

METRIC PROPERTIES OF OUTER SPACE

STEFANO FRANCAVIGLIA AND ARMANDO MARTINO

Abstract

We define metrics on Culler-Vogtmann space, which are an analogue of the Thurston metric and are constructed using stretching factors. In fact the metrics we study are related, one being a symmetrised version of the other. We investigate the basic properties of these metrics, showing the advantages and pathologies of both choices.

We show how to compute stretching factors between marked metric graphs in an easy way and we discuss the behaviour of stretching factors under iterations of automorphisms.

We study metric properties of folding paths, showing that they are geodesic for the non-symmetric metric and, if they do not enter the thin part of Outer Space, quasi-geodesic for the symmetric metric.

CONTENTS

1. Introduction	434
2. Preliminaries	435
3. Calculating stretching factors	437
4. Metrics	446
5. Folding paths and geodesics	455
6. The symmetric metric is not geodesic	460
7. Quasi-geodesics	462
8. Iterating automorphisms	468
9. Some open questions	471
9.1. Existence of quasi-geodesics	471
9.2. Existence of a geodesic axis for an iwip	471
9.3. Hyperbolicity, flats and coarse properties	472
References	472

2010 *Mathematics Subject Classification.* 20F65.

Key words. Outer space, free group, Lipschitz metric, stretching factor, optimal maps.

1. Introduction

Culler-Vogtmann space, or Outer Space as it is sometimes called, has been the subject of intense study. Much of the direction of this work has been to develop a theory for Outer Space, and the Outer Automorphism group of a free group in an analogous way to the theory of Teichmüller space and the mapping class group of a surface.

Our contribution to this effort is the study of a metric which is an analogue of the Thurston metric, with the goal that the important features of both Outer Space, and the automorphisms of a free group are captured by the geometry of this metric.

After recalling the basic definitions in Section 2, we spend some time in Section 3 defining and understanding the “one-sided” metric, from which our metric is obtained by a “symmetrisation”. In fact, a special case of this one-sided metric (where the objects are a rose, and its image under an automorphism) is a quantity that has appeared in the work of Kapovich, [11], where it is shown that the value is computable in double exponential time (in the number of Nielsen generators). As part of our efforts to understand our metric, and simplify many of the proofs of its properties, we show that the calculation is considerably simpler, Proposition 3.15, so that the calculation for a rose is actually linear in the length of the images of the generators (so exponential in the number of Nielsen generators) though it is exponential for more complicated graphs.

We then study the metric itself in Section 4, and show that the metric topology is the same as the usual length function topology, as well as showing that the metric is proper; closed balls are compact in this space. This is one advantage the symmetric version of the metric has over the unsymmetric version, since for the one-sided metric not only are Cauchy sequences not always convergent, but also points which *should* be at infinite distance, namely points on the boundary of Outer Space, are actually at finite distance from points in the interior of Outer Space.

Section 5 is concerned with the connection between the geometry of Outer Space and the properties of the automorphisms of a free group. Specifically, we study the behaviour of “folding paths” and their metric properties. It is fairly straightforward to show that these paths are geodesics for the one-sided metric, but it seems to be much more difficult to show that they are even quasi-geodesics for the actual metric. However, these folding paths are shown to have good properties, such as the “4 point property”, defined in Theorem 7.3.

In Section 6 we show with an example that Outer Space, equipped with the symmetric metric, is not a geodesic space. We want to stress here that such example was suggested to the authors by Bert Wiest and Thierry Coulbois when a previous version of this paper was posted on the arxiv.

In Section 7, we show that if folding paths remain within the “thick part” of Outer Space, then they will be quasi-geodesics which is a result, definitions aside, that is very intuitive. We finish, in Section 8 by showing that for an automorphism of exponential growth, the map from \mathbb{Z} to Outer Space which sends an integer, n to the n^{th} iterate of a given point under the automorphism is a quasi-isometry. Interestingly, while this result is clearly false for automorphisms of polynomial growth, we show that for a particular example of polynomial growth automorphism, the folding path between the rose and a image of the rose under an (arbitrary) iterate of the automorphism is a quasi-geodesic with uniform constants.

Acknowledgements. We warmly, sincerely, and enthusiastically thank M. Bestvina, K. Bromberg, T. Coulbois, V. Guirardel, A. Hilion, P. Hubert, I. Kapovich, G. Levitt, J. Los, M. Lustig, G. Th  ret and B. Wiest. This research was possible thanks to the financial support of the European Research Council (MEIF-CT-2005-010975 and MERG-CT-2007-046557). The second named author gratefully acknowledges partial support from the MEC (Spain) through project BFM2003-06613.

2. Preliminaries

We refer the reader to [13] for an excellent survey and reference article to Culler-Vogtmann space.

Our basic objects are finite marked metric graphs of some given rank n . A graph of this type is represented as a metric graph, A —that is, with a positive length assigned to every edge— and a marking τ_A which is a homotopy equivalence from the rose with n petals, R_n to A ,

$$\tau_A: R_n \rightarrow A.$$

We shall make the standard assumption that vertices have valence at least three. Nonetheless, we notice that it is sometimes convenient to allow vertices of valence two. When it is clear from the context, we will not specify whether we use bi-valent vertices.

Two marked graphs A and B are equivalent if there is a homotopy, $h: A \rightarrow B$, such that the following diagram commutes up to free

homotopy,

$$\begin{array}{ccc}
 A & \xrightarrow{h} & B \\
 & \swarrow \tau_A & \nearrow \tau_B \\
 & R_n &
 \end{array}$$

Alternatively, we could only consider metric graphs of volume 1 and then the equivalence would be given by isometries in place of homotheties. In either case, the resulting space of equivalence classes is called Culler-Vogtmann space of rank n , or CV_n (when bi-valent vertices are allowed, two marked graphs are also equivalent if they have a common finite subdivision).

Remark 2.1. In the following, if there are no ambiguities we will not distinguish between a marked metric graph and its class.

When we will need to be precise we will refer to a metric graph as an element of the unprojectivised CV_n , and to its class as an element of CV_n .

Given any marked graph A , we can look at the universal cover T_A which is an \mathbb{R} -tree on which $\pi_1(R_n)$ acts by isometries, via the marking τ_A . (From now on, we identify the free group of rank n , F_n , with the fundamental group of R_n .) Conversely, given any minimal free action of F_n by isometries on a simplicial \mathbb{R} -tree, we can look at the quotient object, which will be a graph, A , and produce a homotopy equivalence $\tau_A: R_n \rightarrow A$ via the action. Equivalence of graphs in CV_n corresponds to actions which are equivalent up to equivariant homothety.

Thus, points in CV_n , can be thought of as equivalence classes of minimal free isometric actions on simplicial \mathbb{R} -trees. Given an element w of F_n and a point A of the unprojectivised CV_n , with universal cover T_A whose metric we denote by d_A , we may consider,

$$l_A(w) := \inf_{p \in T_A} d_A(p, wp).$$

It is well known that this infimum is always achieved and that, for a free action, it is non-zero for the non-identity elements of the group. In this context, $l_A(w)$ is called the translation length of the element w in the corresponding tree and clearly depends only on the conjugacy class of w in F_n . Thus for any point, A , in CV_n we can associate the sequence $(l_A(w))_{w \in F_n}$ and it is clear that equivalent marked metric graphs will produce two sequences, one of which is a multiple of the other by a positive real number (the homothety constant). Moreover, it is also the case that inequivalent points in CV_n will produce sequences which

are not multiples of each other [8]. Thus, we have an embedding of CV_n into \mathbb{R}^{F_n}/\sim , where \sim is the equivalence relation of homothety. The space CV_n is given the subspace topology induced by this embedding.

Finally it is clear we can realise any automorphism, ϕ , of F_n as a homotopy equivalence, also called ϕ , of R_n . Thus the automorphism group of F_n acts on CV_n by changing the marking. That is, given a point (A, τ_A) of CV_n the image of this point under ϕ is $(A, \tau_A\phi)$.

$$\begin{array}{ccccc}
 & & \tau_A\phi & & \\
 & \curvearrowright & & \curvearrowleft & \\
 R_n & \xrightarrow{\phi} & R_n & \xrightarrow{\tau_A} & A.
 \end{array}$$

Since two automorphisms which differ by an inner automorphism always send equivalent points in CV_n to equivalent points, we actually have an action of $\text{Out}(F_n)$ on CV_n , and this space is often called *Outer Space* for this reason.

3. Calculating stretching factors

Given two marked metric graphs, A and B , both having fundamental group free of rank n , we would like to compute the distance between them and, as a first step, the “right hand distance” between them, defined as follows.

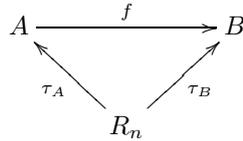
Definition 3.1 (Right hand factor). For any pair A, B of marked graphs we set

$$\Lambda_R(A, B) := \sup_{1 \neq w \in F_n} \frac{l_B(w)}{l_A(w)}.$$

Recall that $l_A(w)$ is the translation length of the element corresponding to w in the tree T_A (and hence is dependent only on the conjugacy class of w). However, it is readily seen that this translation length is the same as the length of the shortest representative in the free homotopy class of loops in A defined by the (conjugacy class of) w . We note that this second definition means that $l_A(w)$ is easy to compute given a particular w : we look at the image of w in A via the marking and we “cyclically reduce” the loop in the graph by performing free cyclic reductions which may, of course, change the basepoint. The length of any cyclically reduced element in this sense, calculated simply by summing the lengths of the edges crossed, will be $l_A(w)$. We shall also use l_A to refer to the lengths of (free homotopy classes of) loops in A in the obvious way. We also note that saying a loop in A is cyclically reduced is equivalent to saying that, if we consider the loop as a map from the circle

to the graph it is an immersion. In the same spirit, a path is reduced if it is an immersion when considered as a map from a closed interval.

While finding lengths of elements with respect to a marked metric graph structure is straightforward, that does not indicate how to calculate the supremum given above. In order to do that, we need to relate one structure to the other. One way to do this is to find an equivariant map from A to B , which we can simply think of as a homotopy equivalence between the graphs, which respects the markings. That is, a map f for which the following triangle commutes up to free homotopy,



In other words, f is a map homotopic to τ_A^{-1} followed by τ_B , $f \simeq \tau_B \tau_A^{-1}$. Here we are abusing notation by writing τ_A^{-1} for a *homotopy inverse* as opposed to the actual inverse.

It is important to note that this is not a graph map in that edges are not necessarily sent to edges nor vertices to vertices. We will therefore restrict to a particular class of maps that are more easy to handle.

Definition 3.2 (PL maps). A map $f: A \rightarrow B$ is a PL-map if it is linear on edges. More precisely, for each edge e of A , if we parameterise $f|_e$ with the segment $[0, l_A(e)]$, then $f|_e$ has constant speed. We denote by $S_{f,e}$ the speed of $f|_e$ (the stretching factor of e).

The stretching factor of a PL-map f , defined as the maximal speed of f , is in fact the Lipschitz constant of f . We denote that quantity by S_f (the notation L_f for the Lipschitz constant is more natural but also more confusing since we already have lengths denoted by the letter l)

$$S_f = \max\{S_{f,e} : e \text{ edge of } A\} = \text{Lip}(f).$$

In general, given f , there is a unique PL-map \bar{f} which is homotopic to f and agrees with f on vertices. It is readily checked that

$$(1) \quad S_{\bar{f}} = \text{Lip}(\bar{f}) \leq \text{Lip}(f).$$

A useful observation one can make here is that $\text{Lip}(f)$ serves as an upper bound for $\Lambda_R(A, B)$. This is because, starting with a loop γ in A , it is clear that

$$l_B(f(\gamma)) \leq \text{Lip}(f)l_A(\gamma).$$

Since we can consider all loops which are cyclically reduced in A this means that,

$$l_B(w) \leq \text{Lip}(f)l_A(w), \text{ for all } w \in F_n,$$

and we hence proved

Lemma 3.3. *For any Lipschitz map $f: A \rightarrow B$ in the homotopy class of $\tau_B\tau_A^{-1}$*

$$\Lambda_R(A, B) = \sup_{1 \neq w \in F_n} \frac{l_B(w)}{l_A(w)} \leq \text{Lip}(f).$$

Since f is arbitrary, and because of (1), we can deduce that

$$(2) \quad \Lambda_R(A, B) = \sup_{1 \neq w \in F_n} \frac{l_B(w)}{l_A(w)} \leq \inf\{S_f : f \text{ is PL and } f \simeq \tau_B\tau_A^{-1}\}.$$

It is fairly clear that the infimum on the right hand side of equation (2) will be realised by an actual map.

Lemma 3.4. *Let A, B two marked metric graphs. Then there exists an $f_\infty \simeq \tau_B\tau_A^{-1}$ such that*

$$S_{f_\infty} = \inf\{S_f : f \text{ is PL and } f \simeq \tau_B\tau_A^{-1}\} = \inf\{\text{Lip}(f) : f \simeq \tau_B\tau_A^{-1}\}.$$

Proof: For any c , the set of c -Lipschitz maps from A to B is precompact by Ascoli-Arzelà theorem because B is compact. Therefore a sequence of maps f_n , whose stretching factors tend to the infimum has a convergent sub-sequence whose limit is f_∞ , and it is easily checked that $S_{f_\infty} = \inf\{S_{f_n}\}$ and that, if each f_n is in the same homotopy class as $\tau_B\tau_A^{-1}$, then so is f_∞ . \square

Remark 3.5. Equations (1) and (2), and Lemma 3.4 tell us that from now on we can, as we do, assume that any map is a PL.

Now note that there are two obstructions to making equation (2) an equality. While we may realise the infimum by a concrete map, f , we may still have that for a given loop γ , not all edges of γ may be stretched by the same amount S_f . Thus we need the collection of edges which are stretched maximally to be large enough as to contain a loop. Furthermore, even if we have such a loop γ , the image $f(\gamma)$ may not be cyclically reduced in B . However, if we have a cyclically reduced loop, γ , in A , all of whose edges are stretched by S_f and such that $f(\gamma)$ is cyclically reduced in B , then $\Lambda_R(A, B) = S_f$. It will turn out that there always exists a map f and a loop γ with these properties. Before going into details, we need some preliminaries.

Definition 3.6. Let A, B be marked metric graphs of rank n . Given a PL-map $f \simeq \tau_B\tau_A^{-1}$, we denote by $A_{\max}(f)$ the subgraph of A whose edges are stretched maximally, by S_f .

Definition 3.7 (Optimal maps). A PL-map $f \simeq \tau_B \tau_A^{-1}$ is NOT optimal if there is some vertex of $A_{\max}(f)$ such that all edges of $A_{\max}(f)$ terminating at that vertex have f -image with a common terminal partial edge.

Otherwise f is optimal.

Remark 3.8. Using the terminology of legal and illegal turns, a PL-map is optimal if each vertex of A_{\max} has at least one legal turn.

Suppose that a map $f \simeq \tau_B \tau_A^{-1}$ is not optimal. Let v be a vertex of $A_{\max}(f)$ such that all edges of $A_{\max}(f)$ terminating at v have f -image with a common terminal partial edge, say α . Let $\text{star}(v)$ denote the set of edges emanating from v . We set $N = \text{star}(v) \cap A_{\max}$ and $K = \text{star}(v) \setminus N$.

Now, let f_t be the homotopy that moves v backward along α . More precisely, we let $F: A \times [0, T] \rightarrow B$ be the homotopy such that $f_t = F(\cdot, t): A \rightarrow B$ is the PL-map that agrees with f outside $\text{star}(v)$ and such that $f_t(v) \in \alpha$ with $d(f_t(v), f(v)) = t$. Such a homotopy exists for small t . Moreover, for small t we have:

- (1) For any $e_0 \in N$ and any $e_1 \in K$, $S_{f_t, e_1} < S_{f_t, e_0}$.
- (2)
 - Either $S_{f_t} = S_f$ and $A_{\max}(f_t) \subset A_{\max}(f)$ (but not equal).
 - Or $S_{f_t} < S_f$ and $A_{\max}(f_t) = A_{\max}(f)$.

Remark 3.9. Note that by this result, an optimal map is one where both S_f and $A_{\max}(f)$ are minimal.

Definition 3.10. Let t_0 be the supremum of times t such that f_t exists and has the above properties. We define $\text{Next}_v(f)$ as f_{t_0} .

Note that $\text{Next}_v(f)$ can be defined only for non-optimal maps. We can now prove that the inequality (2) is an equality, as was first proved by Tad White.

Proposition 3.11. *Let A, B be marked metric graphs of rank n . Then there exists an $f \simeq \tau_B \tau_A^{-1}$ and a cyclically reduced loop γ contained in A_{\max} , the subgraph of maximally stretched edges of A , whose f -image is also cyclically reduced. In particular, $\Lambda_R(A, B) = S_f$ for this map f . In particular, this is true for any optimal map.*

Proof: By Lemma 3.4, we may choose a map f whose stretching factor is minimal. Moreover, we may choose such a map with the least number of edges in $A_{\max}(f)$. Hence, $\text{Next}_v(f)$ cannot exist, and therefore f is optimal. Therefore optimal maps exist.

Now let f be any optimal map. This means that any path, p , in $A_{\max}(f)$ which is mapped to a reduced path by f can be continued to a

longer path, which is also mapped to something reduced. This is because the obstruction to continuing p is exactly non-optimality of f . Starting with a single edge, and since there are only finitely many oriented edges in $A_{\max}(f)$, we can find a reduced path of the form $e q e$ which is mapped to a reduced path by f . It is then clear that $\gamma = e q$ is a cyclically reduced loop, which is mapped to something cyclically reduced. Moreover, $l_B(\gamma) = S_f l_A(\gamma)$, and hence $\Lambda_R(A, B) = S_f$ as required. \square

Actually, one can do better.

Definition 3.12. Let $f: A \rightarrow B$ be a PL-map. For any sub-graph A_0 of A , we define $\partial_f A_0$ the f -boundary of A_0 as the set of vertices v of A_0 such that all edges of A_0 terminating at v have f -image with a common terminal partial edge.

So, for example, a map is optimal if and only if $\partial_f A_{\max} = \emptyset$.

Proposition 3.13. *Let A, B be marked metric graphs of rank n . Then there exists an $f \simeq \tau_B \tau_A^{-1}$ such that, if $\lambda_1 > \dots > \lambda_k$ are the stretching factors of edges, if A_i denotes the sub-graph of edges stretched by λ_i , then for all i*

$$\partial_f A_i \subset A_{i-1}, \quad i \geq 1 \quad \partial_f A_1 = \emptyset.$$

(So, heuristically, A_i is a cycle relative to A_{i-1} .)

Proof: Once one finds optimal maps as in Proposition 3.11, choose between them one that has the smallest λ_2 and A_2 , argue as in Proposition 3.11, and conclude inductively on i . \square

We note that implicit in the proof of Proposition 3.11 is a proof that $\Lambda_R(A, B)$ is computable. Namely, the path γ produced at the end of the proof can be chosen minimally, and so we may assume that it passes through each oriented edge at most once. There are only finitely many such paths, and we may compute their lengths in A and B (without reference to f) as well as the maximum of the ratio of these lengths. By the proposition, this maximum will be exactly $\Lambda_R(A, B)$. However, the number of such γ will be exponential in the number of edges. We will now show that it is always possible to find a “less complicated” loop γ , which will cut down the computational complexity considerably. More precisely, using Proposition 3.15, if A is a rose with n petals, then $\Lambda_R(A, B)$ may be obtained by calculating the lengths of at most $2n$ paths.

We will approach this problem in two steps, and the idea of this result is that we want to reduce the complexity of γ as a loop in A . We always have in mind an optimal map f , and so we will assume that

γ lies in A_{\max} . We shall attempt to simplify by cutting and gluing γ to itself. Since we will only use edges that were already in γ , we ensure that our loops are always contained in A_{\max} . In order for the cutting and pasting to result in loops which still give the value for $\Lambda_R(A, B)$, we need to make sure that the resulting image in B is cyclically reduced. Therefore we always need to keep in mind that we are working at two levels. On the one hand we have a loop, γ , thought of as a map from the circle to A (A_{\max} , in fact). We then compose this map with f and the resulting loop in B is an immersion. For the first step of our result, we prove the following ‘‘Sausages Lemma’’, which says that we may take a γ which realises $\Lambda_R(A, B)$ and whose shape in A is as in Figure 1.

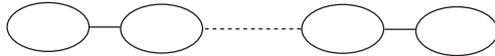


FIGURE 1. Sausages.

For any oriented path γ we denote by $\bar{\gamma}$ its inverse.

Lemma 3.14 (Sausages Lemma). *Let A, B be marked metric graphs of rank n , and let $f \simeq \tau_B \tau_A^{-1}$ be an optimal map. Then there exists a loop γ such that $l_B(\gamma)/l_A(\gamma) = S_f = \Lambda_R(A, B)$. In particular, γ is cyclically reduced in A and in B via f . Furthermore, γ is a sausage, i.e. $\gamma = \gamma_1 \bar{\gamma}_2$ where each γ_i is a path in A that can be parameterised with $[0, 1]$ in such a way that*

- γ_1 and γ_2 are embeddings;
- there exists a finite family of disjoint closed intervals $I_j \subset (0, 1)$, each one possibly consisting of a single point, such that $\gamma_1(t) = \gamma_2(s)$ if and only if $t = s$ and t belongs to $\{0, 1\} \cup_j I_j$.

Proof: The content of the result is that $\gamma = \gamma_1 \bar{\gamma}_2$ has the specified properties, since everything else follows from Proposition 3.11. This will follow from three sublemmas. First we establish some notation. We shall think of γ as a map from S^1 to A and also, via f , as a map from S^1 to B . We shall subdivide S^1 to give it a graph structure and so that edges map to edges in A and vertices map to vertices in A . For simplicity, although it is not really necessary, we shall assume that all the vertices of A map to vertices of B , which we can arrange after a suitable subdivision. Hence edges in S^1 map to edges or vertices in B .

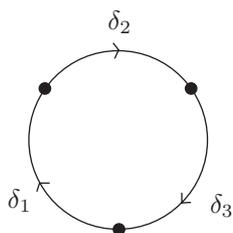


FIGURE 2. Triple points.

Our first sublemma says that if three distinct points in S^1 have the same image in A , then we can choose γ to be shorter (in both A and B). To do this, we look at three points in S^1 mapped to the same point in A . Thus we decompose γ as $\delta_1\delta_2\delta_3$ as in the picture above, where the endpoints of each δ_i map to the same point in A . Our first attempt is to try to replace γ with one of the δ_i , each of which is clearly a shorter path in A , and each of which maps to a reduced path in B . The only way that this can fail is if each δ_i maps to a reduced but not cyclically reduced path in B . This means that we can write,

$$\begin{aligned} \delta_1 &= e_1 \dots \overline{e_1} \\ \delta_2 &= e_2 \dots \overline{e_2} \\ \delta_3 &= e_3 \dots \overline{e_3}, \end{aligned}$$

where we are writing each δ_i as a concatenation of edges *labelled* by the image of that edge in B . Thus we are saying that the image of δ_1 in B begins with an edge e_1 and ends with the inverse edge $\overline{e_1}$. However, we know that γ is immersed in B , so that $e_1 \neq e_2$. In particular, this implies that the loop $\delta_1\delta_2$ is immersed in B , and we are done.

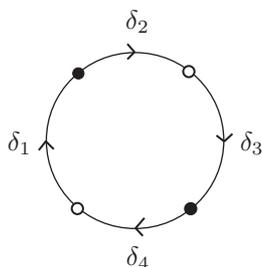


FIGURE 3. Crossing points.

For the second sublemma, we will show that we can avoid crossing double points in γ . That is, if we can write γ as a concatenation $\delta_1\delta_2\delta_3\delta_4$ in S^1 such that the initial points of δ_1 and δ_3 have the same image in A , and the initial points of δ_2 and δ_4 have the same image in A , then we may replace γ by a shorter path (shorter in both A and B).

Now we try to replace γ by one of the paths $\delta_i\delta_{i+1}$ (subscripts taken modulo 4). If any of these map to cyclically reduced loops in B , we are done. Otherwise, we get that,

$$\begin{aligned}\delta_1\delta_2 &= e_1 \dots \overline{e_1} \\ \delta_2\delta_3 &= e_2 \dots \overline{e_2} \\ \delta_3\delta_4 &= e_3 \dots \overline{e_3} \\ \delta_4\delta_1 &= e_4 \dots \overline{e_4},\end{aligned}$$

where, as before, this is a concatenation of edges in S^1 labelled by the images in B . This implies that

$$\delta_i = e_i \dots \overline{e_{i+3}},$$

with subscripts taken mod 4. Since we know that γ is immersed in B , we must have that $e_1 \neq e_3$ and $e_2 \neq e_4$. Thus it is clear that the loop $\delta_1\overline{\delta_3}$ is immersed in B , and hence we have proven the second sublemma.

For our third and final sublemma, we wish to remove all “bad triangles”. This may be slightly confusing terminology, but we wish to avoid the situation where γ is the concatenation of 6 paths, where alternating paths in this decomposition are closed (and the other 3 form a, not necessarily embedded, triangle). Formally, let us assume that we can write

$$\gamma = \delta_1\delta_2\delta_3\delta_4\delta_5\delta_6,$$

where $\delta_1, \delta_3, \delta_5$ are closed paths, and show that this means we can shorten γ . Note that if any of the paths $\delta_1, \delta_3, \delta_5$ are immersed in B then we are done, simply by replacing γ . So let us assume that none of these subpaths are immersed. Using arguments similar to those before, this implies that $\delta_1\delta_2\delta_3\overline{\delta_2}$ is a closed path in A whose image in B is immersed.

We would now be done if this new path were shorter than γ . If this fails, we apply the same argument to the path $\delta_3\delta_4\delta_5\overline{\delta_4}$ which would now be shorter than γ .

Armed with these sublemmas, we may remove all triple points, all crossing points and all bad triangles since there are only finitely many loops less than a given length in A (or B). We subdivide γ (really the source circle S^1) into edges and vertices, labelled by their image in A .

Clearly, the labelling in S^1 may give more than one edge the same label since γ need not be embedded, however since we have removed all triple points each label may occur at most twice. If each label occurs once, γ is embedded and we are done. Otherwise, choose an “innermost” pair of vertices in γ with the same label. These are points in S^1 which map to the same vertex in A (and therefore also in B). That is, choose such a pair u, v and one of the paths between them, δ so that δ embeds in A except at the endpoints. Note that if there is a single pair of points on S^1 mapping to the same vertex in A , then γ will be a figure-eight and we are done, whereas if no such pair exists then γ would be embedded in A and we are also done.

Note that if we have two innermost pairs, we can assume that the paths between them are disjoint as we have removed all crossing points. Also, since we have removed all bad triangles, there are at most two innermost such pairs (in fact there are exactly two, if we also keep track of the path between them and remember that we are assuming that γ is not embedded or a figure-eight). For each innermost pair, choose a point between them (i.e. on the specified path). So we now have two points on γ and therefore two subpaths, γ_1, γ_2 between them and $\gamma = \gamma_1\overline{\gamma_2}$. Since we have no bad triangles and no crossing points, both γ_1 and γ_2 must be embedded in A . We have also divided γ , and hence its subpaths, according to the image in A and use this paramaterisation to finish the lemma. Namely, the disjoint intervals I_j correspond to edges or vertices of A which have more than one pre-image in γ . Since we have eliminated all crossing points in γ , the intervals I_j appear in the same order in both γ_1 and γ_2 are we are done. \square

The final step in simplifying our loop γ is to move from a collection of sausages to at most two.

Proposition 3.15. *Let $A, B \in CV_n$, and let $f \simeq \tau_B\tau_A^{-1}$ be an optimal map. Then there exists a loop γ with $l_B(\gamma)/l_A(\gamma) = \hat{S}_f = \Lambda_R(A, B)$ so that either*

- O: γ is a simple closed curve in A ,
- ∞ : γ is an embedded bouquet of two circles, i.e. $\gamma = \gamma_1\gamma_2$, where γ_i are simple closed curves which do not meet each other, except at a single point, or
- O-O: $\gamma = \gamma_1\gamma_3\gamma_2\overline{\gamma_3}$, where γ_1 and γ_2 are simple closed curves which do not meet, and γ_3 is an embedded path that touches γ_1 and γ_2 at their initial points only.

In particular, there exists a finite set of loops, Γ , in A so that $l_B(\gamma)/l_A(\gamma) = \Lambda_R(A, B)$ for some $\gamma \in \Gamma$ and the set Γ can be chosen independently of B .

Proof: We shall start by taking the loop $\gamma = \gamma_1\overline{\gamma_2}$ supplied by Lemma 3.14. If the family of intervals $\{I_j\}$ is empty, then γ is a simple closed curve; if it consists of a single interval I then γ is either an embedded ∞ - or O-O-curve, depending whether I is a single point or not. In these cases we are done.

Suppose that the family $\{I_j\}$ contains at least two intervals. We show how to reduce to the case of only one interval. Let $[a, b]$ and $[c, d]$ be the two extremal intervals of $\{I_j\}$; namely, such that $0 < a \leq b < c \leq d < 1$ and no I_j in $(0, a) \cup (d, 1)$. We replace the loop γ_2 with the following

$$\delta_2(t) = \begin{cases} \gamma_2(t) & t < b \\ \gamma_1(t) & t \in [b, c] \\ \gamma_2(t) & t > c. \end{cases}$$

Note that δ_2 is embedded in A because $\gamma_1(t) = \gamma_2(s)$ if and only if $t = s$ (by Lemma 3.14). Also, the f -image of δ_2 in B is reduced because of the same reason and because the f -images of both γ_1 and γ_2 are reduced. The new loop $\tilde{\gamma} = \gamma_1\overline{\delta_2}$ is therefore a sausage-loop satisfying $l_B(\tilde{\gamma})/l_A(\tilde{\gamma}) = S_f = \Lambda_R(A, B)$, and the cardinality of the I_j 's is now one. \square

Another interesting consequence of Proposition 3.11 is that Λ_R is always defined and finite. We note that this can also be proved directly using the immersion of paths in the space of geodesic currents (see [9], [12]). Indeed, the space of geodesic currents is compact, and lengths are continuous linear functionals, so the ratio of two length functionals always has maximum and minimum realised by some current. In particular the maximum is finite and the minimum is non-zero, and we have additionally proved that it is realised by a *rational* current.

4. Metrics

We are now in a position to define a metric on CV_n and our starting point will be Definition 3.1. In fact we have both left hand and right hand displacements (whose existence is guaranteed by Proposition 3.11 and the preceding discussion).

Definition 4.1 (Right and left factors). For any pair A, B of marked metric graphs of rank n we set:

$$\Lambda_R(A, B) := \sup_{1 \neq w \in F_n} \frac{l_B(w)}{l_A(w)}, \quad \Lambda_L(A, B) := \sup_{1 \neq w \in F_n} \frac{l_A(w)}{l_B(w)} = \Lambda_R(B, A).$$

Remark 4.2. Since F_n embeds in the space of geodesic currents as a dense sub-space, we could equivalently define Λ_R and Λ_L taking the supremum over the space of currents.

The reason that we wish to use both Λ_R and Λ_L is that they are not symmetric functions and hence if we wish to define a genuine metric on CV_n we will need to use both of them. We are now ready to define the metric on CV_n .

Definition 4.3 (Distance). For all $A, B \in CV_n$, we define

$$\Lambda(A, B) := \Lambda_R(A, B)\Lambda_L(A, B).$$

The *distance* between A and B is then given by,

$$d(A, B) = \log \Lambda(A, B).$$

This is finite by Proposition 3.11.

The first remark is that if we scale the length functions l_A and l_B by positive numbers, $d(A, B)$ remains unchanged. So it is well-defined on CV_n .

Proposition 3.11 shows in fact that $d(A, B)$ is always finite, but we still need to show that it is indeed a metric. We begin with an elementary observation.

Remark 4.4. Given a positive real valued function, f ,

$$\sup \frac{1}{f(x)} = \frac{1}{\inf f(x)}.$$

Moreover, $\sup \frac{1}{f(x)}$ exists if and only if $\inf f(x)$ exists and is non-zero.

This has an easy but interesting consequence for us,

Lemma 4.5.

$$\Lambda(A, B) = \frac{\sup_{1 \neq w \in F} \frac{l_A(w)}{l_B(w)}}{\inf_{1 \neq w \in F} \frac{l_A(w)}{l_B(w)}}.$$

Proof: Apply the previous remark to $\frac{l_A(w)}{l_B(w)}$, noting that $\Lambda_R(A, B)$ always exists. □

It is now immediate that d will be a non-negative function.

Corollary 4.6. *For all $A, B \in CV_n$, $\Lambda(A, B) \geq 1$ and hence $d(A, B) \geq 0$.*

Next we need to show that d is only zero when the two entries are the same point of CV_n .

Lemma 4.7. *Given $A, B \in CV_n$, $d(A, B) = 0$ if and only if $A = B$.*

Proof: Thinking of CV_n as a space of length functions, it is clear that if the two functions, l_A and l_B differ by a multiplicative constant, then $\Lambda(A, B) = 1$ and so $d(A, B) = 0$. Conversely, if $d(A, B) = 0$ then after rescaling (by $\Lambda_R(A, B)$) we get that $l_A = l_B$. \square

Lemma 4.8 (Triangular inequality). *For all marked metric graphs A, B, C of rank n*

$$d(A, C) \leq d(A, B) + d(B, C).$$

Proof: For any $1 \neq g \in F_n$

$$\begin{aligned} \Lambda_R(A, B)\Lambda_R(B, C) &= \sup_{1 \neq w \in F_n} \frac{l_B(w)}{l_A(w)} \sup_{1 \neq w' \in F_n} \frac{l_C(w')}{l_B(w')} \\ &\geq \frac{l_B(g)}{l_A(g)} \frac{l_C(g)}{l_B(g)} \\ &= \frac{l_C(g)}{l_A(g)}. \end{aligned}$$

Thus $\Lambda_R(A, B)\Lambda_R(B, C) \geq \Lambda_R(A, C)$. Using the same argument for Λ_L , we have verified the triangle inequality for d . \square

Since the function d is clearly symmetric, collecting previous lemmata we have a proof of

Theorem 4.9. *The function $d(A, B) = \log \Lambda(A, B)$ defines a metric on CV_n .*

Remark 4.10. It is straightforward that automorphisms of the free group act by isometries on CV_n with respect to d .

Armed with the metric above, we clearly need to verify that the topology it gives is the same as the one we already have on CV_n .

Theorem 4.11 (The topology). *The topology induced by d on CV is the usual one.*

Proof: Note that the Gromov topology, the Axes topology and the Weak topology coincide for Culler-Vogtmann Space, and a fuller discussion may be found here, [10].

Marked metric graphs are characterised by their translation lengths, so elements of CV_n are characterised by the projective classes of their translation lengths.

We show that the two topologies —the metric one and the usual one— have the same convergent sequences, that being enough since both topologies have countable bases.

First, we show that if $d(A_k, A) \rightarrow 0$ then $A_k \rightarrow A$ in CV_n . If $d(A_k, A) \rightarrow 0$, then by Lemma 4.5 the function

$$\frac{\sup}{\inf}(l_{A_k}/l_A)$$

uniformly converges to 1. Therefore, up possibly to rescaling, $l_{A_k} \rightarrow l_A$ pointwise, and thus $A_k \rightarrow A$ as elements of CV_n .

Conversely, if $A_k \rightarrow A$ as elements of CV_n , then, up possibly to rescaling, the length functions of A_k tend to the length function of A . Denote the length function of A_k by $\|\cdot\|_k$ and that of A by $\|\cdot\|$.

Now, consider optimal maps $f_k: A \rightarrow A_k$ as in Proposition 3.11. Now, by definition, the Lipschitz constant for f_k , which we denote h_k , is realised by a loop γ in A , in the sense that $h_k = \|\gamma\|_k/\|\gamma\|$. We need to show that $h_k \rightarrow 1$.

By Proposition 3.15 there are in fact a finite set of loops in A from which we can guarantee that one will realise this Lipschitz constant so we may therefore assume that there is a single γ such that $h_k = \|\gamma\|_k/\|\gamma\|$ for all k . However, by assumption we have that $\|\gamma\|_k \rightarrow \|\gamma\|$ and hence $h_k \rightarrow 1$ as required. \square

Theorem 4.12 (Properness). *For any $X \in CV_n$, any closed d -ball centred at X is compact. Whence (CV_n, d) is complete.*

Proof: Let $\{A_i\}$ be any sequence in CV_n such that $d(X, A_i) \leq e^R$. We show that it has a convergent sub-sequence. By hypothesis we have

$$\frac{\sup}{\inf}(l_{A_i}/l_X) < R$$

and, up to possibly scaling the metric of A_i , we can suppose $\inf(l_{A_i}/l_X) = 1$. Therefore $\{\sup(l_{A_i}/l_X)\}$ is a bounded sequence, and a diagonal argument now shows that, up to possibly passing to subsequences, l_{A_i} has as pointwise limit that we denote by l_∞ . Since the closure of Outer Space is the space of “very small actions”, [8], [2], [6], l_∞ corresponds to a translation length function of a minimal isometric action of the free group F_n on an \mathbb{R} -tree. Since the infimum of functions is upper semicontinuous, l_∞ is bounded below away from zero. We show in Lemma 4.13 that this implies that the action given by l_∞ is actually free on a simplicial tree,

and corresponds therefore to a point A of CV_n which, by Theorem 4.11 is the limit of $\{A_i\}$. \square

Lemma 4.13. *Let l be the translation length function of a minimal isometric action of the free group F_n on a \mathbb{R} -tree T . If $\inf l > c > 0$ then T is simplicial and the action is free.*

Proof: The fact that the action is free is obvious since l is bounded below away from zero. Now suppose, by contradiction, that the action is not simplicial. Then, there is a point $x \in T$ and a sequence of segments σ_k , no three of them co-linear, such that the sequence $\{s_k\}$ of their starting points converges to x . Let \tilde{R}_n denote the universal cover of the standard rose R_n (i.e. \tilde{R}_n is the Cayley graph of F_n) with a marked origin O , and let $f: \tilde{R}_n \rightarrow T$ be a Lipschitz, PL-map which is equivariant with respect to the actions of F_n on \tilde{R}_n and T . Let $y_k \in \tilde{R}_n$ such that $f(y_k) = s_k$. Let $w_k \in F_n$ be elements such that $w_k(y_k)$ stay at distance less than one from O . After passing to a subsequence, we may assume that $w_k(y_k)$ is convergent in \tilde{R}_n , and hence that $w_k(s_k)$ is convergent in T . Looking at distances in T we see that,

$$\begin{aligned} d(w_k(s_k), w_h(s_k)) &\leq d(w_k(s_k), w_h(s_h)) + d(w_h(s_h), w_h(s_k)) \\ &= d(w_k(s_k), w_h(s_h)) + d(s_h, s_k). \end{aligned}$$

Hence, from the remarks above, the translation length of $w_h^{-1}w_k$ in T , tends to zero, as $h, k \rightarrow \infty$. Moreover, since the no three of the σ_k 's are co-linear, the family $\{w_k\}$ is infinite and hence $w_h^{-1}w_k$ cannot always equal the identity. This contradicts the hypothesis that l is bounded away from zero. \square

Since our metric d is the corresponding of a symmetrised version of the Thurston metric on Teichmüller space, it is natural to ask what happens to the non-symmetric pieces.

Definition 4.14. Given $A \in CV_n$ we denote by \bar{A} its representative which has total volume one.

Definition 4.15 (Right and left hand non-symmetric metric). For any $A, B \in CV_n$ we define

$$d_R := \log(\Lambda_R(\bar{A}, \bar{B})), \quad d_L := \log(\Lambda_L(\bar{A}, \bar{B})).$$

Since $d_L(A, B) = d_R(B, A)$ we can restrict our study to the right hand metric d_R . The elementary properties require some more work

than in the case of the symmetric metric. Also note that $d(A, B) = d_R(A, B) + d_R(B, A)$ so is a symmetrisation of the one-sided metric.

Observe that d is well-defined for marked metric graph, and it is scale-invariant, so it descends to a metric on CV_n . This property does not hold for d_R , however, which is why the normalisation to volume one is crucial.

Lemma 4.16. *For any $A, B \in CV_n$ the right hand distance is non-negative and vanishes only if $A = B$:*

$$d_R(A, B) \geq 0 \quad \text{and} \quad d_R(A, B) = 0 \Leftrightarrow A = B \in CV_n.$$

Proof: Let $f: \bar{A} \rightarrow \bar{B}$ be an optimal map (that exists by Proposition 3.11) then

$$(3) \quad 1 = \text{vol}(\bar{B}) = \text{vol}(\text{Im}(f)) \leq \Lambda_R(\bar{A}, \bar{B}) \text{vol}(\bar{A}) = \Lambda_R(\bar{A}, \bar{B})$$

so $d_R(A, B) \geq 0$. If, for any edge e of \bar{A} we denote by $l_{\bar{A}}(e)$ its length (hence $\sum l_{\bar{A}}(e) = 1$) recalling that $S_{f,e}$ denotes the stretching factor of e , we have

$$(4) \quad \text{vol}(\text{Im}(f)) = \sum_{e \text{ edge of } A} S_{f,e} l_{\bar{A}}(e) - C$$

where C is a non-negative quantity that measures the overlappings of f . Therefore, if $\Lambda_R(A, B) = 1$, then the inequality of (3) is an equality, and from (4) we get $S_{f,e} = 1$ for all edges e , and $C = 0$ which together imply that f is an isometry. Thus $\bar{A} = \bar{B}$ as marked graphs, and $A = B$ as elements of CV_n . \square

Ordered triangular inequality is already proven in Lemma 4.8, so we have proved

Theorem 4.17. *The function $d_R(A, B)$ defines a non-symmetric metric on CV_n .*

As for the symmetric case, the topology induced by d_R on CV_n is the usual one.

Theorem 4.18 (The topology). *For any sequence $\{A_k\}$ and $A \in CV_n$*

$$d(A, A_k) \rightarrow 0 \Leftrightarrow d_R(A, A_k) \rightarrow 0 \Leftrightarrow d_R(A_k, A) \rightarrow 0.$$

Proof: Clearly if $d = d_R + d_L \rightarrow 0$ then both d_R and d_L go to zero. Suppose that $d_R(A, A_k) \rightarrow 0$. Let $f_k: \bar{A} \rightarrow \bar{A}_k$ be an optimal map. As in (4) we have

$$1 = \text{vol}(\text{Im}(f_k)) = \left(\sum_{e \text{ edge of } A} S_{f_k,e} l_{\bar{A}}(e) \right) - C_k$$

with $S_{f_k, e} \leq \Lambda_R(A, A_k) \rightarrow 1$ and $\sum l_{\bar{A}}(e) = 1$. Which implies that f_k converges to an isometry and therefore $d(A, A_k) \rightarrow 0$. A similar argument works for when $\Lambda_R(A_k, A) \rightarrow 1$. \square

The first important difference between symmetric and non-symmetric metrics is that the latter are not complete. Therefore, in general, the fact that a sequence is a right hand Cauchy sequence does not guarantee convergence in CV_n .

Theorem 4.19 (Incompleteness). *The space (CV_n, d_R) is not complete. Namely there are sequences $\{A_k\}$ such that $d_R(A_k, A_{k+m}) \rightarrow 0$ as $k \rightarrow \infty$ which have no accumulation point. Moreover, for any $A \in CV_n$ and any $B \in \overline{CV_n} \setminus CV_n$ one has that $\Lambda_R(A, B) < \infty$.*

Proof: Let A_0 be R_n the standard n -petals rose with a uniform metric of volume one. Let A_k be the graph obtained by multiplying the metric of one petal by a factor $1/k$ and normalised to have volume one. Then, a direct calculation shows

$$\Lambda_R(A_k, A_{k+m}) = \frac{((k+m)n-1)k}{(k+m)(kn-1)}$$

which goes to 1 as $k \rightarrow \infty$. Thus, $\{A_k\}$ is a right hand Cauchy sequence, but its only accumulation point is the standard rose with $n-1$ petals which does not belong to CV_n —but it can be viewed as an element of $\overline{CV_n}$.

In order to prove the second statement, one simply constructs a PL, equivariant map from A to B . This is guaranteed to be Lipschitz, since A is in CV_n (for any choice of B). Whence $\Lambda_R(A, B)$ is bounded. \square

Remark 4.20. Theorem 4.19 points out another “pathology” of the non-symmetric metrics. Indeed, consider a volume-one, marked metric graph A , and a sequence B_k of volume-one, marked metric graphs such that $\Lambda_R(A, B_k)$ goes to infinity. This can be easily done using iterations of automorphisms (see for instance Section 8). Then, up to possibly passing to a subsequence, $B_k \rightarrow B$ a point in $\overline{CV_n} \setminus CV_n$. By Theorem 4.19 we have $\Lambda_R(A, B) < \infty$ and $\Lambda_R(A, B_k) \rightarrow \infty$.

On the other hand, right and left hand metrics are more deeply related to folding procedures, this providing an easy description of geodesics.

We note that one interesting consequence of the existence of the metric, is that one can use it to prove the Bounded Cancellation Lemma of [7].

The Bounded Cancellation Lemma, first proved by Cooper, is a key result in the study of automorphisms of free groups. It has many equivalent formulations, of which we state one.

Theorem 4.21 (Bounded Cancellation Lemma, [7]). *Let A, B be marked metric graphs of rank n , and consider $f: A \rightarrow B$, a PL map such that $f \simeq \tau_B \tau_A^{-1}$. Let $|\cdot|_A$ and $|\cdot|_B$ denote the length functions of A and B respectively (the edge lengths of paths). (Note that this is not quite the translation length, since we do not cyclically reduce.) Let α, β be loops in A , at a vertex v , such that $|\alpha\beta|_A = |\alpha|_A + |\beta|_A$. Then, there exists a constant K depending only on A and B (and not on α, β) such that,*

$$|f(\alpha\beta)|_B \geq |f(\alpha)|_B + |f(\beta)|_B - 2K.$$

We call K a bounded cancellation constant for the map, f , which clearly only depends on f up to homotopy relative to vertices.

We observe that the existence of the bounded cancellation constant is related to our (left) distance.

Proposition 4.22. *Given A, B and f as above, let λ be the Lipschitz constant for f . Then if i is not a bounded cancellation constant for f , we may find loops α_i, β_i at a vertex v of A such that*

- (1) $|\alpha_i\beta_i|_A = |\alpha_i|_A + |\beta_i|_A$.
- (2) $|f(\alpha_i\beta_i)|_B < |f(\alpha_i)|_B + |f(\beta_i)|_B - 2(i - \lambda \text{vol}(A))$.
- (3) $|f(\alpha_i)|_B \leq \lambda \text{vol}(A) + i, |f(\beta_i)|_B \leq \lambda \text{vol}(A) + i$.

Moreover, we can ensure that $\alpha_i\beta_i$ is cyclically reduced in A .

Proof: By hypothesis, we may find loops α_i, β_i such that $|f(\alpha_i\beta_i)|_B < |f(\alpha_i)|_B + |f(\beta_i)|_B - 2i$. This means that there is a terminal segment of $f(\alpha_i)$ cancels with an initial segment of $f(\beta_i)$ of length i (though the cancellation may be longer). So consider a terminal segment of α_i and an initial segment of β_i whose images cancel and have length i .

Now, by adding a segment of length not greater than $\text{vol}(A)$ to each of these pre-images, we may replace α_i, β_i by paths which are loops, (which we continue to call α_i, β_i) so that $\alpha_i\beta_i$ is cyclically reduced in A .

By construction, $f(\alpha_i)$ is a loop in B which is the original cancellation segment of length i , followed by a path which is the image of something of length at most $\lambda \text{vol}(A)$. Since the image of this terminal segment has length at most $\lambda \text{vol}(A)$, we know that a terminal segment of $f(\alpha_i)$ of length at least $i - \lambda \text{vol}(A)$ survives (and is a terminal segment of the original cancellation segment). By a similar argument for $f(\beta_i)$, we may

deduce that a segment of length at least $i - \lambda \text{vol}(A)$ must cancel in $f(\alpha_i\beta_i)$. Therefore, $|f(\alpha_i\beta_i)|_B < |f(\alpha_i)|_B + |f(\beta_i)|_B - 2(i - \lambda \text{vol}(A))$.

Moreover, by construction, $|f(\alpha_i)|_B \leq \lambda \text{vol}(A) + i$, $|f(\beta_i)|_B \leq \lambda \text{vol}(A) + i$ and we are done. \square

Now, consider two loops in A , α , β , which are based at the same vertex of A , such that $\alpha\beta$ is cyclically reduced and $|\alpha\beta|_A = |\alpha|_A + |\beta|_A$, and with the additional condition that $|\alpha|_A, |\beta|_A \leq 4\lambda \text{vol}(A)\Lambda_L(A, B)$. Let

$$K_{\alpha,\beta} = \frac{|f(\alpha)|_B + |f(\beta)|_B - |f(\alpha\beta)|_B}{2}$$

since there are only finitely many pairs, α , β with the above properties, we may find a maximum K of the numbers $K_{\alpha,\beta}$.

Corollary 4.23. *With the above notation, the number $K + \lambda \text{vol}(A)$ is a bounded cancellation constant for f .*

Proof: Recall that

$$\frac{1}{\Lambda_L(A, B)} = \inf_w \frac{\|w\|_B}{\|w\|_A}$$

and that $\|w\| \leq |w|$ with equality if and only if w is cyclically reduced. In particular, whenever $\alpha\beta$ is cyclically reduced, we have

$$\frac{|f(\alpha\beta)|_B}{|\alpha\beta|_A} \geq \frac{\|f(\alpha\beta)\|_B}{\|\alpha\beta\|_A} \geq \frac{1}{\Lambda_L(A, B)}.$$

By Proposition 4.22, if i is not a bounded cancellation constant for f , we may find α , β such that $\alpha\beta$ is cyclically reduced, the cancellation in $f(\alpha\beta)$ is greater than $i - \lambda \text{vol}(A)$, and $|f(\alpha)|_B, |f(\beta)|_B \leq \lambda \text{vol}(A) + i$.

So we get $|f(\alpha\beta)|_B \leq 4\lambda \text{vol}(A)$ and

$$|f(\alpha\beta)|_B \Lambda_L(A, B) \geq |\alpha\beta|_A$$

whence $|\alpha\beta|_A \leq 4\lambda \text{vol}(A)\Lambda_L(A, B)$ and thus $K_{\alpha,\beta} \leq K$.

Since the cancellation in $f(\alpha\beta)$ is greater than $i - \lambda \text{vol}(A)$

$$|f(\alpha\beta)|_B \leq |f(\alpha)|_B + |f(\beta)|_B - 2(i - \lambda \text{vol}(A))$$

whence $i - \lambda \text{vol}(A) \leq K$. \square

For deeper bounded cancellation results, see [3].

5. Folding paths and geodesics

In this section we study properties of geodesics and metric properties of folding paths for the symmetric and the non-symmetric metrics.

First we define a geodesic for a non-symmetric metric.

Definition 5.1. Let d_R be a non-symmetric metric on a space X . Then a path $\gamma: [a, b] \rightarrow X$ is called a geodesic if for all $a \leq s \leq t \leq b$, $t - s = d_R(\gamma(s), \gamma(t))$. More generally, call a path γ a geodesic if it can be re-parameterised so that it satisfies the equality above.

The following lemma provides an easy characterisation of geodesics.

Lemma 5.2. *Let γ be a continuous path from an interval $[a, b]$ to a (possibly non-symmetric) metric space. If for any three points $x < y < z \in [a, b]$, the path γ realises the triangular equality*

$$d(\gamma(x), \gamma(y)) + d(\gamma(y), \gamma(z)) = d(\gamma(x), \gamma(z)),$$

then γ is a geodesic.

Proof: Given a subdivision $a = t_0 < t_1 < \dots < t_n = b$ of $[a, b]$, the sum $\sum_{i=1}^n d(\gamma(t_{i-1}), \gamma(t_i))$ approximates the length of γ as the subdivision is finer and finer. By the triangular equality we get

$$\sum_{i=1}^n d(\gamma(t_{i-1}), \gamma(t_i)) = d(\gamma(t_0), \gamma(t_2)) + \sum_{i=3}^n d(\gamma(t_{i-1}), \gamma(t_i))$$

and inductively we conclude that γ is rectifiable and that its length realises the distance between $\gamma(a)$ and $\gamma(b)$. □

Corollary 5.3. *Let $A_t, t \in [a, b]$ denote a continuous path in CV_n . Suppose that for each $x, y, z \in [a, b]$ there is a loop γ which is maximally stretched both from A_x to from A_y and A_y to A_z . More precisely, suppose that*

$$\max_w \frac{l_{A_y}(w)}{l_{A_x}(w)} = \frac{l_{A_y}(\gamma)}{l_{A_x}(\gamma)}, \quad \max_w \frac{l_{A_z}(w)}{l_{A_y}(w)} = \frac{l_{A_z}(\gamma)}{l_{A_y}(\gamma)}.$$

Then A_t is a d_R -geodesic.

Proof: It is immediate to check that A_t realises the (oriented) triangular equality. □

The very same argument gives the following.

Corollary 5.4. *Let $A_t, t \in [a, b]$ denote a continuous path in CV_n . Suppose that for each $x, y, z \in [a, b]$ there are loops γ and η which are respectively maximally and minimally stretched both from A_x to A_y and A_y to A_z . More precisely, suppose that*

$$\begin{aligned} \max_w \frac{l_{A_y}(w)}{l_{A_x}(w)} &= \frac{l_{A_y}(\gamma)}{l_{A_x}(\gamma)} & \max_w \frac{l_{A_z}(w)}{l_{A_y}(w)} &= \frac{l_{A_z}(\gamma)}{l_{A_y}(\gamma)}; \\ \min_w \frac{l_{A_y}(w)}{l_{A_x}(w)} &= \frac{l_{A_y}(\eta)}{l_{A_x}(\eta)} & \min_w \frac{l_{A_z}(w)}{l_{A_y}(w)} &= \frac{l_{A_z}(\eta)}{l_{A_y}(\eta)}. \end{aligned}$$

Then A_t is a d -geodesic.

Remark 5.5. Since $d = d_R + d_L$, a path is d -geodesic if and only if it is both d_R - and d_L -geodesic.

We are now ready to construct d_R -geodesics using scalings and folding paths.

Theorem 5.6 (Right hand geodesics). *For each A, B in CV_n there is a d_R -geodesic path between them, that is to say a continuous path $t \mapsto A_t$ such that $d_R(A, A_t) = t$ and $A_{d_R(A, B)} = B$.*

Proof: Recall that \bar{A} and \bar{B} denote the volume-one representatives in their respective projective classes. Let $f: \bar{A} \rightarrow \bar{B}$ be an optimal map, let $\gamma \subset \bar{A}_{\max}$ be a loop realising $\Lambda_R(\bar{A}, \bar{B})$. Namely, γ is a geodesic in \bar{A} (i.e. a reduced loop) whose f -image is geodesic (i.e. reduced) in \bar{B} , and such that γ is uniformly stretched by f exactly by $\Lambda_R(\bar{A}, \bar{B})$. The existence of such f and γ is ensured by Proposition 3.11.

Let A' be the marked metric graph obtained from \bar{A} by shrinking each edge so that it is stretched by f exactly by $\Lambda_R(\bar{A}, \bar{B})$, and let A_0 be the graph homothetic to A' so that $\Lambda_R(A_0, \bar{B}) = 1$. We still denote by f the induced map $f: A_0 \rightarrow \bar{B}$.

Note that we still have that γ is a reduced loop in A_0 whose f -image is reduced, and that it realises the maximal stretching factor $\Lambda_R(A_0, \bar{B}) = 1$. Also, note that now f stretches each edges of A_0 exactly by 1 (that is to say, f is an isometry on edges).

We describe now a folding procedure that will produce our geodesic. (The idea is that we never touch γ , so that it will realise the maximally stretched loop between any two points of the folding, so that we can invoke Corollary 5.3.)

First, we subdivide —allowing valence-two vertices— both A_0 and \bar{B} so that f is simplicial (i.e. vertices to vertices, edges to edges). For each vertex v of A_0 and $t \geq 0$ let $\sim_{t,v}$ be the equivalence relation on A_0

defined by:

$$x \sim_{t,v} y$$

if and only if $f(x) = f(y)$ and both x and y lie at distance less or equal than t from v . Let \sim_t be the union of all relations $\sim_{t,v}$ as v varies on the set all vertices of A_0 . For $t \geq 0$ we define

$$A_t := A_0 / \sim_t$$

we denote by p_t the projection $A_0 \rightarrow A_t$, and we denote by f_t the map $A_t \rightarrow \bar{B}$ induced by f , which is well-defined since $x \sim_{t,v} y$ implies $f(x) = f(y)$.

For small times t , A_t is obtained from A_0 just identifying germs of edges having the same image under f (local folding). Let t_1 be the smallest time t such that a pair of edges of A_0 is completely identified in A_t .

Our first claim is that, for $t \in [0, t_1]$, A_t is a metric graph and that f_t is an homotopy equivalence, whence A_t is a marked metric graph. The fact that A_t is a graph is because for any segment σ in \bar{B} , $f^{-1}(\sigma)$ is a finite union of segments, and therefore A_t is the result of identifications of a finite number of segments. The fact that f_t is a homotopy equivalence follows from the fact that f factorises as

$$f: A_0 \xrightarrow{p_t} A_t \xrightarrow{f_t} \bar{B}$$

and from the fact that $(p_t)_*: \pi_1(A_0) \rightarrow \pi_1(A_t)$ is surjective.

Our second claim is now that γ realises both $\Lambda_R(A_0, A_t)$ and $\Lambda_R(A_t, \bar{B})$. First note that, as the f -image of γ is geodesic, then also its p_t -image is. Thus $\Lambda_R(A_0, A_t)$ is greater or equal to the ratio $l_{A_t}(\gamma)/l_{A_0}(\gamma)$, which is one because p_t is a local isometry on edges. A fact that also implies that $\Lambda_R(A_0, A_t) \leq 1$. Thus $\Lambda_R(A_0, A_t) = 1$ and it is realised by γ . A similar argument shows that $\Lambda(A_t, \bar{B}) = 1$ is realised by γ .

We argue now by induction. As above, we define relations $\sim_{t-t_1,v}$ for each vertex v of A_1 , and \sim_{t-t_1} as their union. For $t \geq t_1$, we set $A_t := A_{t_1} / \sim_{t-t_1}$, we let $p_t: A_0 \rightarrow A_t$ be the projection, and f_t be map induced by f .

As above, it is easy to check that we have that $\Lambda_R(A_0, A_t) = \Lambda_R(A_t, \bar{B}) = 1$ are both realised by γ .

Our third claim is that such a process ends in a finite time. Indeed, since our folding is isometric on edges, for $t < s$ we can bound below the difference of volumes

$$\text{vol}(A_t) - \text{vol}(A_s)$$

by $s - t$.

So we must stop at a time, say \bar{t} . Since we stopped, at each vertex the folding relations are trivial, but this simply means that $f_{\bar{t}}$ is an isometry.

Summarising, we have constructed a path A_t for $t \in [0, \bar{t}]$ with the property that, A_0 is in the class of A' as element of CV_n , $A_{\bar{t}} = \bar{B}$ is in the class of B as element of CV_n and for each t $\Lambda_R(A_0, A_t)$ and $\Lambda_R(A_t, \bar{B})$ are realised by the same γ . This last property does not change if we rescale each A_t to its volume-one multiple \tilde{A}_t . Therefore, for each t we have

$$d_R(A', B) = d_R(A', A_t) + d_R(A_t, B).$$

Now, note that for any $0 \leq s < t \leq \bar{t}$, if we construct a folding path from A_s to A_t following the above rules, we find exactly the restriction of the folding path we build so far. Therefore, the path A_t from A_0 to B realises the triangular equality, and is therefore d_R -geodesic by Lemma 5.2.

The shrinking procedure from A to A' also realises the triangular equality because everything is shrank and γ is not touched. Finally, if we consider a point X between A and A' and a point Y on the geodesic between A' and B , again we have that every loop is stretched less than $\Lambda_R(A, B)$ and γ is stretched exactly by $\Lambda_R(A, B)$. In conclusion, γ always realises the maximum stretching factor between any two points in the path we constructed. Such a path is then d_R -geodesic by Corollary 5.3. \square

Since it is of independent interest, we formalise the precise definition and notation the folding procedure described in the proof of Theorem 5.6.

Definition 5.7 (Fast folding paths and turns). Let A, B be two marked metric graphs, let $f: A \rightarrow B$ be an optimal map, and let A_0, \bar{B} as in the proof of Theorem 5.6.

A *fast folding path* is a path $t \mapsto A_t$ constructed following the procedure described in the proof of Theorem 5.6.

A fast folding path comes with the simplicial subdivisions and the sequence of times $0 = t_0 < t_1 < \dots < \bar{t}$ such that in each $[t_i, t_{i+1}]$ a whole segment is identified.

A *turn* τ at a time t of a fast folding path is a pair of edges having a common end-point and whose germs are identified for $t' > t$. We say that the turn is folded, or that τ is a folding turn.

Remark 5.8. The folding path we constructed in the proof of Theorem 5.6 is not unique in general, as in general we can start with many different optimal maps. Also, we could start folding separately at every single

vertex, still obtaining a geodesic. This shows that d_R -geodesics between points of CV_n are not unique.

We analyse now the local structure of geodesics in the PL structure of CV_n .

Definition 5.9 (Simplices of Outer Space). A simplex of CV_n is a subset of CV_n consisting of all marked metric graphs with fixed topological type and marking.

Given a marked graph with edges e_1, \dots, e_k , the corresponding simplex σ is identified with the positive cone of \mathbb{R}^k just by assigning the metric, i.e. a length for each edge:

$$A \in \sigma \longleftrightarrow (l_A(e_1), \dots, l_A(e_k)).$$

Similarly, we can assign to each loop, its counting vector. Namely, for a loop ξ let $\xi(e_i)$ be the number of occurrences of the edge e_i in ξ ; then

$$\xi \mapsto (\xi(e_1), \dots, \xi(e_k)).$$

This viewpoint generalises immediately to the setting of geodesic currents (see [11], [12], [9]) and in fact it is in that setting that linear structures arises naturally. Nevertheless, since the use of currents is not strictly necessary for our purposes, we stick to the world of loops.

The local linear structures of CV_n and the space of loops have as consequence that we can handle length as a linear function

$$L_A(\xi) = \langle A, \xi \rangle := \langle (l_A(e_1), \dots, l_A(e_k)), (\xi(e_1), \dots, \xi(e_k)) \rangle$$

where the last scalar product is the standard one of \mathbb{R}^k .

Proposition 5.10. *Segments in simplices of CV_n are d_R - and d_L -, whence d -, geodesics.*

Proof: Let A, B marked metric graphs in the same simplex. Let ξ be a loop that realises $\sup_w l_B(w)/l_A(w)$. The segment between A and B is parameterised by $A_t = (1 - t)A + tB$ (as vectors of \mathbb{R}^k). For any $0 \leq s < t \leq 1$ we have

$$\frac{l_{A_t(w)}}{l_{A_s(w)}} = \frac{(1 - t)\langle A, w \rangle + t\langle B, w \rangle}{(1 - s)\langle A, w \rangle + s\langle B, w \rangle} = \frac{(1 - t) + t(\langle B, w \rangle / \langle A, w \rangle)}{(1 - s) + s(\langle B, w \rangle / \langle A, w \rangle)}.$$

The function

$$x \mapsto \frac{1 - t + tx}{1 - s + sx}$$

is monotone increasing for $t > s$. Thus, for any $s < t$ the stretching factor $l_{A_t(w)}/l_{A_s(w)}$ is maximal on ξ . The thesis now follows from Corollary 5.3. □

Example 5.11. Two points of the same simplex are connected by several geodesics.

Proof: In CV_2 , consider the simplex of the trivalent graph with a disconnecting edge (i.e. a O-O graph). Let A be the vector $(1, 1, 1)$ where the middle coordinate is referred to the disconnecting edge. Let $B = (\lambda, 1, \lambda^{-1})$ with $\lambda > 1$, and let $C = (\lambda, 1, 1)$. Let γ_1 be the segment between A and B . Let γ_2 be the union of the segment between A and C and the one between C and B . Using Corollary 5.3 it is readily checked that γ_1 and γ_2 are different geodesics between A and B . \square

6. The symmetric metric is not geodesic

In this section, we describe an example of two points in CV_2 which are not connected by a d -geodesic.

This example is due to Bert Wiest and Thierry Coulbois.

Consider the Outer Space in rank two, with graphs normalised to have volume one, and where we denote the generators of the free group of rank two by a and b . Consider two simplices of maximal dimension in CV_2 corresponding to graphs without disconnecting edges (theta-graphs) such that they touch along a 1-dimensional simplex corresponding to a rose with two petals. Let X and Y be two points metric graphs, one in each simplex, as shown in Figure 4.

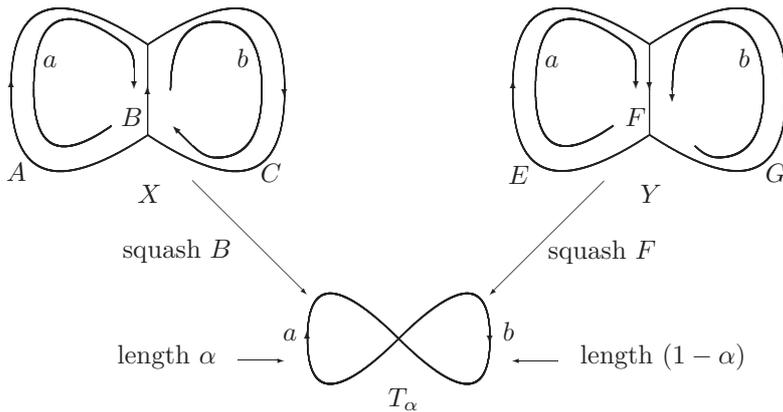


FIGURE 4. The graphs X, Y and T_α .

Since each 1-simplex disconnects CV_2 , any path between X and Y must cross the edge common to the two simplices. We parameterise such

edge by a number α , so that in the graph T_α the petal corresponding to a has length α , and the one corresponding to b has length $1 - \alpha$ (see Figure 4).

By proposition 5.10 a d_R -geodesic between X and Y reduces to the union of two segments XT_α and $T_\alpha Y$, for some α . By Remark 5.5, if there is a d -geodesic between X and Y , there exists α such that $XT_\alpha \cup T_\alpha Y$ is both d_R - and d_L -geodesic. It is readily checked that $XT_\alpha \cup T_\alpha Y$ is d_R -geodesic if and only if there is a loop which is maximally stretched from X to T_α , from T_α to Y and from X to Y (so that the triangular inequality become equality). The same holds for d_L .

We choose now X and Y in a suitable way, we compute the α so that $XT_\alpha \cup T_\alpha Y$ is d_R -geodesic and we show that for such α , $XT_\alpha \cup T_\alpha Y$ is not d_L -geodesic.

We choose X and Y in a symmetric way with respect the common edge:

$$\begin{aligned} X : A = 1/6 \quad B = 1/3 \quad C = 1/2 \\ Y : E = 1/2 \quad F = 1/3 \quad G = 1/6. \end{aligned}$$

We compute now the right factors $\Lambda_R(X, T_\alpha)$ and $\Lambda_R(X, Y)$. By Proposition 3.15 we have to check only the lengths of the loops AB , BC , AC .

Loop in X	AB	BC	AC
Length in X	$1/2$	$5/6$	$2/3$
Length in T_α	α	$1 - \alpha$	1
l_{T_α}/l_X	2α	$6(1 - \alpha)/5$	$3/2$
Corresponding loop in Y	EF	GF	$EFGF$
Length in Y	$5/6$	$1/2$	$4/3$
l_Y/l_X	$5/3$	$3/5$	2
Loop maximally stretched from X to Y			*

It follows that AC must be the maximally stretched also from X to T_α , whence we get

$$3/2 \geq 2\alpha \quad \text{and} \quad 3/2 \geq 6(1 - \alpha)/5$$

that is

$$\alpha \leq 3/4.$$

We compute now $\Lambda_R(T_\alpha, Y)$. By Proposition 3.15 we have only to check the loops a , b , ab , ab^{-1}

Loop in T_α	a	b	ab	ab^{-1}
Length in T_α	α	$1 - \alpha$	1	1
Corresponding loop in Y	EF	GF	EG^{-1}	$EFGF$
Length in Y	$5/6$	$1/2$	$2/3$	$4/3$
l_Y/l_{T_α}	$5/6\alpha$	$1/2(1 - \alpha)$	$2/3$	$4/3$

thus, since ab^{-1} must be the maximally stretched loop, we get

$$4/3 \geq 5/6\alpha \quad \text{and} \quad 4/3 \geq 1/2(1 - \alpha)$$

that is

$$\alpha \geq 5/8 \quad \text{and} \quad \alpha \leq 5/8.$$

We therefore conclude that any d_R -geodesic between X and Y must cross the 1-simplex at the point $T_{5/8}$. The completely symmetric calculation shows that any d_L -geodesic must cross the central edge at the point $T_{3/8}$. Thus no path from X to Y can be simultaneously d_R - and d_L -geodesics. It follows that no d -geodesic in CV_2 joins X and Y .

7. Quasi-geodesics

In Section 5 we have seen how to construct folding paths that are d_R -geodesic. In this section we address the question of whether such paths are quasi-geodesic for the symmetric metric, with constants depending only on the rank. In other words, we ask whether two points of Outer Space can be joined by a quasi-geodesic with uniform constants.

To start, we recall the definition of a quasi-geodesic path.

Definition 7.1. A path, $\alpha: I \rightarrow X$, where I is a real interval and (X, d) is a metric space, is called a (λ, ϵ) quasi-geodesic if for every $x, y \in I$,

$$\frac{1}{\lambda}|x - y| - \epsilon \leq d(\alpha(x), \alpha(y)) \leq \lambda|x - y| + \epsilon.$$

The following lemma is tautological.

Lemma 7.2. *Let α be a path from an interval to a metric space. Suppose that there is a constant C such that*

$$d(\alpha(x), \alpha(y)) > C \cdot \text{length}(\alpha|_{[x,y]}).$$

Then the arc-length reparameterisation of α is bi-lipschitzian with constants $C, 1$. In particular, it is a $(C^{-1}, 0)$ quasi-geodesic.

Theorem 7.3 (4 point property). *Let A, B be two marked metric graph of the same rank. Let α be a d_R -geodesic from A to B constructed as in Theorem 5.6. Then for every $s \leq x \leq y \leq t$ we have*

$$d(\alpha(s), \alpha(t)) \geq d(\alpha(x), \alpha(y)).$$

Proof: Let us denote by l_p the length function of the point $\alpha(p)$. We consider the folding paths constructed before the rescaling to volume 1, so that while volume is not constant along the path, for every $p < q$, $\Lambda_R(\alpha(p), \alpha(q)) = 1$. Thus, the distance between $\alpha(p)$ and $\alpha(q)$ is exactly the logarithm of $\Lambda_L(\alpha(p), \alpha(q)) = \sup l_p/l_q$. Now we look at the points $s \leq x \leq y \leq t$. As in Proposition 3.11, there exists a μ which realises $\Lambda_L(\alpha(x), \alpha(y))$. Next we realise μ as an immersed path in $\alpha(s)$. The folding path itself has two parts, one in which we shrink the lengths of certain edges, and another in which we isometrically identify edges-folding. In either of these parts it is clear that the length of μ can never increase as we travel along the path. Thus,

$$l_s(\mu) \geq l_x(\mu) \geq l_y(\mu) \geq l_t(\mu).$$

In particular,

$$\sup l_s/l_t \geq l_s(\mu)/l_t(\mu) \geq l_x(\mu)/l_y(\mu) = \sup l_x/l_y,$$

and thus $d(\alpha(s), \alpha(t)) \geq d(\alpha(x), \alpha(y))$, as required. □

Proposition 7.4. *Let γ be a path with the 4 point property. Suppose that γ is a finite union of pieces which are quasi-geodesic. Then γ is a quasi-geodesic with constants depending on the constants of the pieces and on the number of the pieces.*

More precisely, if γ is the path with the 4 point property which is the concatenation of n (λ, ϵ) quasi-geodesics, then γ is a $(n\lambda, n\epsilon)$ quasi-geodesic.

Proof: By hypothesis, there exist numbers $x_0 \leq x_1 \leq \dots \leq x_n$ such that γ is a map from the interval $[x_0, x_n]$ and that each restriction, $\gamma|_{[x_i, x_{i+1}]}$ is a (λ, ϵ) quasi-geodesic (we assume that $n > 1$ since otherwise there is nothing to prove). Now consider $p \leq q \in [x_0, x_n]$, and find i, j such that $p \leq x_i \leq x_j \leq q$ so that i is minimal and j is maximal (note that $i \geq 1$

and $j \leq n - 1$). It is clear that,

$$\begin{aligned} d(\gamma(p), \gamma(q)) &\leq d(\gamma(p), x_i) + \sum_{k=0}^{j-i-1} d(x_{i+k}, x_{i+k+1}) + d(\gamma(x_j), q) \\ &\leq \lambda(x_i - p) + \lambda \sum_{k=0}^{j-i-1} (x_{i+k+1} - x_{i+k}) \\ &\quad + \lambda(q - x_j) + (2 + j - i)\epsilon \\ &\leq \lambda(q - p) + n\epsilon. \end{aligned}$$

For the other inequality we note that, using the x_r , we have divided the interval $[p, q]$ into at most n pieces. Thus, one of these pieces is of length at least $(q - p)/n$. Now, suppose that $x_{i+k+1} - x_{i+k} \geq (q - p)/n$. Then, by the 4 point property,

$$\begin{aligned} d(\gamma(p), \gamma(q)) &\geq d(x_{i+k}, x_{i+k+1}) \\ &\geq (x_{i+k+1} - x_{i+k})/\lambda - \epsilon \\ &\geq (q - p)/n\lambda - \epsilon. \end{aligned}$$

Clearly, the same argument works if either $x_i - p \geq (q - p)/n$ or $q - x_j \geq (q - p)/n$. \square

Example 7.5. There are metric spaces with no rectifiable, non-constant paths having the 4 point property.

Proof: Consider the space $L^2([0, 1])$ of the square-summable functions on $[0, 1]$. Let $f: [0, 1] \rightarrow L^2([0, 1])$ be the embedding

$$t \mapsto \chi_{[0, t]},$$

where $\chi_{[0, t]}$ denotes the characteristic function of the set $[0, t]$. Let d the f -pull-back metric on $[0, 1]$:

$$d(s, t) = \sqrt{t - s}.$$

It is straightforward to check that $([0, 1], d)$ has the 4 point property and no rectifiable, non-constant paths. \square

By Theorem 7.3 and Proposition 7.4, to check whether a right geodesic between two points A and B , constructed as in Theorem 5.6, is a quasi-geodesic (with uniform constants not depending on A and B ,) it is enough to check whether the fast folding path from A_0 to \bar{B} is a quasi-geodesic.

Definition 7.6 (Multiplicities). Let $A_t \neq B$ be any point in a fast folding path. The *multiplicity* of a turn τ in a loop γ is the number $\mu_{\tau,t}(\gamma)$ of occurrences of τ turn in γ (counted without any orientation).

The *folding multiplicity* of γ is the sum $\mu_t(\gamma)$ of the multiplicities of all folding turns (see Definition 5.7) in γ :

$$\mu_t(\gamma) = \sum_{\tau} \mu_{\tau,t}(\gamma).$$

In order to use Lemma 7.2, we need to estimate the local speed of a fast folding path. A folding path is PL, and therefore smooth in all but finitely many points (w.r.t. the PL-structure of CV_n). In particular, the right-derivative of the (non-symmetric) distance is always defined, and its integral gives the total length of the path. In a situation where $d = d_L$, such derivative represents the local speed of the fast folding path.

Lemma 7.7 (Local speed of a folding path). *Let $t \mapsto A_t$ be a fast folding path. Then, its local speed is*

$$\frac{2\mu_t(\gamma)}{l_{A_t}(\gamma)}$$

where γ is a folded loop minimising $l_{A_t}(\gamma)/\mu_t(\gamma)$.

Proof: Recall that in our situation (isometric folding as in Theorem 5.6,) we have $d = d_L$. This is because we are defining paths with non-constant volume, so it is possible to maintain $d_R = 0$ while traversing the path. Obviously, in order to do this we apply the definition of the non-symmetric metric directly without rescaling.

Therefore, for small enough ε , the distance between $A_{t+\varepsilon}$ and A_t is given by

$$d(A_{t+\varepsilon}, A_t) = \log \left(\sup_{\xi} \frac{l_{A_t}(\xi)}{l_{A_{t+\varepsilon}}(\xi)} \right) = \log \left(\sup_{\xi} \frac{l_{A_t}(\xi)}{l_{A_t}(\xi) - 2\mu_t(\xi)\varepsilon} \right)$$

which is thus realised by a loop γ minimising $l_{A_t}(\gamma)/\mu_t(\gamma)$. Note that γ can be always chosen to be simple, because if $\gamma = \alpha + \beta$ we would have $l_{A_t}(\gamma) \geq l_{A_t}(\alpha) + l_{A_t}(\beta)$ and $\mu_t(\gamma) = \mu_t(\alpha) + \mu_t(\beta)$, whence

$$\min \left\{ \frac{l_{A_t}(\alpha)}{\mu_t(\alpha)}, \frac{l_{A_t}(\beta)}{\mu_t(\beta)} \right\} \leq \frac{l_{A_t}(\gamma)}{\mu_t(\gamma)}.$$

Therefore, the speed (as right-derivative) is given by

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \frac{d(A_{t+\varepsilon}, A_t)}{\varepsilon} &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \log \left(\frac{l_{A_t}(\gamma)}{l_{A_{t+\varepsilon}}(\gamma)} \right) \\ &= \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \log \left(\frac{l_{A_t}(\gamma)}{l_{A_t}(\gamma) - 2\mu_t(\gamma)\varepsilon} \right) = \frac{2\mu_t(\gamma)}{l_{A_t}(\gamma)}. \quad \square \end{aligned}$$

Another quantity we need to estimate during a folding procedure, is the speed we are approaching the final point B , defined as minus the right-derivative of the distance from B .

Lemma 7.8 (Local speed toward B). *Let $t \mapsto A_t$ be a fast folding path. Then, the speed at which A_t is approaching B is given by*

$$\frac{2\mu_t(\gamma)}{l_{A_t}(\gamma)}$$

where γ is a loop that realises the maximal stretching factor from B to A_t .

Proof: As above, since $t \mapsto A_t$ is an isometric folding path constructed as in Theorem 5.6, we are interested only in d_L . We have

$$d(A_t, B) = d_L(A_t, B) = \log \left(\frac{l_{A_t}(\gamma)}{l_B(\gamma)} \right).$$

During the folding procedure, in the marked graph A_t , the length of γ decrease twice the number of occurrences of the folding turns in γ . Whence the claim follows. \square

Now, the aim is to show that the ratio between the speed toward B and the local speed is bounded below by a given constant. Indeed, if so, one could deduce that the hypothesis of Lemma 7.2 is satisfied, this providing quasi-geodesics with uniform constants.

Lemma 7.9. *Let $A_t \neq B$ be any point in a fast folding path. Let γ be a loop that realises the maximal stretching factor from B to A_t . Then*

$$\mu_t(\gamma) \geq 1.$$

Proof: Otherwise γ would be immersed via the optimal map f used for defining the folding procedure, which would imply $l_{A_t}(\gamma) = l_B(\gamma)$, whence $A_t = B$. \square

Lemma 7.10. *In a fast folding path, for any loop γ , the quantity $\mu_t(\gamma)$, as a function of t , is monotone non-increasing.*

Proof: Let $0 = t_0 < t_1 \dots$ be the subdivision of times. Clearly, nothing change for t different from the t_i 's. We show that the multiplicity cannot increase passing trough any t_i 's. Let τ be a turn where the segments a and b are identified during the interval of time $[t_{i-1}, t_i]$. Denote such a turn τ by (a, b) . The segments a and b have one extreme in common, say the starting point. On the other hand, the ending points of a and b , say x and y respectively, must be different, otherwise the folding procedure would decrease the rank of our marked metric graphs, which is not possible.

The multiplicity, in γ , of the turns that already exist for $t \in (t_{i-1}, t_i)$ is unchanged. So we have to check what happens to the new turns created by the folding. Those are pair of segments a' and b' having x and y as starting points, and identified by the optimal map. Let $\{(a_j, b_j)\}$ be the set of turns folded for $t \in (t_{i-1}, t_i)$ whose ending points are x and y .

The multiplicity of the turn (a', b') counts how many times γ passes trough the turn. But any times that γ passes trough (a', b') must passes trough one of the (a_j, b_j) 's as well. So the total sum is not increased. \square

Definition 7.11 (ε -thin part). The ε -thin part of CV_n is the set of marked metric graphs having a loop shorter than ε in the volume-one-representative. In other words, the class a marked metric graph A in CV_n lies in the ε -thin part if

$$\frac{l_A \text{ (shortest loop of } A)}{\text{vol } A} < \varepsilon.$$

Otherwise, we say that A lies in the ε -thick part.

Lemma 7.12. *There is a constant $C > 0$ such that for any fast folding path $t \mapsto A_t$, if A_t never enters the ε -thin part, then the ratio between the speed approaching toward B and the local speed is bounded below by $C \cdot \varepsilon$.*

Proof: Since our folding procedure is isometric, if, starting from A_t , we fold during a time T , then the volume of A_t is decreased at least by T :

$$T \leq \text{vol}(A_t) - \text{vol}(B) = \text{vol}(A_t) - 1.$$

On the other hand, the length of a given loop is decreased by

$$l_{A_t}(\gamma) - l_B(\gamma) = 2 \int_t^{t+T} \mu_s(\gamma) ds \leq 2T \mu_t(\gamma)$$

where the inequality follow from Lemma 7.10.

Now, let γ be a loop realising the maximal stretching factor from B to A_t . Since $\text{vol}(\bar{B}) = 1$, the length of γ in \bar{B} is less than 2 (because of

Proposition 3.15). By the above inequalities it follows that

$$l_{A_t}(\gamma) \leq 2\mu_t(\gamma) \operatorname{vol}(A_t).$$

Let γ_1 be a simple loop minimising $l_{A_t}(w)/\mu_t(w)$. The ratio between the approaching speed toward B and the local speed is, by Lemmata 7.7 and 7.8

$$\frac{l_{A_t}(\gamma_1)\mu_t(\gamma)}{\bar{l}_{A_t}(\gamma)\mu_t(\gamma_1)}$$

which is therefore bounded below by

$$\frac{l_{A_t}(\gamma_1)}{2 \operatorname{vol} A_t \mu_t(\gamma_1)} \geq C \frac{l_{A_t}(\text{shortest loop of } A_t)}{\operatorname{vol}(A_t)}$$

where C is a constant depending only on the rank n . Actually, the constant C depends on the fact that $\mu_t(\gamma_1)$ is bounded above, depending on the rank, because γ_1 is a simple loop.

Therefore, the ratio between the approaching speed toward B and the local speed is bounded below by $C \cdot \varepsilon$ if A_t lies in the ε -thick part of CV_n . □

An immediate corollary is the following

Theorem 7.13 (Folding paths are quasi-geodesic). *For any $\varepsilon > 0$ there are constants K, L depending only on ε and the rank of CV_n such that for any two marked metric graph A and B whose corresponding fast folding path $t \mapsto A_t$ from A_0 to \bar{B} (notation as in Theorem 5.6) stay in the ε -thick part, there is right-geodesic between A and B which is a (K, L) -quasi-geodesic.*

Proof: Lemma 7.12 implies that the hypothesis of Lemma 7.2 is satisfied. By Theorem 7.3 and Proposition 7.4 the claim follows. □

This can be used to give an alternative proof of quasi-symmetry of d_R in the thick part (see [1] for the proof using quasi-symmetry of the combinatorial metric and the fact that that metric is quasi-isometric to d_R in the thick part).

8. Iterating automorphisms

Here, we study the behaviour of the orbits of automorphisms with respect to our metrics.

Theorem 8.1. *Let $\Phi \in \operatorname{Aut}(F_n)$ be an automorphism of exponential growth. Then for any $A \in CV_n$ the sequence $\Phi^h A$ is a d -quasi-geodesic as a map from $\mathbb{Z} \rightarrow CV_n$. Moreover, if A is a train-track for Φ , then it is a d_R -geodesic.*

Proof: If Φ has exponential growth so does Φ^{-1} (this is a consequence of the existence of the relative train track representatives of [4]). That means that $\sup_{1 \neq w \in F_n} l(\Phi^h(w))/l(w) > kc^h$ for some $k > 0$ and $c > 1$, where the length l is calculated in any fixed rose (and the same holds for Φ^{-1}). We have

$$\sup_{1 \neq w \in F_n} \frac{l_A(\Phi^{h+m}w)}{l_A(\Phi^m w)} = \sup_{1 \neq w \in F_n} \frac{l_A(\Phi^h w)}{l_A(w)} = \sup_{1 \neq w \in F_n} \frac{l_A(\Phi^h w)}{l(\Phi^h w)} \cdot \frac{l(\Phi^h w)}{l(w)} \cdot \frac{l(w)}{l_A(w)}.$$

In the last term of above equality, the first and the last factors are bounded below by constants because A lies at finite distance from the rose used for calculating l . The middle term is bounded below by kc^h by our hypothesis of exponential growth. Similarly, using that also Φ^{-1} has exponential growth, we can show that

$$\Lambda(\Phi^{h+m} A, \Phi^m A) > kc^h$$

for some constants $k > 0$ and $c > 1$, this giving

$$d(\Phi^{h+m} A, \Phi^m A) > \log k + h \log c.$$

The other inequality is even easier, and does not need any assumption on Φ :

$$\sup_{1 \neq w \in F_n} \frac{l_A(\Phi^{h+m}w)}{l_A(\Phi^m w)} = \sup_{1 \neq w \in F_n} \frac{l_A(\Phi^{h+m}w)}{l_A(\Phi^{h+m-1}w)} \cdot \frac{l_A(\Phi^{h+m-1}w)}{l_A(\Phi^{h+m-2}w)} \cdots \frac{l_A(\Phi^{1+m}w)}{l_A(\Phi^m w)}$$

which is bounded above by

$$\left(\sup_{1 \neq w \in F_n} \frac{l_A(\Phi w)}{l_A(w)} \right)^h$$

whence (arguing the same way for Φ^{-1})

$$\Lambda(\Phi^{h+m} A, \Phi^m A) \leq \Lambda(\Phi A, A)^h$$

and

$$d(\Phi^{h+m} A, \Phi^m A) \leq hd(\Phi A, A).$$

Suppose now that A is a train track for Φ . Then every edge is stretched exactly by λ , the Perron-Frobenius eigenvalue associate to the transition matrix for Φ (see [4]). It follows that $\Lambda_R(\Phi^{h+m}, \Phi^m) = \lambda^h$, and the second claim follows. \square

The fact that train tracks for Φ and Φ^{-1} are in general different, and that also the Perron-Frobenius eigenvalues for Φ and Φ^{-1} may differ, tells us that we cannot follow this approach for building a d -geodesic axis for Φ .

Now, Theorem 8.1 clearly would fail if the automorphism in question is of polynomial growth as $d(\Phi^h A, \Phi^k A)$ will be bounded by a constant multiple of $\log |h - k|$. However it is important to note that, nevertheless, the various folding paths from a point to the points in its orbit may still be quasi-geodesics (with the unit speed parametrisation) as in the following example.

Example 8.2. Let R be the rose of rank 2, with loops labelled A, B and let ϕ be the automorphism which sends A to A and B to BA . Then, for any k , the folding path from R to $\phi^k(R)$ is a $(4, 0)$ quasi-geodesic.

Proof: In the rose, the petals have the same length, but since our metric is scale invariant, we may choose that length—we choose it to be $k + 1$. We let R_k denote $\phi^k(R)$, which then also has two loops of the same length, which we label A_k and B_k , and give them both length 1. By definition, A maps to the loop A_k in R_k and B maps to $B_k(A_k)^k$.

In the folding path we start, first of all, by shrinking all the edges so that (after scaling, which we have already done) the map from the left to the right is isometric on edges. This means that we shrink the loop A until it has length 1. We call this new graph R_0 ; it has one vertex and two loops, $A_0 \rightarrow A_k$ and $B_0 \rightarrow B_k(A_k)^k$. The length of A_0 is 1 and the length of B_0 is $k + 1$.

The folding path then proceeds by folding A_0 into B_0 . If one imagines this as a discrete process, after the i^{th} stage we will obtain a graph R_i , with a single vertex and two loops, $A_i \rightarrow A_k$ and $B_i \rightarrow B_k(A_k)^{(k-i)}$; the length of A_i is 1 and the length of B_i is $K + 1 - i$.

If we then fold a part of A_i , of length δ , into B_i we travel to a point in the folding path which we shall call $R_{i,\delta}$. This has two vertices, \bullet and \circ , and three edges, $A_{i,\delta}$, $B_{i,\delta}$ and $C_{i,\delta}$.

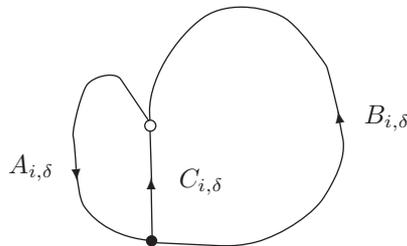


FIGURE 5. The graph $R_{i,\delta}$.

Here, we can map the vertex \bullet to the unique vertex of R_k and then map the loop $A_{i,\delta}C_{i,\delta}$ to A and $B_{i,\delta}C_{i,\delta}$ to $B_k(A_k)^{(k-i)}$ (this is enough to specify the marking up to homotopy equivalence); the length of $A_{i,\delta}$ is $1 - \delta$, the length of $B_{i,\delta}$ is $k + 1 - i - \delta$ and the length of $C_{i,\delta}$ is δ . This marked metric graph represents an arbitrary point on the folding path from R_0 to R_k . Now, following Lemma 7.7, the local speed is realised by the loop $B_{i,\delta}\overline{A_{i,\delta}}$, whereas the distance to R_k is realised by the loop B_k which is realised by $B_{i,\delta}\overline{A_{i,\delta}}(\overline{C_{i,\delta}}\overline{A_{i,\delta}})^{k-i-1}$ in $R_{i,\delta}$. Both of these loops pass through the unique folding turn (Definition 5.7) of $R_{i,\delta}$ exactly once.

Hence, by Lemmas 7.7 and 7.8, the ratio of the speed toward R_k and the local speed is,

$$\frac{k + 2 - i - 2\delta}{2k + 1 - 2i - 2\delta} \geq \frac{1}{2}.$$

Thus, by Lemma 7.2, the path from R_0 to R_k is a $(2, 0)$ quasi-geodesic and thus by Proposition 7.4, the whole path is a $(4, 0)$ quasi-geodesic. \square

9. Some open questions

In this section we address some questions which arose during the many conversations we had with colleagues, principally during the coffee breaks of conferences, about the metric properties of Outer Space.

9.1. Existence of quasi-geodesics. As we have seen, folding paths that do not fold into the thin part provide quasi-geodesics for the symmetric metric. Here we address mainly two questions. First, whether a folding path will always produce a quasi-geodesic or not, with constants depending only on the rank. Second, whether it is in general possible to connect any two marked metric graphs with a path which is a quasi-geodesic, with constants depending only on the rank of the graphs.

For the latter question, there is an heuristic argument: suppose the answer is no. Then, letting blowing up the constants, one would get a counter-example-sequence that contradicts Lemma 7.2. Then, following the arguments of Theorem 7.13 one gets that the folding paths of the counter-example-sequence will eventually enter any ε -thin part, but explicit computations show that a folding path that enters the thin part cannot stay for too long inside that part (one has perhaps to understand how many times a folding path can enter the thin part). Thus suggesting an affirmative answer to our questions.

9.2. Existence of a geodesic axis for an iwip. We have seen that iterates of automorphisms produce quasi-geodesics (and geodesics for the

non-symmetric metric). The natural question here is whether automorphisms have an axis and whether can such an axis be described in terms of metric properties. Also, one can ask whether one can compute the “geometric rank” of such axis. Is there any analogue of the bounded projection Lemma? (see [5] and the recent preprint [1]).

9.3. Hyperbolicity, flats and coarse properties. It is natural to ask whether some subset of Outer Space (some thick-part?) is hyperbolic or presents hyperbolicity phenomena. On the other hand, it would be interesting to study the (quasi-) flats of Outer Space, if any. In general coarse properties of Outer Space are still unknown (for instance, what do its asymptotic cones look like?).

References

- [1] Y. ALGOM-KFIR, Strongly contracting geodesics in Outer Space (2008), arXiv:0812.1555.
- [2] M. BESTVINA AND M. FEIGN, Outer limits, Preprint (1994).
- [3] M. BESTVINA, M. FEIGN, AND M. HANDEL, Laminations, trees, and irreducible automorphisms of free groups, *Geom. Funct. Anal.* **7(2)** (1997), 215–244.
- [4] M. BESTVINA AND M. HANDEL, Train tracks and automorphisms of free groups, *Ann. of Math. (2)* **135(1)** (1992), 1–51.
- [5] M. BESTVINA AND K. FUJIWARA, A characterization of higher rank symmetric spaces via bounded cohomology, *Geom. Funct. Anal.* **19(1)** (2009), 11–40.
- [6] M. M. COHEN AND M. LUSTIG, Very small group actions on \mathbf{R} -trees and Dehn twist automorphisms, *Topology* **34(3)** (1995), 575–617.
- [7] D. COOPER, Automorphisms of free groups have finitely generated fixed point sets, *J. Algebra* **111(2)** (1987), 453–456.
- [8] M. CULLER AND J. W. MORGAN, Group actions on \mathbf{R} -trees, *Proc. London Math. Soc. (3)* **55(3)** (1987), 571–604.
- [9] S. FRANCAVIGLIA, Geodesic currents and length compactness for automorphisms of free groups, *Trans. Amer. Math. Soc.* **361(1)** (2009), 161–176.
- [10] V. GUIRADEL AND G. LEVITT, Deformation spaces of trees, *Groups Geom. Dyn.* **1(2)** (2007), 135–181.
- [11] I. KAPOVICH, The frequency space of a free group, *Internat. J. Algebra Comput.* **15(5–6)** (2005), 939–969.

- [12] I. KAPOVICH, Currents on free groups, in: *“Topological and asymptotic aspects of group theory”*, Contemp. Math. **394**, Amer. Math. Soc., Providence, RI, 2006, pp. 149–176.
- [13] K. VOGTMANN, Automorphisms of free groups and Outer Space, in: *“Proceedings of the Conference on Geometric and Combinatorial Group Theory”*, Part I (Haifa, 2000), *Geom. Dedicata* **94** (2002), 1–31.

Stefano Francaviglia:
Dipartimento di Matematica
Università Bologna
P.zza Porta S. Donato 5
40126 Bologna
Italy
E-mail address: `francavi@dm.unibo.it`

Armando Martino:
School of Mathematics
University of Southampton
SO17 1BJ Southampton
UK
E-mail address: `a.martino@soton.ac.uk`

Primera versió rebuda el 17 de setembre de 2010,
darrera versió rebuda el 27 de gener de 2011.