, without loss of generality, suppose $\langle \bar{e}_1, \bar{e}_2 \rangle$ is legal, and $\langle \bar{e}_2, \bar{e}_1 \rangle$ is the only illegal turn in x , occurring at the vertex v . Consider the graph that is left after deleting e_1 and e_2 (but keeping the end vertices). If this graph is disconnected, consider the component that contains v . Since $n \geq 3$, the graph $x \setminus \{e_1, e_2\}$ is non-trivial and our chosen component has to have a nontrivial loop β which is legal. Let ω be a path connecting v to β not containing e_1 or e_2 . The the loop $\alpha = \bar{e}_1 \bar{e}_2 \omega \beta \omega^{-1}$ is a legal loop containing P .

For the fourth case, we can assume e_1 starts and ends at v and e_2 starts at v and ends in w . Then there has to be another edge e_3 connecting v to a vertex u otherwise, $\{e, \bar{e}\}$ forms a yo-yo. If u can be connected to w , then this path and P form a legal loop. Otherwise, we find an immersed loop passing through w and an immersed loop passing through u . Then a concatenation of these two loops, two copies of e_2 and e_3 each and the loop e_1 forms an immersed legal loop containing P . This finishes the proof. \triangle

Let z be a point that lies on a possibly different geodesic connecting x to y , that is

$$d(x, z) + d(z, y) = d(x, y). \tag{2}$$

Let $\phi: x \rightarrow z$ be a difference of markings map that gives rise to a standard path γ_{std} . We decomposed $\phi = \phi_1 \circ \phi_2$, where $\phi_1: x \rightarrow w$ represents the scaling segment of the standard path. Assume that the tension graph of ϕ is not all of x and consider an edge $e \notin x_\phi$.

By the claim, there exists a ϕ_r -legal immersed loop α containing e . Then

$$|\alpha|_x e^{d(x,y)} = (|\alpha|_x e^{d(x,z)}) e^{d(z,y)} < |\alpha|_z e^{d(z,y)} \leq |\alpha|_y. \tag{3}$$




Figure 2

But α is ϕ_r -legal, hence

$$|\alpha|_x e^{d(x,y)} = |\alpha|_y.$$

This is a contradiction. Thus $x_\phi = x$, which implies ϕ_1 is degenerate, $w = x$ and $\phi = \phi_2$. That is, there is a folding path γ connecting x to z .

Let $u = \gamma(s)$ be the last point along γ where γ_r and γ agree, let $\psi_{r,s}: u \rightarrow y$ be the leftover difference of markings map associated to γ_r and $\psi_s: u \rightarrow z$ be the left over difference of markings map associated to γ . That is, the path $\gamma|_{[0,s]}$ is a (possibly degenerate) sub-path of γ_r , but at u , there is a ψ_s -illegal turn τ that is $\psi_{r,s}$ -legal. Consider the segment P consists of the pair of edges that form τ . This segment is legal in $\psi_{r,s}$, and hence by the claim, there exists a $\psi_{r,s}$ -legal immersed loop α containing P .

That is, α is not stretching maximally from $\psi_s(s)$ to $\psi_s(s + \epsilon)$ and an identical to the argument above gives a contradiction. Thus, the path ψ_s is subpath of γ_r and z lies on γ_r .

We now show the other direction, that is, we establish that the uniqueness of a geodesic implies that it is a rigid folding path. Consider a standard geodesic $[x, y]_{\text{std}}$ from x to y . Again, the path $[x, y]_{\text{std}}$ is by definition a concatenation of a rescaling path γ_1 and a folding path γ_2 .

Suppose that, on γ_2 there are two illegal turns at some point. Then [9] shows that folding the two illegal turns at different speeds renders different geodesic paths, hence obstructing uniqueness. Otherwise, suppose γ_2 contains a yo-yo at some time s_0 . Then a segment $\gamma_2|_{[s_0,s_1]}$ can be replaced with a different geodesic $\hat{\gamma}_2$. Referring to Figure 3, the geodesic $\gamma_2|_{[s_0,s_1]}$ is obtained from folding the yo-yo, labeled τ from $\gamma_2(s_0)$ to $\gamma_2(s_1)$. However, one can also fold the edge e first at σ_1 to the point $\hat{\gamma}_2(s)$, $s \in (s_0, s_1)$, and then fold σ_2 to reach $\gamma_2(s_2) = \hat{\gamma}_2(s_2)$. Thus, the geodesic $\gamma_2|_{[s_0,s_1]}$ is not a unique geodesic connecting its end points.

Consider now the segment γ_1 . Suppose there is more than one edge that is not in x_ϕ . Then we can choose how fast to rescale the lengths of these edges rendering multiple geodesics with same endpoints as γ_1 . Otherwise, suppose e is the only edge that is not in x_ϕ . Similar to the paths illustrated in Figure 3, e can be folded onto one of its neighboring edges in a zigzag manner such that the end graph is isomorphic to the endpoint of γ_1 . Thus γ_1 is never a unique geodesic connecting its endpoints, unless it is degenerate, that is $\gamma = \gamma_2$.




Figure 3. A yo-yo illegal turn gives rise to two geodesic paths.

To sum up, for a standard geodesic to be unique, γ_1 is necessarily degenerate and γ_2 contains only one non-yo-yo illegal turn at any point. That is to say, it is a rigid folding path. \square

Remark 2.6. Note that the first part of the proof still works for rank $n = 2$. That is, if points x and y are connected via a rigid folding path, that path is the unique geodesic from x to y . However, in CV_2 , even when there is a yo-yo, the geodesic is unique. In fact, the paths γ_2 and $\hat{\gamma}_2$ described in Figure 3 are identical in CV_2 .

3. Weak convexity

The purpose of this section is to prove the following:

Theorem 3.1. *Given a difference of markings map $\phi: x \rightarrow y$, $x, y \in CV_n$, where $x_\phi = x$, there exists a speed assignment \mathcal{S} defining a folding path $\gamma: [0, t_1] \rightarrow CV_n$ starting at x towards y so that, for every loop α and every time $t \in [0, t_1]$,*

$$|\alpha|_t \leq \max(|\alpha|_x, |\alpha|_y).$$

Recall that a speed assignment is a set $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$ of speeds assigned to all sub-gates of size 2. For t small enough, there is a quotient map $\bar{\phi}_t: x \rightarrow \bar{x}_t$ (that is, $\bar{\phi}_t$ is an isometry along the edges of x) where the edges in gate τ are identified along a subsegment of length $t s_\tau$. Let $|\mathcal{S}|$ be the speed at which \bar{x}_t is losing length, that is,

$$|\mathcal{S}| = \frac{1 - \text{total length of } \bar{x}_t}{t}$$

Note that this is a constant for small values of t . Also, if S^c is the associated coherent speed assignment, then $|\mathcal{S}| = |S^c|$ because they define the same quotient graph \bar{x}_t .

Lemma 3.2. *Let $[x, y]_f$ be a folding path associated to a difference of markings map $\phi: x \rightarrow y$ and a coherent speed assignment $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$. Then, for every loop α , the derivative of the length of α along this path equals*

$$|\dot{\alpha}|_t = |\alpha|_t - 2 \sum_{\tau \in \mathcal{T}_\phi(\alpha)} \frac{s_\tau}{|\mathcal{S}|} \tag{4}$$

where the derivative is taken with respect to distance in CV_n .

Proof. For every $\tau \in \mathcal{T}_\alpha$, there are two sub-edges of α of length $t s_\tau$ that are identified under the quotient map $\bar{\phi}_t: x \rightarrow \bar{x}_t$. And x_t is obtained from \bar{x}_t by a scaling of factor $\frac{1}{1-t|\mathcal{S}|}$. Since \mathcal{S} is coherent, the length loss of the loop α associated for to every $\tau \in \mathcal{T}_\alpha$ is exactly $2t s_\tau$. Hence

$$|\alpha|_t = \frac{|\alpha|_x - 2t \sum_{\tau \in \mathcal{T}_\alpha} s_\tau}{1 - t |\mathcal{S}|} \tag{5}$$

and, for $s > t$,

$$|\alpha|_s - |\alpha|_t = \frac{(s - t) |\mathcal{S}| |\alpha|_x - 2(s - t) \sum_{\tau \in \mathcal{T}_\alpha} s_\tau}{(1 - s |\mathcal{S}|)(1 - t |\mathcal{S}|)}.$$

Also $d(x_t, x_s) = \log \frac{(1-s|\mathcal{S}|)}{(1-t|\mathcal{S}|)}$. That is, when $(s - t)$ is small,

$$d(x_t, x_s) = \log \left(\frac{1 - t|\mathcal{S}|}{1 - s|\mathcal{S}|} \right) = \log \left(1 + \frac{(s - t)|\mathcal{S}|}{1 - s|\mathcal{S}|} \right) \sim \frac{(s - t)|\mathcal{S}|}{1 - s|\mathcal{S}|}.$$

Therefore,

$$|\dot{\alpha}|_t = \lim_{s \rightarrow t} \frac{|\alpha|_s - |\alpha|_t}{d(x_t, x_s)} = \frac{|\alpha|_x - 2 \sum_{\tau \in \mathcal{T}_\alpha} \frac{s_\tau}{|\mathcal{S}|}}{1 - t |\mathcal{S}|}. \tag{6}$$

On the other hand, replacing, $|\alpha|_t$ in the right-hand side of (4) with the expression in equation (5), we get

$$|\alpha|_t - 2 \sum_{\tau \in \mathcal{T}_\phi(\alpha)} \frac{s_\tau}{|\mathcal{S}|} = \frac{|\alpha|_x - 2t \sum_{\tau \in \mathcal{T}_\alpha} s_\tau}{1 - t |\mathcal{S}|} - 2 \sum_{\tau \in \mathcal{T}_\phi(\alpha)} \frac{s_\tau}{|\mathcal{S}|} = \frac{|\alpha|_x - 2 \sum_{\tau \in \mathcal{T}_\alpha} \frac{s_\tau}{|\mathcal{S}|}}{1 - t |\mathcal{S}|}. \tag{7}$$

The right hand sides of equation (6) and equation (7) are the same, hence the left hand sides are equal. But this is what was claimed. \square

For given x and y , our goal is to find an appropriate speed assignment so that, for every loop α , if $|\alpha|_y \leq |\alpha|_x$ then $|\dot{\alpha}|_{t=0} \leq 0$. To this end, we will define values ℓ_τ that quantifies the contribution of sub-gate τ to the total length loss from x to y and then use these to define a speed assignment.

Notation 3.3. For the remainder of this section, let $\bar{y} \in \text{cv}_n$ be the representative in the projective class of y so that the associated change of markings map $\bar{\phi}: x \rightarrow \bar{y}$ restricted to every edge is a length preserving immersion. Also, let $\Phi: T_x \rightarrow T_{\bar{y}}$ be a lift of $\bar{\phi}$.

Consider a point $p \in T_{\bar{y}}$ and let $\text{Pre}(p) \subset T_x$ denote the set of pre-images of p under the map Φ and let $\text{CH}(p)$ denote the convex hull of $\text{Pre}(p)$ in T_x . Assume p is generic, and thus, $\text{Pre}(p)$ does not contain any vertex in T_x . We give $\text{CH}(p)$ a tree structure where there are no degree 2 vertices; some edges of $\text{CH}(p)$ may consist of several edges in T_x . The tree $\text{CH}(p)$ also inherits its illegal turn structure from T_x , however, an edge of $\text{CH}(p)$ may contain one or more illegal turns. Also, note that since all endpoints of $\text{CH}(p)$ map to p , $\text{CH}(p)$ does not contain any legal path connecting its end vertices.

We denote the set of all sub-gates of $\text{CH}(p)$ by Θ . For each sub-gate $\sigma \in \Theta$, we assign a weight $c(\sigma, p)$ to σ which measures how much of the branching of $\text{CH}(p)$ is due to σ .

Proposition 3.4. For x, \bar{y}, p and Θ as above, and for any $\sigma \in \Theta$, there exists a weight assignment $c(\sigma, p) > 0$ such that

$$\sum_{\sigma \in \Theta} c(\sigma, p) = |\text{Pre}(p)| - 1 \tag{8}$$

Furthermore, let α be a path in $\text{CH}(p)$ with end points in $\text{Pre}(p)$ and let $\Theta_\alpha \subset \Theta$ be the set of gates associated to the illegal turns appearing along α . Then

$$\sum_{\sigma \in \Theta_\alpha} \sum_{\hat{\sigma} \supseteq \sigma} \frac{c(\hat{\sigma}, p)}{|\hat{\sigma}| - 1} \geq 1. \tag{9}$$

Proof. Note that all vertices of degree 1 in $\text{CH}(p)$ are in $\text{Pre}(p)$. But some points in $\text{Pre}(p)$ may lie on the interior of an edge of $\text{CH}(p)$. First, we cut $\text{CH}(p)$ along these points to decompose $\text{CH}(p) = \sqcup_{i=1}^m T_i$. The vertices of degree 1 in T_i are exactly $\text{Pre}(p) \cap T_i$ and T_i is the convex hull of $\text{Pre}(p) \cap T_i$. This decomposes $\Theta = \sqcup_{i=1}^m \Theta_i$.

For a tree $T_i, i = 1, \dots, m$, a vertex is an *outer vertex* if it has degree 1 and an edge is an *outer edge* if one of its vertices has degree one. All other vertices and edges are called inner vertices and inner edges. It follows that each tree T_i has the property that no legal path joins outer vertices.

Note that if equation (9) holds for each T_i it also holds for $\text{CH}(p)$. Indeed, if α is a path in $\text{CH}(p)$ with endpoints in $\text{Pre}(p)$ then the restriction of α to some T_i is non-empty. Then, the sum of the weights along the whole path is larger than the sum of the weights along the subpath that is contained in T_i .

Similarly, if we show

$$\sum_{\sigma \in \Theta_i} c(\sigma, p) = |\text{Pre}(p) \cap T_i| - 1. \tag{10}$$

we can conclude equation (8). In fact,

$$\sum_{i=1}^m |\text{Pre}(p) \cap T_i| = |\text{Pre}(p)| + (m - 1) \tag{11}$$

This is because the number of points in the interior of $\text{CH}(p)$ is $(m - 1)$ and these points are counted twice on the left hand side of the above equation. We now have

$$\begin{aligned} \sum_{\sigma \in \Theta} c(\sigma, p) &= \sum_{i=1}^m \sum_{\sigma \in \Theta_i} c(\sigma, p) \\ &= \sum_{i=1}^m (|\text{Pre}(p) \cap T_i| - 1) && \text{(equation (10))} \\ &= \sum_{i=1}^m (|\text{Pre}(p) \cap T_i|) - m \\ &= |\text{Pre}(p)| + (m - 1) - m = |\text{Pre}(p)| - 1. && \text{(equation (11))} \end{aligned}$$

Therefore it is sufficient to show that there is a weight assignment for each Θ_i . We do this inductively. Namely, we will show:

Claim 1. *Let T be a finite tree with a given illegal turn structure with the property that there are no legal paths joining two different outer vertices of T . Let Θ be the set of gates in T . Then there is a weight assignment such that equation (9) and equation (10) hold.*

The base case is when T has one inner vertex v which is a one-gate vertex (denote the gate by σ) and all the outer edges contain no illegal turns. In this case, we define $c(\sigma, p) = |\sigma| - 1$ and $c(\tau, p) = 0$ for all sub-gates τ of σ . In general, for any sub-gate where we do not specifically assign a weight, the weight assignment is assumed to be zero. In this case, equation (9) and equation (10) clearly hold.

Note that if there is only one inner vertex, it is necessarily a one-gate vertex otherwise there is a legal path joining outer vertices of T . Hence, if we are not in the base case, we either have more than one inner vertex or some outer edge has an illegal turn.

Case 1. Assume there is an edge e of T that contains illegal turns $\sigma_1, \dots, \sigma_k$. Then we define $c(\sigma_i, p) = \frac{1}{k}$ and we remove this edge from T . We also remove any vertices that have degree two to obtain a tree T' . In T' each edge is again a topological edge, consistent with our initial condition. Also, T' still has the property that it does not contain any legal paths connecting its end vertices. To see this, notice that since we removed exactly one topological edge, the degree of the vertex at which we removed this edge is still two or higher. This implies

we did not create a new leaf by removing one edge, which means a legal path the would have appeared after this step already exists before the step. But that contradicts our assumption. By induction, there is a weight assignment for gates in T' satisfying equation (10) and equation (9). Equation (10) still holds for T because total weights assigned was 1 and the number of end vertices of T is reduced exactly by one. To see equation (9), let α be a path joint end vertices of T . If α traverses e then equation (9) holds since the weight along e already add up to 1. Otherwise, α is a path in T' and equation (9) holds by induction.




Figure 4. two sub-cases of Case 2.

Case 2. Assume that no outer edge of T contains any illegal turn. We claim that there is a vertex v of T and a sub-gate σ at v so that σ contain all but one edge incident to v and all edges of σ are outer edges (the remaining edge can be either inner or outer). To see this, consider the longest embedded path v_0, v_1, \dots, v_m in T . Then, all but one of the edges incident to v_1 are outer edges. Otherwise, the path can be made longer. In fact, all these outer-edges have to be in some sub-gate σ . Otherwise, there is a legal path connecting two outer vertices. This proves the claim.

There are two sub-cases. If all edges incident to v are in the same gate τ (which contains σ and the remaining edge) then we define $c(\tau, p) = |\tau| - 1$ and we remove v and all edges incident to v to obtain T' (this is consistent with what we did in the base case in which case T' would be empty). The tree T' still has the property it does not contain any legal path joining its end vertices because any such path could be extended to a legal path in T hence the assumption of induction applies. The number of ends of T goes down by $|\sigma| = |\tau| - 1$. Therefore, since equation (10) holds for T' by induction, it also holds for T . Also, for any path α that passes through v , the sum in equation (9) is at least 1 and any other path is contained in T' . Hence, equation (9) also holds for T .

Otherwise, there are two gates at v , namely σ and a gate with one edge. In this case, we define $c(\sigma, p) = |\sigma| - 1$ and delete all the edges associated to σ from T to obtain T' . The vertex v survives in T' and the number of ends of the tree is reduced by $|\sigma| - 1$. Again, since equation (10) holds for T' by induction, it also holds for T and the tree T' still has the property it does not contain any legal path joining its end vertices because any such path could be extended to a legal path in T . Let α be a path in T joining its end points. If α is a two edge path consisting

of two edges in σ , the equation (9) holds. If α does not traverse v then α is in T' and equation (9) holds by induction. If α passes through v but contains one edge in σ , then let α' be α minus this edge. Then α' is a path in T' joining its end points and, by induction, the sum in equation (9) associated to α' is at least 1. But this is the lower bound for the sum associated to α since α' is a sub-path of α . This finishes the proof. \square

Remark 3.5. The result of the algorithm is not unique. Different gates may be assigned different values depending on the order in which we remove the outer edges that contain any illegal turn. However, we make these choices for every p once and for all so that the following holds.

- (1) Since p is assumed to be generic, there is an open interval containing p so that $\text{CH}(q)$ has the same combinatorics as $\text{CH}(p)$ for all q near p . We apply the algorithm for all points in this interval simultaneously to ensure $c(\sigma, p)$ is a locally constant function almost everywhere.
- (2) We make choices so that $c(\sigma, p)$ is equivariant. If p' is in the orbit of p , $\text{CH}(p)$ and $\text{CH}(p')$ have the same combinatorics. We make sure the algorithm is applied the same way for an equal length interval around p and every point in the orbit of p .

As a result, $c(\sigma, p)$ is well defined for almost every point p , and it is equivariant and locally constant almost everywhere. In particular, for every σ , $c(\sigma, p)$ is an integrable function.




Figure 5. $c(\sigma, p)$ is computed iteratively by applying Step 1 as many times as possible and then apply Step 2, and then repeat.

Example 3.6. The example in Figure 5 illustrate the definition of $c(\sigma, p)$. Suppose $\text{CH}(p)$ is as shown in the leftmost graph, with seven outer edges and seven gates, marked $\{\sigma_1, \sigma_2, \dots, \sigma_7\}$. First apply Step 1. There are two possible candidates for Step 1. The illegal turns that occur as part of a topological edge are σ_6 and σ_7 , which belong to different topological edges. It does not matter which one of them is assigned first. After applying Step 1 twice, we get

$$c(\sigma_6, p) = 1 \text{ and } c(\sigma_7, p) = 1.$$

After deleting these two topological edges, σ_5 no longer exists, therefore

$$c(\sigma_5, p) = 0.$$

Next, Step 2 picks out either σ_1 or σ_4 . Suppose we start with σ_4 , which contains 3 outer edges, thus

$$c(\sigma_4, p) = 3 - 1 = 2.$$

After deleting these three outer edges of σ_4 , we apply Step 1 again and observe that there is a new topological edge with two illegal turns σ_2 and σ_3 . Thus

$$c(\sigma_2, p) = c(\sigma_3, p) = \frac{1}{2}.$$

And the topological edge is deleted at the end of the step. Lastly we have only one gate with two outer edge, thus

$$c(\sigma_1, p) = 1.$$

It can be verified that

$$\sum_{\sigma \in \Theta} c(\sigma, p) = 1 + 1 + 0 + 2 + \frac{1}{2} + \frac{1}{2} + 1 = 6 = |\text{Pre}(p)| - 1.$$

We now use the $c(\sigma, p)$ functions to define the length loss functions:

$$\ell_\sigma = \int_{T_{\bar{y}}} c(\sigma, p) dp$$

where the integral is taken with respect to the length in $T_{\bar{y}}$ (the universal cover of \bar{y} , see Notation 3.3). Note that $c(\sigma, p)$ is defined only for a generic p but the integral is still defined. Even though this is an integral over the whole tree which has infinite measure, for every σ , the set of points p where $c(\sigma, p)$ is non-zero is contained in a compact set and hence has finite measure. Also, since our construction is equivariant, for any sub-gate $\tau \in \mathcal{T}$ (Recall \mathcal{T} is the set of all sub-gates in x under the map ϕ .) we can define $\ell_\tau = \ell_\sigma$ where σ is any lift of τ to T_x .

The number ℓ_τ represents how much of the length loss from x to \bar{y} we are attributing to the sub-gate τ . In particular, we have

Lemma 3.7. For $\bar{\phi}: x \rightarrow \bar{y}$ and ℓ_τ defined as above,

$$\sum_{\tau \in \mathcal{T}} \ell_\tau = |x| - |\bar{y}|.$$

Proof. We denote points in \bar{y} by q and $\text{Pre}(q)$ represents the pre-image of q under $\bar{\phi}$. Since the map $\bar{\phi}$ is locally a length preserving immersion, we have

$$1 = |x| = \int_{\bar{y}} |\text{Pre}(q)| dq.$$

Let $T_0 \subset T_{\bar{y}}$ be a tree that is a fundamental domain of action \mathbb{F}_n on $T_{\bar{y}}$ and let Θ_0 be a finite subset of Θ that contains exactly one lift for every $\tau \in \mathcal{T}$. Now,

$$\begin{aligned} \sum_{\tau \in \mathcal{T}} \ell_{\tau} &= \sum_{\sigma \in \Theta_0} \ell_{\sigma} = \sum_{\sigma \in \Theta_0} \int_{T_{\bar{y}}} c(\sigma, p) dp \\ &= \sum_{g \in \mathbb{F}_n} \sum_{\sigma \in \Theta_0} \int_{T_0} c(\sigma, g(p)) dp && (T_{\bar{y}} = \cup_g g(T_0)) \\ &= \sum_{g \in \mathbb{F}_n} \sum_{\sigma \in \Theta_0} \int_{T_0} c(g^{-1}(\sigma), p) dp && (c(\cdot, \cdot) \text{ is equivariant}) \\ &= \int_{T_0} \sum_{\sigma \in \Theta} c(\sigma, p) dp && (\Theta = \cup_g g(\Theta_0)) \\ &= \int_{T_0} (|\text{Pre}(p)| - 1) dp && (\text{equation (11)}) \\ &= \int_{\bar{y}} (|\text{Pre}(q)| - 1) dq = |x| - |\bar{y}|. \end{aligned}$$

which is what was claimed in the lemma. □

Next, we use length loss contributions ℓ_{τ} to define a coherent speed assignment. For a sub-gate τ with $|\tau| = 2$, define

$$s_{\tau} = \sum_{\hat{\tau} \supseteq \tau} \frac{\ell_{\hat{\tau}}}{|\hat{\tau}| - 1} \tag{12}$$

where the sum is over all sub-gates $\hat{\tau}$ containing τ . We are dividing $\ell_{\hat{\tau}}$ by $(|\hat{\tau}| - 1)$ because if you fold edges of $\hat{\tau}$ along a segment of length l , the length loss is larger by factor $(|\hat{\tau}| - 1)$. Let \mathcal{S}^c be the coherent speed assignment induced by $\{s_{\tau}\}_{|\tau|=2}$. Then \mathcal{S}^c is our desired speed assignment. Since we only work with coherent speed assignments, we omit the superscript in the rest of the paper and simply use \mathcal{S} emphasizing each time that \mathcal{S} is coherent.

Example 3.8. We illustrate the computation of s_τ with the example mentioned in the introduction. Consider $\mathbb{F}_3 = \langle a, b, c \rangle$. Let x be a rose with three petals. The three edges we refer to as e_1, e_2, e_3 . The edge e_1 is labeled ac^2 , the edge e_2 is labeled bc , the edge e_3 is labeled c . The edge lengths are $\frac{1}{2}, \frac{1}{3}, \frac{1}{6}$, respectively. The graph \bar{y} is a rose of three petals with labels $\{a, b, c\}$ and lengths $\{\frac{1}{6}, \frac{1}{6}, \frac{1}{6}\}$. The construction is such that $\phi: x \rightarrow y$ satisfies $x_\phi = x$. \bar{y} is obtained from x by wrapping ac^2 around c twice and bc around c once.

$T_{\bar{y}}$ contains three types of edges. If the point p is on an edge labeled a or b , then the pre-image contains only one copy of p , and $c(\tau, p) = 0$ for all τ . If p is on the c -edge, then $\text{CH}(p)$ is as shown in Figure 6, where the four pre-images of p are marked with a circle.




Figure 6. The $\text{CH}(p)$ of Example 3.8 where p is a point on the edge labeled c in y

$\text{CH}(p)$ has two gates. One contains a black and green edge, which we denote $\sigma_{e_1 e_3}$. The other gate contains three edges, black, green and blue, and we denote the gate $\sigma_{e_1 e_2 e_3}$. At Step 1,

$$c(\sigma_{e_1 e_3}, p) = 1;$$

at Step 2,

$$\sigma_{e_1 e_2 e_3} = 2.$$

Next, we compute the length loss function by integrating $c(\cdot, \cdot)$ over $T_{\bar{y}}$. In this case, the only non-zero component of the integral is when integrating over edge labeled c , which has length $\frac{1}{6}$, therefore:

$$\begin{aligned} \ell_{e_1 e_3} &= 1 \times \frac{1}{6} = \frac{1}{6}, \\ \ell_{e_1 e_2 e_3} &= 2 \times \frac{1}{6} = \frac{1}{3}. \end{aligned}$$

Indeed, it is the case that

$$\frac{1}{3} + \frac{1}{6} = \sum_{\sigma} \ell_{\sigma} = x - \bar{y} = 1 - \frac{1}{2} = \frac{1}{2}.$$

Based on the length loss functions we compute the folding speed of all sub-gates in x . Again we can denote a sub-gate in x by the edges in the gate, so we have

$$s_{e_1e_3} = l_{e_1e_3} + \frac{1}{2}l_{e_1e_2e_3} = \frac{1}{6} + \frac{1}{2} \times \frac{1}{3} = \frac{1}{3},$$

$$s_{e_1e_2} = s_{e_2e_3} = \frac{1}{2}l_{e_1e_2e_3} = \frac{1}{2} \times \frac{1}{3} = \frac{1}{6}.$$

That is to say, since ac^2 wraps over c twice while bc wraps over c once, infinitesimally, the folding associated with the former is twice as fast.

Lemma 3.9. *For the coherent speed assignment \mathcal{S} above, we have*

$$|\mathcal{S}| \leq \sum_{\tau \in \mathcal{T}} \ell_{\tau}.$$

Proof. We organize the argument by considering one maximal gate τ and its sub-gates only. The length losses at different gates add up, hence, it is sufficient to prove the lemma one gate at the time. For the rest of the argument, let τ be a fixed gate in \mathcal{T} .

By an ϵ -neighbourhood of a gate τ we mean the intersection of an ϵ -ball around the vertex associated to τ with the edges associated to τ . For ϵ small enough, the ϵ -neighbourhood of τ is a tree with one vertex v and $|\tau|$ edges $e_1, \dots, e_{|\tau|}$ of size ϵ . Choose $t > 0$ small enough so that, for every $1 \leq i, j \leq |\tau|$, $ts_{i,j} < \epsilon$.

The image of this ϵ -neighbourhood in \bar{x}_t is a quotient of the ϵ -neighbourhood of τ after identifying e_i and e_j along a segment of length $ts_{\{e_i, e_j\}}$ starting from the vertex v . In fact, it is enough to do $|\tau| - 1$ identifications. Namely, we choose the pair of edges e_1, e'_1 that are identified along the longest segments (that is, for $\tau_1 = \{e_1, e'_1\}$, s_{τ_1} is maximal among all sub-gates of τ of size 2). We put e_1 and e'_1 in the same group and all other edges in separate individual groups. Continue in this way for $i = 2, \dots, (|\tau| - 1)$, we choose a pair of edges e_i, e'_i from different groups so that for $\tau_i = \{e_i, e'_i\}$, s_{τ_i} is maximal among all such pairs and combine the groups associated to e_i and e'_i into one group. The maximality implies that the amount of identification along τ_i is not caused by identifications along other gates. That is, we have not changed the speed s_{τ_i} to make \mathcal{S} coherent and we still have

$$s_{\tau_i} = \sum_{\hat{\tau} \supseteq \tau_i} \frac{\ell_{\hat{\tau}}}{|\hat{\tau}| - 1}.$$

After $|\tau| - 1$ steps, we have only one group (See Figure 7).

The image of the ϵ -neighborhood of τ in x to \bar{x}_t can be obtained from the ϵ -neighborhood of τ in x by identifying e_i and e'_i along a segment of length ts_{τ_i} because any other identification between the two groups is along a smaller




Figure 7. We identify the edges of τ_i along a segment of length $t s_{\tau_i}$ to obtain the image of the ϵ -neighborhood of τ in x to \bar{x}_t .

segment. That is, setting

$$|\mathcal{S}_\tau| = \sum_{i=1}^{|\tau|-1} s_{\tau_i}$$

we have $t|\mathcal{S}_\tau|$ is the length loss in the ϵ -neighborhood of τ .

The way τ_i are chosen, every sub-gate $\tau' \subset \tau$ contains at most $(|\tau'| - 1)$ of these τ_i . This is because, after $(|\tau'| - 1)$ sub-gates of τ' are chosen in the above process, every edge in τ' is already in the same group. Therefore, letting \mathcal{J}_τ to be the set of sub-gates of τ , we have

$$\sum_{i=1}^{|\tau|-1} s_{\tau_i} = \sum_{i=1}^{|\tau|-1} \sum_{\tau' \supseteq \tau_i} \frac{\ell_{\tau'}}{|\tau'| - 1} = \sum_{\tau' \in \mathcal{J}_\tau} |\{i \mid \tau_i \subseteq \tau'\}| \cdot \frac{\ell_{\tau'}}{|\tau'| - 1} \leq \sum_{\tau' \in \mathcal{J}_\tau} \ell_{\tau'}.$$

Combining the last two equations, we have

$$|\mathcal{S}_\tau| \leq \sum_{\tau' \subseteq \tau} \ell_{\tau'}.$$

But $|\mathcal{S}| = \sum_{\tau \text{ maximal}} |\mathcal{S}_\tau|$ where the sum is over maximal gates. This finishes the proof. □

A *vanishing path* for $\Phi: T_x \rightarrow T_y$ is an immersion $v: [0, 1] \rightarrow T_x$ such that $\Phi \circ v$ is homotopic to a point relative to the endpoints. Abusing notation, we sometimes refer to the image of v as a vanishing path and write $v \subset T_x$.

Lemma 3.10. *Let $v \subset T_x$ be a vanishing path and let Θ_v be the set of sub-gates in T_x that appear along v .*

$$|v|_x \leq 2 \sum_{\sigma \in \Theta_v} s_\sigma.$$

Proof. Let $p \in \Phi(v) \subset T_{\bar{y}}$ be a generic point. Define

$$m(p) = \#\{\Phi^{-1}(p) \cap v\}.$$

Recall that, to compute the weights $c(\sigma, p)$, we first write $\text{CH}(p) = \sqcup T_i$ (Step 0) so that only the end points of each T_i are mapped to p . Then $v \cap \text{CH}(p)$ consists of $m(p) - 1$ separate segments $\alpha_1, \alpha_2, \dots, \alpha_{m(p)-1}$ where each α_i is the intersection of v with T_i . If Θ_i are the gates appearing along α_i , we claim that for each i ,

$$1 \leq \sum_{\sigma \in \Theta_i} \sum_{\hat{\sigma} \supseteq \sigma} \frac{c(\hat{\sigma}, p)}{|\hat{\sigma}| - 1}.$$

This is because, if α_i passes through an edge e of T_i with an illegal turn, then the sum

$$1 = \sum_{\sigma \in \Theta_e} c(\sigma, p) \leq \sum_{\sigma \in \Theta_i} c(\sigma, p)$$

and the claim follows. Otherwise, a sub-gate $\sigma \in \Theta_i$ is contained in a sub-gate $\bar{\sigma}$ that appears in $\text{CH}(p)$, where all the associated edges are legal, and hence, following the algorithm, $c(\hat{\sigma}, p) = |\hat{\sigma}| - 1$. Again the claim follows. This gives

$$m(p) - 1 \leq \sum_{\sigma \in \Theta_v} \sum_{\hat{\sigma} \supseteq \sigma} \frac{c(\hat{\sigma}, p)}{|\hat{\sigma}| - 1}.$$

On the other hand, we know that $|v|_x = \int_{\Phi(v)} m(p) dp$ and, since p is a generic point, we have

$$m(p) \geq 2 \quad \text{and} \quad m(p) \leq 2(m(p) - 1).$$

We now have

$$\begin{aligned} |v|_x &= \int_{\Phi(v)} m(p) dp \\ &\leq \int_{\Phi(v)} 2(m(p) - 1) dp \\ &\leq 2 \int_{\Phi(v)} \sum_{\sigma \in \Theta_v} \sum_{\hat{\sigma} \supseteq \sigma} \frac{c(\hat{\sigma}, p)}{|\hat{\sigma}| - 1} dp && \text{(using the claim)} \\ &\leq 2 \sum_{\sigma \in \Theta_v} \sum_{\hat{\sigma} \supseteq \sigma} \int_{T_{\hat{\sigma}}} \frac{c(\hat{\sigma}, p)}{|\hat{\sigma}| - 1} dp && \text{(enlarging the domain of integration)} \\ &= 2 \sum_{\sigma \in \Theta_v} \sum_{\hat{\sigma} \supseteq \sigma} \frac{\ell_{\sigma}}{|\hat{\sigma}| - 1} = 2 \sum_{\sigma \in \Theta_v} s_{\sigma}. \end{aligned}$$

And we are done. □

We are now ready to prove Theorem 3.1. Intuitively, the proof is as follows. If α passes through an illegal turn τ , folding along τ contributes to length loss along α but it also contributes to length loss in other parts of the graph x which could be larger. Since the folding speed at each illegal turn τ is proportional to the total length loss contribution associated to this gate, the sum of folding speeds at illegal gates along α is larger than total length loss along α . This is essentially the content of equation (13) below. Of course, the length of α increases along the portion of its length that does not pass through any illegal turns but, along a balanced folding path, the speed of length of α is less than the slope of length α as measured from x to y . The rest of the proof is essentially just re-parametrizing the balanced folding path with arc length to obtain the correct estimates.

Proof of Theorem 3.1. Let S be the coherent speed assignment defined after equation (12), let $t_1 > 0$ be a time for which the folding with the speed S is defined (see Proposition 2.4) and let α be any loop. Denote the geodesic representative of α in x with α_x and in \bar{y} with $\alpha_{\bar{y}}$. Also for small enough t_1 , α_x can be sub-divided to segments $u_1 \cup w_1 \cup \dots \cup u_m \cup w_m$ so that, for $i = 1, \dots, m$,

- The segments u_i are vanishing paths in x .
- The segments $\Phi(w_i)$ are immersed and $\alpha_{\bar{y}} = \cup_i \Phi(w_i)$.

In particular,

$$|\alpha|_x = |\alpha|_{\bar{y}} + \sum_i |u_i|_x.$$

Let v_i be a lift of u_i to T_x . We can assume v_i are completely disjoint from each other. Since w_i are all legal, there is a one-to-one correspondence between sub-gates in \mathcal{T}_α and in $\cup_i \Theta_{v_i}$. Thus, Lemma 3.10 implies

$$\sum_i |u_i|_x = \sum_i |v_i|_x \leq 2 \sum_i \sum_{\sigma \in \Theta_{v_i}} s_\sigma = 2 \sum_{\tau \in \mathcal{T}_\alpha} s_\tau. \tag{13}$$

We have

$$\frac{|\alpha|_x - \sum_i |u_i|_x}{|\bar{y}|} = |\alpha|_y = |\alpha|_x + (|\alpha|_y - |\alpha|_x),$$

$$|\alpha|_x - \sum_i |u_i|_x = |\alpha|_x |\bar{y}| + (|\alpha|_y - |\alpha|_x) |\bar{y}|.$$

By equation (13), we can replace $\sum_i |u_i|_x$ with $2 \sum_{\tau \in \mathcal{T}_\alpha} s_\tau$ to get

$$|\alpha|_x (1 - |\bar{y}|) - 2 \sum_{\tau \in \mathcal{T}_\alpha} s_\tau \leq (|\alpha|_y - |\alpha|_x) |\bar{y}|.$$

But $(1 - |\bar{y}|) = \sum_{\tau \in \mathcal{T}} \ell_\tau \geq |\mathcal{S}|$, therefore

$$|\alpha|_x - 2 \sum_{\tau \in \mathcal{T}_\alpha} \frac{s_\tau}{|\mathcal{S}|} \leq (|\alpha|_y - |\alpha|_x) \frac{|\bar{y}|}{1 - |\bar{y}|}.$$

Note that

$$d(x, y) = \log \frac{1}{|\bar{y}|} \implies |\bar{y}| = e^{-d(x,y)} \implies \frac{|\bar{y}|}{1 - |\bar{y}|} = \frac{1}{e^{d(x,y)} - 1}.$$

Hence, letting $V_\alpha = 2 \sum_{\tau \in \mathcal{T}_\phi(\alpha)} \frac{s_\tau}{|\mathcal{S}|}$, we have

$$|\alpha|_x - V_\alpha \leq \frac{|\alpha|_y - |\alpha|_x}{e^{d(x,y)} - 1}. \tag{14}$$

Now re-parametrize the folding path with arc-length and denote the new parameter with s . Solving the differential equation given in Lemma 3.2, we have

$$|\dot{\alpha}|_s = |\alpha|_s - V_\alpha \implies |\alpha|_s = (|\alpha|_x - V_\alpha)e^s + V_\alpha. \tag{15}$$

Note that if $|\alpha|_y \leq |\alpha|_x$, then by equation (14) the rate of change of the length is negative and $|\alpha|_t \leq |\alpha|_x$. If $|\alpha|_y \geq |\alpha|_x$ then,

$$\begin{aligned} |\alpha|_s &= (|\alpha|_x - V_\alpha)e^s + V_\alpha \\ &= (|\alpha|_x - V_\alpha)(e^s - 1) + |\alpha|_x \\ &\leq (|\alpha|_y - |\alpha|_x) \frac{e^s - 1}{e^{d(x,y)} - 1} + |\alpha|_x && \text{(equation (14))} \\ &\leq |\alpha|_y. && (s \leq d(x, y)) \end{aligned}$$

That is, in either case, $|\alpha|_t \leq \max(|\alpha|_x, |\alpha|_y)$. This finishes the proof. \square

4. Decorated difference of markings map

In Section 3, we constructed a balanced folding path starting from x towards y assuming that there is an optimal difference of markings map $\phi: x \rightarrow y$ such that $x_\phi = x$. In general, such a difference of markings map does not exist. In this section, we decorate the graph y and modify the map ϕ in a way that the tension graph becomes all of x . This modified difference of markings map is called a *decorated difference of markings map*. Recall that, given $\phi: x \rightarrow y$, one can construct a standard geodesic path $[x, y]_{\text{std}} \in \text{CV}_n$ where first every edge e outside of x_ϕ is shortened. For a decorated difference of markings map, we instead create an illegal turn in the interior of every such edge. The folding of that illegal turn effectively shortens the length of e . But this can be done simultaneously with folding of other illegal turns hence ensuring that a version of Theorem 3.1 still holds.

4.1. Decorating the graphs. Consider a pair of points $x, y \in \text{CV}_n$. Consider an optimal map $\phi: x \rightarrow y$. Recall, from Section 2.5, that $\phi = \phi' \circ \phi_{\text{scale}}$ where $\phi_{\text{scale}}: x \rightarrow x'$ is a scaling map between x and some $x' \in \Delta_x$ and $x'_{\phi'}$ is the whole graph x' . First, we equip x with the train track structure coming from ϕ' . We then add further illegal turns as follows (See Figure 8).

Let e be an edge outside of x_ϕ , say connecting v_0 to v_3 . Add two subdividing vertices v_1 and v_2 to e such that the following holds: Let $e_{i,j}$ denote the edge with end vertices v_i and v_j . Then, we require that

$$|e_{0,1}|_x = |e_{1,2}|_x, \quad \text{and} \quad \lambda |e_{2,3}|_x = |\phi(e)|_y.$$

Where $\lambda = L_\phi$ is such that $\log \lambda = d(x, y)$. This is always possible since $\lambda |e|_x > |\phi(e)|_y$. We call v_1 and v_2 *pseudo-vertices* and refer to the graph x with all the pseudo-vertices added as x^d .




Figure 8. Decoration of edges in x and y .

Next we decorate y . For every $e \in x$ and vertex $v_0 \in x$ as above, we attach a new edge of length $\lambda |e_{0,1}|_x$ to y at the point $\phi(v_0) \in y$. The other end of this new edge is incident to an added degree-one vertex. Thus we have added a *hair* at $\phi(v_0)$. We denote the resulting decorated graph by y^d . By *core of y^d* we mean the graph obtained from y^d after removing all the hair. Note that the core of y^d still has volume 1.

We now modify the optimal map $\phi: x \rightarrow y$ to a map $\phi^d: x^d \rightarrow y^d$. For every edge e outside of x_ϕ and vertices v_0, v_1, v_2 and v_3 as above, we label the associated hair, oriented away from $\phi(v_0) \in y$, by ϵ_e , and the same edge with opposite orientation by ϵ_e^{-1} . We map $e_{0,1}, e_{1,2}, e_{2,3}$ to $\epsilon_e, \epsilon_e^{-1}, \phi(e)$, respectively. Note that the vertex v_1 is a one-gate vertex. By construction

$$L(\phi^d) = \lambda.$$

and the tension subgraph $x^d_{\phi^d} = x^d$. Let \mathcal{T} be the set of sub-gates of ϕ^d , let $\mathcal{T}_H \subset \mathcal{T}$ be set of sub-gates that fold to obtain all the hairs and let $\mathcal{T}_C = \mathcal{T} - \mathcal{T}_H$. Lastly let \bar{y}^d denote the graph obtain by scaling y^d by factor λ so that the core of \bar{y}^d has volume $1/\lambda$. Then the map

$$\bar{\phi}^d: x^d \longrightarrow \bar{y}^d$$

which is a composition of ϕ^d and the scaling map, is a 1-Lipschitz map.

Remark 4.1. We have placed the hair at the beginning of the edge e near v_0 . However, this is not essential and the hair could be placed at the end or anywhere in the middle. The placement of the hair may result on a different folding path but they all satisfy the properties we shall claim below.

4.2. Folding paths under the decorated difference of markings maps. We can use the decorated graphs to construct a folding path from x to y .

As done previously, for small enough t , we can fold x^d according to a given speed assignments $\mathcal{S} = \{s_\tau\}_{|\tau|=2}$ by identifying edges associated to τ along a sub-segment of length $t s_\tau$ for t small enough, to obtain a quotient map

$$\bar{\phi}_t^d : x^d \longrightarrow \bar{x}_t^d.$$

A hair in \bar{x}_t^d is any segment associated to folding for gates in \mathcal{T}_H .

We also define a map Cor that removes hair of a graph. Specifically, let \bar{x}_t be the graph obtained from \bar{x}_t^d after removing the hair and let

$$\text{Cor}: \bar{x}_t^d \longrightarrow \bar{x}_t$$

be the map that sends each hair in \bar{x}_t^d to the attaching vertex. Note that the map Cor can be used to mark the graph \bar{x}_t and we can consider \bar{x}_t as a point in cv_n . Let x_t be the associated point in CV_n which has volume one. Composition of $\text{Cor} \circ \bar{\phi}_t^d$ with the normalization map defines the map

$$\phi_t : x \longrightarrow x_t.$$

Similarly, define $\bar{y} = \text{Cor}(\bar{y}^d)$ which we consider as a point in cv_n , and define

$$\bar{\phi} : x \longrightarrow \bar{y} \quad \text{by} \quad \bar{\phi} = \text{Cor} \circ \bar{\phi}^d.$$

The maps $\bar{\psi}_t^d$ and ψ_t are the leftover difference of markings maps defined as

$$\bar{\psi}_t^d = \bar{\phi}^d \circ (\bar{\phi}_t^d)^{-1} \quad \text{and} \quad \psi_t = \phi \circ \phi_t^{-1}.$$

We refer to the discussion in Section 2.4 and note that while $(\phi_t^d)^{-1}, \phi_t^{-1}$ are only homotopy equivalences, the maps $\bar{\psi}_t^d$ and ψ_t are well defined.

Let $L(\phi_t)$ denote the Lipschitz constant of ϕ_t . By the definition of the distance in Outer space,

$$\log L(\phi_t) \geq d(x, x_t).$$

But also, $L(\phi_t)$ is the scaling factor from \bar{x}_t to x_t and the length of any loop that is legal in x^d increases by this factor from x to x_t . Thus,

$$\log L(\phi_t) = d(x, x_t).$$

Also, since the Lipschitz constant of ψ_t is constant everywhere along x_t , we have

$$L(\psi_t) = \frac{L(\phi)}{L(\phi_t)} \implies d(x_t, y) \leq \log L(\psi_t) = d(x, y) - d(x, x_t).$$

But $d(x, x_t) + d(x_t, y) \geq d(x, y)$ by the triangle inequality. Therefore,

$$d(x, x_t) + d(x_t, y) = d(x, y)$$

and hence the path $\{x_t\}$ is a geodesic starting from x towards y . To summarize, similar to Proposition 2.4 we have:

Proposition 4.2. *Given any two points x, y in CV_n there exists a decorated difference of markings map $\phi^d: x^d \rightarrow y^d$ such that $x_{\phi^d}^d = x^d$. Furthermore, any speed assignment \mathcal{S} defines a geodesic $\gamma: [0, t_1] \rightarrow CV_n$ starting from x towards y , for some $t_1 > 0$.*

We now prove an analogue of Theorem 3.1 for decorated folding paths following closely the constructions and arguments of Section 3. We always assume $t \in [0, t_1]$ (Proposition 4.2), in particular, x_t is in the same simplex at x .

Similar to Section 3, for any coherent speed assignment \mathcal{S} , we define

$$|\mathcal{S}| = \frac{1 - |\bar{x}_t|}{t}. \tag{16}$$

Then Lemma 3.2 still holds. Consider the lift of the map $\bar{\phi}^d$,

$$\Phi: T_x \longrightarrow T_{\bar{y}^d}$$

where $T_{\bar{y}^d}$ is the universal cover of \bar{y}^d . As before, for any $p \in T_{\bar{y}^d}$ let $\text{CH}(p)$ be the convex hull of $\text{Pre}(p)$ and Θ denote the set of sub-gate of $\text{CH}(p)$. We define the branching contributions $c(\sigma, p)$ as follows. If σ is associated to a hair and p is a point on this hair, then we set $c(\sigma, p) = 1$ and if p is any other point then we set $c(\sigma, p) = 0$. Note that if p is on a hair associated to σ , then $\text{CH}(p)$ contains only one illegal turn and $\Theta = \{\sigma\}$. For any σ not associated to a hair and a point p that is not on a hair, we ignore all illegal turns in $\text{CH}(p)$ associated to hairs and apply the construction of Proposition 3.4 to obtain a value for $c(\sigma, p)$ for every σ that is not associated to a hair.

To summarize, the length loss associated to hairs are assigned to illegal turns in \mathcal{T}_H and the rest of the length loss is assigned to illegal turns in \mathcal{T}_C according to Proposition 3.4. Equations (8) and (9) still hold.

We now define the length loss function ℓ_σ by

$$\ell_\sigma = \int_{T_{\bar{y}^d}} c(\sigma, p) dp$$

and $\ell_\tau = \ell_\sigma$ where σ is any lift of τ . We have, for $\tau \in \mathcal{T}_H$, ℓ_τ is the length of the hair associated to τ since $c(\sigma, p) = 1$ exactly when p is on the hair and zero otherwise. The proof of Lemma 3.7, show that

$$\sum_{\tau \in \mathcal{T}} \ell_\tau = |x| - |\bar{y}^d|.$$

But, \bar{y} is obtained from \bar{y}^d by removing the hair, hence,

$$|\bar{y}^d| - |\bar{y}| = \sum_{\tau \in \mathcal{T}_H} \ell_\tau.$$

Therefore,

Lemma 4.3. *For ℓ_τ defined as above, we have*

$$\sum_{\tau \in \mathcal{T}_C} \ell_\tau + 2 \sum_{\tau \in \mathcal{T}_H} \ell_\tau = |x| - |\bar{y}|.$$

Exactly as in equation (12) we define

$$s_\tau = \sum_{\hat{\tau} \geq \tau} \frac{\ell_{\hat{\tau}}}{|\hat{\tau}| - 1} \tag{17}$$

The proof of Lemma 3.9 still works to show

$$|\mathcal{S}| \leq \sum_{\tau \in \mathcal{T}_C} \ell_\tau + 2 \sum_{\tau \in \mathcal{T}_H} \ell_\tau, \tag{18}$$

since the proof works one gate at the time and for an illegal turn τ associated to a hair, the length loss associated to τ is $2\ell_\tau$.

Now we are ready to prove the main theorem of this section:

Theorem 4.4. *Given a difference of markings map $\phi^d: x^d \rightarrow y^d$, $x, y \in \text{CV}_n$, where $x_\phi^d = x^d$, there exists a speed assignment \mathcal{S} defining a folding path $\gamma: [0, t_1] \rightarrow \text{CV}_n$ starting at x towards y so that, for every loop α and every time $t \in [0, t_1]$,*

$$|\alpha|_t \leq \max(|\alpha|_x, |\alpha|_y).$$

Proof. We follow the proof of Theorem 3.1. Let \mathcal{S} be the speed assignment describe above and γ be the associated geodesic starting from x towards y coming from Proposition 4.2. Recall that a decorated difference of markings map fold x into a graph with hairs \bar{x}_t^d and where the core graph of \bar{x}_t^d is \bar{x}_t . By Lemma 3.2, for a given loop α ,

$$|\dot{\alpha}|_{t=0} = |\alpha|_x - 2 \sum_{\tau \in \mathcal{T}_\phi(\alpha)} \frac{s_\tau}{|\mathcal{S}|}.$$

By Lemma 4.3,

$$\sum_{\tau \in \mathcal{T}_C} \ell_\tau + 2 \sum_{\tau' \in \mathcal{T}_H} \ell_{\tau'} = |x| - |\bar{y}|.$$

Therefore, for every loop α such that $|\alpha|_y \leq |\alpha|_x$,

$$\begin{aligned} |\alpha|_y &= \frac{|\alpha|_x - \sum_i |u_i|_x}{|\bar{y}|} \leq |\alpha|_x \\ |\alpha|_x - \sum_i |u_i|_x &\leq |\alpha|_x |\bar{y}| = |\alpha|_x \left(1 - \sum_{\tau \in \mathcal{T}_C} \ell_\tau - 2 \sum_{\tau' \in \mathcal{T}_H} \ell_{\tau'}\right) \\ |\alpha|_x \left(\sum_{\tau \in \mathcal{T}_C} \ell_\tau + 2 \sum_{\tau' \in \mathcal{T}_H} \ell_{\tau'}\right) &\leq \sum_i |u_i|_x \leq 2 \sum_{\tau \in \mathcal{T}_\alpha} s_\tau. \end{aligned}$$

Recall, from equation (18), that

$$|\mathcal{S}| \leq \sum_{\tau \in \mathcal{T}_C} \ell_\tau + 2 \sum_{\tau' \in \mathcal{T}_H} \ell_{\tau'}.$$

Therefore

$$|\alpha|_x \leq 2 \sum_{\tau \in \mathcal{T}_\alpha} \frac{s_\tau}{|\mathcal{S}|}.$$

It follows that $|\alpha|_x < 2 \sum_{\tau \in \mathcal{T}_\alpha} \frac{s_\tau}{|\mathcal{S}|}$ and $|\dot{\alpha}|_{t=0} \leq 0$. Therefore the length $|\alpha|$ decreases and is smaller than $|\alpha|_x$, satisfying the claim of the theorem. On the other hand, if $|\alpha|_y > |\alpha|_x$, whatever the derivative may be, the claim of the theorem is satisfied until we have a time t where $|\alpha|_t \geq |\alpha|_y$, which falls into the case we address in the proof. \square

5. Construction of the balanced folding paths

In this section we prove Theorem 1.1 restated below:

Theorem 5.1. *Given points $x, y \in \text{CV}_n$, there exists a geodesic $[x, y]_{\text{bf}}$ from x to y so that, for every loop α , and every time t ,*

$$|\alpha|_t \leq \max(|\alpha|_x, |\alpha|_y).$$

Proof. Given x and y , we consider the decorated graphs x^d and y^d and the decorated difference of markings map $\phi^d: x^d \rightarrow y^d$. Applying Theorem 4.4, we obtain a geodesic $\gamma: [0, t_1] \rightarrow \text{CV}_n$ starting from x towards y . Now we consider the pair of points x_{t_1} and y and apply Theorem 4.4 again to continue the geodesic to an interval $[t_1, t_2]$. Continuing in this way, we either reach y after finitely many

steps or limit to a point $x' \in CV_n$. Note that every point x_t along this path has the property that

$$d(x, x_t) + d(x_t, y) = d(x, y) \quad (19)$$

and the set of such points is a compact subset of CV_n . Hence, the same holds for x' . In particular x' is a point in CV_n and the geodesic does not exit CV_n .

Now, we can apply Theorem 4.4 to the pair x' and y and continue the geodesic even further, getting closer to y . This results in a geodesic connecting x to y because if the process stops at some points x'' before y , then x'' is still in the compact set defined by equation (19) and we could apply Theorem 4.4 again to go further.

We re-parametrize this geodesic by arc-length to obtain $\gamma: [0, d] \rightarrow CV_n$, $d = d(x, y)$ (and use the parameter s to emphasize this fact). Let $\Sigma \subset [0, d]$ be the closure of set of times, each of which is an endpoint of an interval coming from an application of Theorem 4.4. Note that if $s \in \Sigma$, then an interval to the right of s is not in Σ . Hence, Σ is a well-ordered set. That is, $[0, d]$ is a union of the interiors of countably many intervals coming from Theorem 4.4 and the countable, well ordered set Σ which includes 0 and d .

For any loop α , we prove the theorem using transfinite induction on Σ . That is, for every time $s \in \Sigma$, we show

$$|\alpha|_s \leq \max(|\alpha|_x, |\alpha|_y). \quad (20)$$

The theorem for other times then follows from Theorem 4.4.

Equation (20) is clearly true for $s = 0$. For any $s \in \Sigma$, assume equation (20) holds for every $s' \in \Sigma$ with $s' < s$. We need to show that it also holds for s . There are two cases. If s is an endpoint of an interval $[s', s]$ coming from Theorem 4.4, then, by Theorem 4.4

$$|\alpha|_s \leq \max(|\alpha|_{s'}, |\alpha|_y)$$

and by the assumption of induction

$$|\alpha|_{s'} \leq \max(|\alpha|_x, |\alpha|_y)$$

and the conclusion follows.

Otherwise, there is a sequence $s_i \in \Sigma$, with $s_i < s$, so that $s_i \rightarrow s$. By the assumption of induction, we have

$$|\alpha|_{s_i} \leq \max(|\alpha|_x, |\alpha|_y).$$

But the length of α is a continuous function over CV_n . Taking a limit, we obtain the theorem. \square

6. Non-convexity

In this section we present examples that combine to prove Theorem 1.4. Some of the examples are done in low-rank free groups, however, they can easily be generalized to a higher rank. First, we show that there are points in Outer space such that no geodesic between them gives rise to convex length functions for all curves.

Proposition 6.1. *There are points $x, y \in CV_n$ and a loop α so that along any geodesic connecting x to y , the length of α is not a convex function of distance in CV_n .*

Proof. We construct a simple example in CV_2 . Let a and b be generators for \mathbb{F}_2 . Let $x \in CV_2$ be a graph that consists of two simple loops labeled a and b , wedged at a vertex v where each loop has length $\frac{1}{2}$ (a rose with two pedals). Let \bar{y} be a quotient of x obtained by identifying two subsegments of length $\frac{1}{8}$ in the loop labeled a . Then \bar{y} is a rank 2 graph in the shape of a dumbbell with total length $\frac{7}{8}$, where the b -loop has length $\frac{1}{2}$ and the a -loop has a length $\frac{1}{4}$. Let y be \bar{y} rescaled to have length 1 (by a factor $\frac{8}{7}$). There is a rigid folding path from x to y because \bar{y} was obtained from x by identifying two sub-edges. Hence, this path $[x, y]_f$ is the unique geodesic connecting x to y (see Remark 2.6). Let α be the loop representing element $a \in \mathbb{F}_2$. Then

$$|\alpha|_x = \frac{1}{2} \quad \text{and} \quad |\alpha|_y = \frac{8}{7} \cdot \frac{1}{4} = \frac{2}{7}.$$

Consider the length of $|\alpha|_t$ of α along this folding path. By Lemma 3.2, the derivative of the length of α at x is:

$$|\dot{\alpha}|_t|_{t=0} = |\alpha|_x - 2 = \frac{1}{2} - 2 < 0.$$

And since the length of α is decreasing the derivative stays negative. In fact, for $s > 0$

$$|\dot{\alpha}|_t|_{t=s} = |\alpha|_s - 2$$

is a decreasing function as well. Therefore $|\alpha|_t$ is concave along this folding path. That is, there is no geodesic between x and y on which the length of α is a convex function of distance. □

We now examine if Theorem 1.1 holds for other geodesics connecting two points in CV_n . We start by looking at a general folding path and we show that a folding path with endpoints in a small ball can still go arbitrarily far away from the center of the ball.

Proposition 6.2. *For any $R > 0$, there are points $x, y, z \in CV_3$ and there is a folding path $[y, z]_{ng}$ connecting y to z so that*

$$y, z \in B_{out}(x, 2) \quad \text{and} \quad [y, z]_{ng} \not\subset B_{out}(x, R).$$

Proof. The example presented below is in CV_3 . For higher rank Outer spaces, one can modify the example to roses with more loops such that the optimal map outside of the simple loops labeled a, b, c is identity.

Consider constants $\epsilon > 0, \delta > 0$ and an integer $m > 0$ so that

$$\epsilon \leq \delta \ll 1 \quad \text{and} \quad m\delta = 1 - 3\delta.$$

Assume \mathbb{F}_3 is generated by elements a, b and c and let x, y, z and w be points in CV_3 which are wedges of simple loops with lengths and labels summed up in Table 1.

Table 1

	x		y		w		z	
	label	length	label	length	label	length	label	length
Edge 1	a	$\frac{\epsilon}{2}$	ab	$\delta + \delta^2$	ab	$\frac{1+\delta}{3}$	a	$\frac{\delta}{2}$
Edge 2	b	$\frac{1}{2}$	b	δ	b	$\frac{1}{3}$	b	$\frac{1}{2}$
Edge 3	c	$\frac{1-\epsilon}{2}$	cb^m	$1 - 2\delta - \delta^2$	c	$\frac{1-\delta}{3}$	c	$\frac{1-\delta}{2}$

Note that if, in y , we fold the edge labeled cb^m m -times around b (without rescaling), we obtain a graph \bar{w} with labels ab, b and c and lengths $(\delta + \delta^2), \delta$ and

$$(1 - 2\delta - \delta^2) - m\delta = (1 - 2\delta - \delta^2) - (1 - 3\delta) = \delta - \delta^2$$

which is a graph that is projectively equivalent to w (by a factor of $\frac{1}{3\delta}$). Similarly, if in w , we fold the edge labeled ab once around b (without rescaling), we obtain a graph \bar{z} with labels a, b and c and lengths $\frac{\delta}{3}, \frac{1}{3}$ and $\frac{1-\delta}{3}$ which a graph that is projectively equivalent to z by a factor $\frac{3}{2}$. Therefore, there is a folding path from y to z that passes through w . But this is not a greedy folding path since the edge labeled ab is not folded around b in the segment $[y, w]$.

Let α be the loop representing the element $a \in \mathbb{F}_3$. Then

$$d(x, y) = \log \frac{|\alpha|_y}{|\alpha|_x} = \log \frac{2(2\delta + \delta^2)}{\epsilon},$$

$$d(x, z) = \log \frac{|\alpha|_z}{|\alpha|_x} = \log \frac{\delta}{\epsilon},$$

$$d(x, w) = \log \frac{|\alpha|_w}{|\alpha|_x} = \log \frac{\frac{2}{3}(2 + \delta)}{\epsilon} \geq \log \frac{4}{3\epsilon}.$$

If, for example, we let $\epsilon = \delta$ then $y, z \in B_{\text{out}}(x, 2)$, but w can be made arbitrarily far away by making δ small. □

Next, we consider standard geodesic paths connecting two points which are the type of geodesics most often considered to connect two arbitrary points in CV_n (not every pair of points can be connected via a folding path). The situation is improved somewhat but, by taking the ball large enough, one can construct examples where a standard geodesic with its endpoints in a ball goes arbitrarily far from the ball.

Proposition 6.3. *There exists a constant $c > 0$ such that, for every $R > 0$, there are points $x, y, z \in CV_n$ and a standard geodesic $[y, z]_{\text{std}}$ connecting y to z such that*

$$y, z \in B_{\text{out}}(x, R) \quad \text{and} \quad [y, z]_{\text{std}} \not\subset B_{\text{out}}(x, 2R - c).$$

That is, the standard geodesic path can travel nearly twice as far from x as y and z are from x .

Proof. As before, we construct the example in CV_3 . Let $\psi \in \text{Out}(\mathbb{F}_3)$ be defined as follows:

$$\begin{aligned} \psi(a) &= ab, & \psi^{-1}(a) &= b, \\ \psi(b) &= a, & \psi^{-1}(b) &= b^{-1}a, \\ \psi(c) &= c, & \psi^{-1}(c) &= c. \end{aligned}$$

It is known (and easy to see) that, for any integer $m > 0$, the word length of $\psi^m(a)$ is F_{m+3} and the word length of $\psi^m(b)$ is F_{m+2} , where F_i is the i -th Fibonacci number. Similarly, the word length of $\psi^{-m}(a)$ is F_{m+2} and the word length of $\psi^{-m}(b)$ is F_{m+3} . For a large integer $m > 0$, let

$$\delta = \frac{1}{F_{m+2} + F_{m+3} + 1}$$

and consider points $x, y, z, w \in CV_3$ which are wedges of simple loops and where the lengths and edge labels are summed up in Table 2.

If we let \bar{z} be the rose with labels a, b and c and all edge lengths δ , then there is a quotient map $\bar{\phi}: w \rightarrow \bar{z}$ that maps the edge of w labeled $\psi^m(a)$ to an edge path containing F_{m+3} edges and maps the edge of w labeled $\psi^m(b)$ to an edge path containing F_{m+2} edges. The graph z is obtained from \bar{z} by scaling by a factor $\frac{1}{3\delta}$. Hence, w can be connected to z using a folding path and tension graph of $\phi: w \rightarrow z$ is all of w . The map from y to w scales two of the edges and contracts the third. The loop $\psi^m(a)$ is maximally stretched from y to w because

Table 2

	<i>x</i>		<i>y</i>		<i>w</i>		<i>z</i>	
	label	length	label	length	label	length	label	length
Edge 1	<i>a</i>	δ	$\psi^m(a)$	δ	$\psi^m(a)$	$F_{m+3} \delta$	<i>a</i>	$\frac{1}{3}$
Edge 2	<i>b</i>	δ	$\psi^m(b)$	δ	$\psi^m(b)$	$F_{m+2} \delta$	<i>b</i>	$\frac{1}{3}$
Edge 3	<i>c</i>	$1 - 2\delta$	<i>c</i>	$1 - 2\delta$	<i>c</i>	δ	<i>c</i>	$\frac{1}{3}$

$F_{m+3} \geq F_{m+2}$. Both loops $\psi^m(a)$ and $\psi^m(b)$ are maximally stretched from *w* to *z*. Since the same loop $\psi^m(a)$ is stretched maximally from *y* to *w* and from *w* to *z*, we have $d(y, z) \geq d(y, w) + d(w, z)$. Hence, the standard geodesic from *y* to *z* constructed above passes through *w*.

Next, we compute the distance from *x* to these points. Let α be the loop representing the element $a \in \mathbb{F}_3$ and β be the loop representing $b \in \mathbb{F}_3$. The loop α has a combinatorial length F_{m+2} (which is the word length of $\psi^{-m}(a)$) in both *y* and *w* and β has a combinatorial length F_{m+3} (which is the word length of $\phi^{-m}(b)$) in both *y* and *w*. In particular,

$$|\beta|_w \geq F_{m+3} \cdot (F_{m+2} \delta)$$

because the geodesic representative of β in *w* consists of F_{m+3} edges each having a length of at least $F_{m+2} \delta$. We have

$$d(x, y) = \log \frac{|\beta|_y}{|\beta|_x} = \log \frac{F_{m+3} \delta}{\delta} = \log F_{m+3},$$

$$d(x, z) = \log \frac{|\alpha|_z}{|\alpha|_x} = \log \frac{|\beta|_y}{|\beta|_x} = \log \frac{1/3}{\delta} = \log \frac{1}{3\delta},$$

$$d(x, w) \geq \log \frac{|\beta|_y}{|\beta|_x} > \log \frac{F_{m+3} \cdot (F_{m+2} \delta)}{\delta} = \log(F_{m+3} F_{m+2}).$$

We now set $R = \log F_{m+3}$ which is larger than $\log \frac{1}{3\delta}$. Then, $y, z \in B_{\text{out}}(x, R)$. There is a constant *c*, (slightly larger than the logarithm of the golden ratio) so that

$$\log(F_{m+2} F_{m+3}) \geq 2 \log(F_{m+3}) - c = 2R - c$$

which implies $w \notin B_{\text{out}}(x, 2R - c)$. This finishes the proof. □

The most well-behaved geodesic often considered is a greedy folding path. In fact, as mentioned in the introduction, the lengths of curves are quasi-convex function of distance along a greedy folding path. However, we show that a greedy folding path with endpoint inside of a ball may exit the ball.

Proposition 6.4. *For $n \geq 4$ and every $R > 0$, there are points $x, y, z \in \text{CV}_{n+2}$ where y and z are connected by a greedy folding path $[y, z]_{\text{gf}}$ such that*

$$y, z \in B_{\text{out}}(x, R) \quad \text{but} \quad [y, z]_{\text{gf}} \not\subset B_{\text{out}}(x, R).$$

Proof. Let $x, y, z, w \in \text{CV}_{n+2}$ be four graphs that are each a bouquets of $n + 2$ simple loops. Consider \mathbb{F}_{n+2} as being generated by a, b and c_i , for $i = 1, \dots, n$. The lengths and the labels of these graphs are described in Table 3 below where ϵ is a small positive number.

Table 3

	x		y		z		w	
	label	length	label	length	label	length	label	length
Edge 1	a	ϵ	ab^2	$3/(2n + 4)$	a	$1/(n + 2)$	ab	$2/(n + 3)$
Edge 2	b	$\frac{(1-\epsilon)}{2}$	b	$1/(2n + 4)$	b	$1/(n + 2)$	b	$1/(n + 3)$
Edge i	c_i	$\frac{(1-\epsilon)}{2n}$	$c_i b$	$2/(2n + 4)$	c_i	$1/(n + 2)$	c_i	$1/(n + 3)$

Note that there is an obvious optimal map from y to z that linearly expands each edge of y around the edges in z according to the labeling. The greedy folding path from y to z passes through w . In fact, $[y, z]_{\text{gf}}$ consists of two subsegments, in the first part $c_i b$ and ab^2 wrap around b simultaneously to reach w , and in the second part the edge labeled ab wraps around b to reach z . The distance $d(y, w) = \log \frac{2n+4}{n+3}$ and the associated stretch factors of edges

$$\lambda(ab^2) = \frac{3/(n + 3)}{3/(2n + 4)} = \frac{2n + 4}{n + 3},$$

$$\lambda(b) = \frac{1/(n + 3)}{1/(2n + 4)} = \frac{2n + 4}{n + 3},$$

$$\lambda(c_i b) = \frac{2/(n + 3)}{2/(2n + 4)} = \frac{2n + 4}{n + 3},$$

are all the same. Likewise, the distance $d(w, z) = \log \frac{n+3}{n+2}$ and associated stretch factors of edges are

$$\lambda(ab) = \frac{2/(n + 2)}{2/(n + 3)} = \frac{n + 3}{n + 2},$$

$$\lambda(b) = \frac{1/(n + 2)}{1/(n + 3)} = \frac{n + 3}{n + 2},$$

$$\lambda(c_i) = \frac{1/(n+2)}{1/(n+3)} = \frac{n+3}{n+2},$$

which again are the same for every edge.

Next, we measure distances from the center of the ball x . Let α be the loop associated with the element $a \in \mathbb{F}_3$. For ϵ small enough, all three distances are realized by the stretch factor associated to α . That is,

$$d(x, y) = \log \frac{|\alpha|_y}{|\alpha|_x} = \log \frac{5/2n+4}{\epsilon} = \log \frac{5}{2n\epsilon+4\epsilon},$$

$$d(x, z) = \log \frac{|\alpha|_z}{|\alpha|_x} = \log \frac{1/n+2}{\epsilon} = \log \frac{1}{n\epsilon+2\epsilon},$$

$$d(x, w) = \log \frac{|\alpha|_w}{|\alpha|_x} = \log \frac{3/n+3}{\epsilon} = \log \frac{3}{n\epsilon+3\epsilon}.$$

But, for all $n \geq 4$, we have

$$\frac{3}{n\epsilon+3\epsilon} > \max\left(\frac{1}{n\epsilon+2\epsilon}, \frac{5}{2n\epsilon+4\epsilon}\right).$$

Thus, if we set $R = \frac{1}{n\epsilon+2\epsilon}$, we have an example of a greedy folding path with endpoint in $B_{\text{out}}(x, R)$ that travels outside the ball. \square

7. In-coming balls

In contrast with out-going balls, we prove that in-coming balls are not weakly quasi-convex:

Theorem 7.1. *For any constant $R > 0$, there are points $x, y, z \in \text{CV}_n$ such that, $y, z \in B_{\text{in}}(x, 2)$ but, for any geodesic $[y, z]$ connecting y to z ,*

$$[y, z] \not\subset B_{\text{in}}(x, R).$$

Proof. We show that there exists a family of balls and pairs of points y_m and z_m in these balls such that the geodesic connecting y_m to z_m is unique and it travels arbitrarily far away from the center of the corresponding ball. Since the geodesic is unique, this can be restated as: every geodesic connecting y_m to z_m travels arbitrarily far from the center of the balls.

Fix an integer $m > 0$ and, as usual, let a, b and c be generators for \mathbb{F}_3 . Examples in higher dimension can be adapted from this example by adding loops on which the map is identity along the path. Let $x = x_m, y = y_m$ and $z = z_m$ be roses with labels and lengths specified in Table 4.

Table 4

	x		y		w		z	
	label	length	label	length	label	length	label	length
Edge 1	a	$\frac{1}{2} - \frac{1}{m}$	ab^m	$\frac{m+1}{2m+4}$	a	$\frac{1}{m+4}$	a	$\frac{1}{3}$
Edge 2	b	$\frac{1}{m}$	b	$\frac{1}{2m+4}$	b	$\frac{1}{m+4}$	b	$\frac{1}{3}$
Edge 3	c	$\frac{1}{2}$	$cb^m a$	$\frac{m+2}{2m+4}$	$cb^m a$	$\frac{m+2}{m+4}$	c	$\frac{1}{3}$

Note that w is obtained from y by wrapping the edge labeled ab^m around the edge labeled b m -times and then scaling by a factor $\frac{2m+4}{m+4}$. Throughout this portion, the illegal turn $\langle ab^m, b \rangle$ is the only illegal turn. Similarly, z is obtained from w wrapping the edge labeled ab^m around the edge labeled a once, then around the edge labeled b m -times and finally scaling by $\frac{m+4}{3}$. Again, during each sub-segment, there is exactly one non-yo-yo illegal turn; first $\langle cb^m a, a \rangle$ and next $\langle cb^m, b \rangle$. The illegal turn is never a yo-yo since there is no cut edge in the graphs along the paths. The loop labeled b in y is legal throughout and hence is maximally stretched from y to w and from w to z . Therefore w lies on a rigid folding path from y to z . By Theorem 1.5 the folding path is the unique (up to re-parametrization) geodesic from y to z .

We now compute distance to the center of the ball. For large enough m , we have

$$d(y, x) = \log \frac{|cb^m a|_x}{|cb^m a|_y} = \log \frac{\frac{1}{2} + 1 + \frac{1}{2} - \frac{1}{m}}{\frac{m+2}{2m+4}} = \log \frac{4m^2 + 6m - 4}{m^2 + 2} < \log 5 < 2,$$

$$d(w, x) = \log \frac{|a|_x}{|a|_w} = \log \frac{\frac{1}{2} - \frac{1}{m}}{\frac{1}{m+4}} = \log \frac{m^2 + 2m - 8}{2m} \geq \log \frac{m}{2},$$

$$d(z, x) = \log \frac{|c|_x}{|c|_z} = \log \frac{\frac{1}{2}}{\frac{1}{3}} = \log \frac{3}{2} < 2.$$

That is, $y, z \in B_{\text{in}}(x, 2)$ and the distance $d(w, x)$ can be made to be arbitrarily large. □

References

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

- [2] Y. Algom-Kfir and M. Bestvina, Asymmetry of outer space. *Geom. Dedicata* **156** (2012), 81–92. [Zbl 1271.20029](#) [MR 2863547](#)
- [3] M. Bestvina and M. Feighn, Subfactor projections. *J. Topol.* **7** (2014), no. 3, 771–804. [Zbl 1346.20054](#) [MR 3252963](#)
- [4] M. Bestvina and M. Feighn, Hyperbolicity of the complex of free factors. *Adv. Math.* **256** (2014), 104–155. [Zbl 1348.20028](#) [MR 3177291](#)
- [5] M. Culler, Finite groups of outer automorphisms of a free group. In K. I. Appel, J. G. Ratcliffe, and P. E. Schupp (eds.), *Contributions to group theory*. Papers dedicated to R. C. Lyndon on the occasion of his 65th birthday. Contemporary Mathematics, 33. American Mathematical Society, Providence, R.I., 1984. [Zbl 0552.20024](#) [MR 0767107](#)
- [6] M. Culler and K. Vogtmann, Moduli of graphs and automorphisms of free groups. *Invent. Math.* **84** (1986), no. 1, 91–119. [Zbl 0589.20022](#) [MR 0830040](#)
- [7] S. Dowdall and S. J. Taylor, Hyperbolic extensions of free groups. *Geom. Topol.* **22** (2018), no. 1, 517–570. [Zbl 1439.20034](#) [MR 3720349](#)
- [8] M Fortier Bourque and K. Rafi, Non-convex balls in the Teichmüller metric. *J. Differential Geom.* **110** (2018), no. 3, 379–412. [Zbl 1404.30049](#) [MR 3880229](#)
- [9] S. Francaviglia and A. Martino, Metric properties of outer space. *Publ. Mat.* **55** (2011), no. 2, 433–473. [Zbl 1268.20042](#) [MR 2839451](#)
- [10] M. Gromov, *Hyperbolic groups*. in S. M. Gersten (ed.), *Essays in group theory*. Mathematical Sciences Research Institute Publications, 8. Springer-Verlag, New York, 1987, 75–263. [Zbl 0634.20015](#) [MR 0919829](#)
- [11] S. P. Kerckhoff, The Nielsen realization problem. *Ann. of Math. (2)* **117** (1983), no. 2, 235–265. [Zbl 0528.57008](#) [MR 0690845](#)
- [12] D. G. Khramtsov, Finite groups of automorphisms of free groups. *Mat. Zametki* **38** (1985), no. 3, 386–392, 476. In Russian. English translation, *Math. Notes* **38** (1985), no. 3–4, 721–724. [Zbl 0595.20036](#) [MR 0811572](#)
- [13] S. Kravetz, On the geometry of Teichmüller spaces and the structure of their modular groups. *Ann. Acad. Sci. Fenn. Ser. A I* **278** (1959), 35 pp. [Zbl 0168.04601](#) [MR 0148906](#)
- [14] A. Lenzhen and K. Rafi, Length of a curve is quasi-convex along a Teichmüller geodesic. *J. Differential Geom.* **88** (2011), no. 2, 267–295. [Zbl 1243.30089](#) [MR 2838267](#)
- [15] W. P. Thurston, Minimal stretch maps between hyperbolic surfaces. Preprint, 1998. [arXiv:math/9801039](#)
- [16] K. Vogtmann, On the geometry of outer space. *Bull. Amer. Math. Soc. (N.S.)* **52** (2015), no. 1, 27–46. [Zbl 1368.20052](#) [MR 3286480](#)
- [17] T. White, Fixed points of finite groups of free group automorphisms. *Proc. Amer. Math. Soc.* **118** (1993), no. 3, 681–688. [Zbl 0798.20021](#) [MR 1164152](#)

Received December 11, 2017

Yulan Qing, Shanghai Center for Mathematical Sciences, Fudan University,
2005 Songhu Road, Shanghai, 200438, China

e-mail: yulan.qing@gmail.com

Kasra Rafi, Department of Mathematics, University of Toronto, Toronto, Ontario,
Canada M5S 2E4

e-mail: rafi@math.toronto.edu