

Length Functions and Outer Space

JOHN SMILLIE* & KAREN VOGTMANN**

1. Introduction

Let F_n be the free group of rank n . The outer automorphism group of F_n acts on a certain contractible space called *Outer Space* which was introduced in [CV1]. The role played by Outer Space in the study of the group of outer automorphisms of F_n is analogous to the role played by Teichmüller space in the study of the mapping class group of a surface. In [T] Thurston constructs an embedding of Teichmüller space into a *finite-dimensional* projective space by means of length functions. In this paper we show that a similar construction for Outer Space is not possible.

We say a graph has genus n if its fundamental group is isomorphic to F_n . An \mathbf{R} -graph is a graph which is a metric space where each edge is isometric to an interval of \mathbf{R} . Outer Space can be thought of as a normalized collection of marked \mathbf{R} -graphs of genus n with no free edges, where a *marking* on a graph is an identification of the fundamental group of the graph with F_n , and where two graphs are equivalent if there is an isometry between them which preserves the marking. There are two methods of implementing the normalization. Define the total length of a graph to be the sum of the lengths of the edges. The first method of normalization is to consider only graphs of total length 1. An alternative method of normalization is to consider equivalence classes of graphs where two graphs are equivalent if there is a homeomorphism preserving the marking taking one to the other which multiplies lengths by a constant.

The term *length function* is used in two different ways in combinatorial group theory. There are the length functions introduced by Lyndon (in [CM] these are called “based length functions”) and there are the “hyperbolic length functions” of [AB] which are called “translation length functions” in [CM]. In our context an element of Outer Space determines a hyperbolic length function on F_n , where the length of $g \in F_n$ is the length of the shortest (un-based) loop representing the homotopy class corresponding to g . (The definition of the Lyndon length function involves choosing a base point and

Received March 1, 1991.

*Partially supported by NSF grant #DMS-9003101.

**Partially supported by NSF grant #DMS-8702070.

Michigan Math. J. 39 (1992).

considering the length of based loops.) Note that the hyperbolic length function is constant on conjugacy classes so it can be thought of as a function on the set \mathcal{C} of conjugacy classes in F_n . It is the hyperbolic length function which is the analogue of the length function of Thurston. In what follows we will concern ourselves only with hyperbolic length functions, which we will refer to simply as “length functions”.

Instead of thinking of a length function as a function on \mathcal{C} determined by a marked graph, we can think of a length function as a function on the set of marked graphs determined by an element of \mathcal{C} . In this way length functions give natural coordinates for the space of marked graphs. The set of all length functions gives an embedding of the set of marked graphs into $\mathbf{R}^{\mathcal{C}}$ (see [AB] or [CM]). If we consider Outer Space to be a set of marked graphs with total length 1 then we get an embedding of Outer Space into $\mathbf{R}^{\mathcal{C}}$. If we use the second method of normalization we get an embedding of Outer Space into the projective space $P(\mathbf{R}^{\mathcal{C}})$.

The above construction mimics the construction of Thurston for Teichmüller space. In that situation, one can choose a *finite* set of conjugacy classes whose length functions give an embedding of Teichmüller space. In this paper we are interested in the corresponding question for Outer Space and for a deformation retract, which we now define.

In rank 2, Outer Space is homeomorphic to a disk with “fins” attached (see [CV2]). The disk corresponds to the space of graphs with no disconnecting edges. In this case we might suppose that Teichmüller space is more closely related to the disk than to the disk with fins. Let us call the subset of Outer Space consisting of graphs with no disconnecting edges *Reduced Outer Space*. In any rank there is an equivariant deformation retract from Outer Space to Reduced Outer Space. In particular, Reduced Outer Space is contractible.

It is shown in [CLS] that no finite set of length functions can be used to embed the (unnormalized) space of marked graphs. That paper leaves open two questions related to embedding Outer Space. First, the construction of [CLS] does not show that the restriction to marked graphs with total length 1 is not an embedding. Second, the construction of [CLS] does not answer the embedding question for Reduced Outer Space. In Section 2 of this paper we give a different construction that answers these questions in rank 3. In Section 3 we show that embedding fails in any rank. Each of the following theorems demonstrates a different aspect of the failure of the embedding. Each is based on a modification of the construction in Section 2.

Recall that a *rose* is a marked graph (g, G) , where G is a graph with one vertex and n edges. The first result shows that there exist families of marked graphs which agree on a given finite set of words and which have arbitrarily large diameter, in the sense that they include an arbitrarily large number of different roses.

THEOREM 1. *Let F_n be the free group of rank $n > 2$. Let Σ be a finite subset of \mathcal{C} . Then there is a 1-parameter family S of marked graphs, so that the length of each element of Σ is the same in every graph in S , and so that S contains an arbitrarily large number of roses. Furthermore, each element of S has total length 1 and has no separating edges.*

The next result follows from the observation that the homeomorphism type of the graph plays a relatively minor role in the construction of the previous section.

THEOREM 2. *Let F_n be the free group of rank $n > 2$. Let Σ be a finite subset of \mathcal{C} , and let (g, G) be any marked graph of genus n such that G has no separating edges. Then there is a different marked graph (g', G') , with G isometric to G' , such that the length of each element of Σ is the same in (g, G) and in (g', G') .*

A further modification shows that one can find an open set of marked graphs which agree on a given finite set of words and which has fairly large dimension.

THEOREM 3. *Let F_n be the free group of rank $n > 2$. Let Σ be a finite subset of \mathcal{C} . Then there is a $(2n - 5)$ -parameter family S of marked graphs so that the length of each element of Σ is the same in every graph in S . Furthermore, each element of S has total length 1 and has no separating edges.*

REMARK. Note that these theorems are all false in rank 2. In this case it is possible to find a set Σ consisting of five words so that the lengths of words in Σ embed Reduced Outer Space in a 2-dimensional linear subspace of \mathbf{R}^5 [CV2].

If we use the second method of normalization then the set of length functions gives an embedding in an infinite-dimensional projective space. By [CM], the closure of the image of this embedding is compact (with respect to the finite-open topology) and the frontier is called the boundary of Outer Space. When this construction is made for surface groups the resulting boundary is Thurston's compactification of Teichmüller space, and Thurston shows that it is finite-dimensional (see [H]). In fact, the lengths of a finite set of conjugacy classes embed the boundary in a finite-dimensional projective space. The following result shows that the analogous statement is not true for free groups.

THEOREM 4. *Let $\Sigma \subset F_n$ be any finite set of words. Then there is a $(2n - 5)$ -parameter family S_∞ of points contained in the boundary of Outer Space so that the corresponding embedding into the projective space $P(\mathbf{R}^\Sigma)$ maps S_∞ to a point.*

2. Construction in Rank 3

In this section we describe the basic construction in detail for rank 3. The first step in the construction involves choosing a basis for F_3 so that all the elements of Σ will have a special form when written as words in that basis.

We start with an arbitrary basis $\{a, b, c\}$, and write the elements of Σ as words in a, b , and c . Since conjugate elements of F_n have the same length in any marked graph, we may assume that the words in Σ are cyclically reduced. Let M be the maximum over all $w \in \Sigma$ of $|k|$, where c^k occurs in w . Define an automorphism ϕ as follows: $\phi a = c^m a c^m$, $\phi b = b$, and $\phi c = c$, where m is chosen to be greater than M .

CLAIM. Each reduced word in $\phi\Sigma$ has the property that a appears only with exponent ± 1 , and whenever a does appear, it is both immediately preceded and immediately followed (in the cyclic sense) by a nonzero power of c .

Proof. Let w be a word in Σ . If a does not appear in w then a does not appear in ϕw .

Now assume that a appears in w . By cyclically reordering the letters we may assume that w begins with a . Write w so that a only occurs with exponent ± 1 , so that w is a concatenation of segments of the form $a^{\pm 1}u$, where u is a (possibly empty) reduced word in b 's and c 's. Then ϕw is cyclically equivalent to a concatenation of segments of the form $a^{\pm 1}c^{\pm m}uc^{\pm m}$ or $a^{\pm 1}c^{\pm m}uc^{\mp m}$. We want to show that some power of c remains to the right and (cyclic) left of a when ϕw is reduced. If u is empty then successive occurrences of a have the same exponent, so the reduced segment is $a^{\pm 1}c^{\pm 2m}$. If $u = c^r$ with $r \neq 0$, then $a^{\pm 1}c^{\pm m}uc^{\pm m}$ reduces to $a^{\pm 1}c^{\pm 2m+r}$ and $a^{\pm 1}c^{\pm m}uc^{\mp m}$ reduces to $a^{\pm 1}c^r$. If $u = c^s v c^t$, where v begins and ends with a power of b and s, t can be 0, then $a^{\pm 1}c^{\pm m}uc^{\pm m}$ reduces to $a^{\pm 1}c^{\pm m+s}vc^{\pm m+t}$ and $a^{\pm 1}c^{\pm m}uc^{\mp m}$ reduces to $a^{\pm 1}c^{\pm m+s}vc^{\mp m+t}$. Our assumptions on m and r imply that the exponents $\pm 2m$, $\pm 2m+r$, r , $\pm m+s$, and $\pm m+t$ of c are nonzero, proving our claim. \square

Now, by replacing the basis $\{a, b, c\}$ by the basis $\{\phi^{-1}a, \phi^{-1}b, \phi^{-1}c\}$ and relabelling this basis as $\{a, b, c\}$, we may assume that the words in Σ have the property that any occurrence of $a^{\pm 1}$ is preceded and followed (in the cyclic sense) by a nonzero power of c .

We now apply the automorphism γ to Σ , where γ is defined by $\gamma a = b^k a b^{-k}$, with $k > 0$, $\gamma b = b$, and $\gamma c = c$. Each segment $c^r a^{\pm 1} c^s$ of a word $w \in \Sigma$ is replaced by $c^r b^k a^{\pm 1} b^{-k} c^s$. This introduces no opportunities for cancellation into the word w .

The words in $\gamma\Sigma$ now have the following property:

- (*) Any occurrence of $a^{\pm 1}$ is immediately preceded (in the cyclic sense) by b and followed by b^{-1} ; furthermore, any two occurrences of a or a^{-1} are separated by a nonzero power of c .

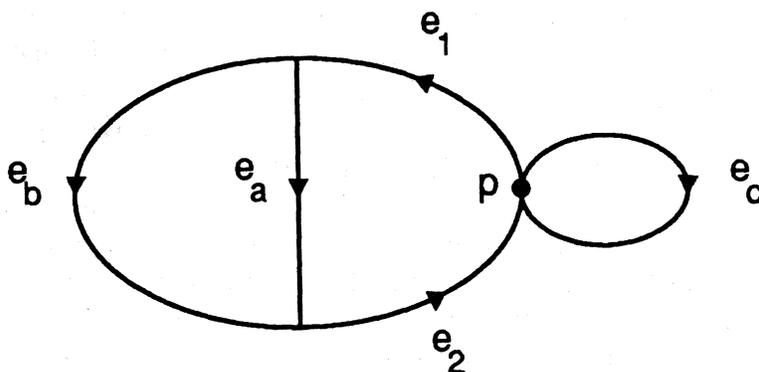


Figure 1

As before, replace the basis by its image under γ^{-1} and relabel, so we may assume that the words in Σ have property (*) with respect to the basis a , b , and c .

Let G be the graph pictured in Figure 1. We choose a basis for its fundamental group $\pi_1(G, p)$ as follows:

$$a = e_1 e_a e_2, \quad b = e_1 e_b e_2, \quad c = e_c.$$

The words in Σ are represented by reduced edge paths which are loops in G . Observe that the reduced edge path representing any word $w \in \Sigma$ passes through p . This is clear if w involves the letter c , and the only words that can be in Σ which do not involve c are powers of b , which also pass through p . Thus we can compute the lengths of w by dividing the reduced edge path representing w into loops which begin and end at p .

CLAIM. A reduced edge path representing a word in Σ does not make the turns $e_a e_2$, $\bar{e}_2 \bar{e}_a$, $\bar{e}_a e_b$, and $\bar{e}_b e_a$. In other words, if we express the word as a reduced sequence of edges then the sequences $e_a e_2$, $\bar{e}_2 \bar{e}_a$, $\bar{e}_a e_b$, and $\bar{e}_b e_a$ do not occur.

Proof. Any occurrence of a or a^{-1} in a word $w \in \Sigma$ is preceded (cyclically) by $c^r b^k$ and followed by $b^{-k} c^s$, with $r, s \neq 0$. The sequence $c^r b^k a b^{-k} c^s$ is represented by the loop

$$e_c^r (e_1 e_b e_2)^k e_1 e_a e_2 (\bar{e}_2 \bar{e}_b \bar{e}_1)^k e_c^s,$$

which reduces to

$$e_c^r (e_1 e_b e_2)^k e_1 e_a \bar{e}_b \bar{e}_1 (\bar{e}_2 \bar{e}_b \bar{e}_1)^{k-1} e_c^s.$$

No further cancellations are possible. The sequence $c^r b^k a^{-1} b^{-k} c^s$ reduces to

$$e_c^r (e_1 e_b e_2)^k e_1 e_b \bar{e}_a \bar{e}_1 (\bar{e}_2 \bar{e}_b \bar{e}_1)^{k-1} e_c^s.$$

No other sequences in the word w involve the edge e_a at all, so the turns $e_a e_2$, $\bar{e}_2 \bar{e}_a$, $\bar{e}_a e_b$, and $\bar{e}_b e_a$ do not occur. \square

We now change the lengths on the edges of G as follows. We increase the lengths of the edges e_1 and e_2 by a small constant ϵ and decrease the length of the edge e_b by 2ϵ . Call the resulting graph G_ϵ . By the observation above, we may calculate the edge-path length of a path representing a word $w \in \Sigma$ by dividing it into subpaths which begin and end at p . If w doesn't involve the letter a , its length in G_ϵ is clearly the same as its length in G , since the circle $e_1e_be_2$ has not changed length. Now suppose that a occurs in w . By the above claim, when the reduced edge path representing w traverses the edge e_a , it must be preceded by the edge e_1 and followed by the edges $\bar{e}_b\bar{e}_1$; thus the total length of this segment of the edge path is unchanged. Similarly, if the reduced edge path representing w traverses \bar{e}_a , it must be preceded by the edges e_1e_b and followed by \bar{e}_1 , and the total length is again unchanged.

Varying ϵ gives us a 1-parameter family of marked graphs. We have just shown that the words in Σ have the same length in all of the marked graphs represented by this family. On the other hand, these marked graphs are not the same; for instance, the word a has different length in each marked graph.

3. Proof of the Theorems

We begin with the proof of Theorem 1, which states that Outer Space contains 1-parameter families of marked graphs which all agree on a given finite set Σ of cyclic words in F_n , and which contain arbitrarily many roses.

Proof of Theorem 1 for $n=3$. We refer to the basic construction given in the previous section and to Figure 1. As we vary the parameter ϵ in the basic construction, the endpoints of the edge e_a move around the circle $C_b = e_1e_be_2$ in opposite directions. It is useful to picture this in the cover of G which is obtained by unwrapping C_b ; in this cover b has a translation axis A_b . The condition that a path representing a word in Σ never makes the turns e_ae_2 , $\bar{e}_2\bar{e}_a$, \bar{e}_ae_b , and \bar{e}_be_a means that a lift of such a path to the cover never turns from e_a or \bar{e}_a onto A_b in the translation direction of b . The fact that any occurrence of a in a word of Σ is preceded by b^k and followed by b^{-k} means that a lifted path representing that word travels k units along A_b , in the translation direction of b , before it crosses a lift of e_a , and then travels at least k units in the opposite direction along A_b . Therefore, if we move the endpoints of e_a less than k units, the lifted path still has the property that it never turns from e_a or \bar{e}_a onto A_b in the direction of b , and the length of a path representing a word in Σ does not change.

In the graph G in Figure 1, suppose that $\text{length}(e_1) = \text{length}(e_2)$. If we move the endpoints of e_a equal amounts in opposite directions, they will at times coincide at the point p ; that is, the graph will be a rose at those times. By the above remarks, the lengths of words in Σ will not change as long as the total distance each endpoint is moved is less than k . By choosing k sufficiently large, we can thus get a path of length functions which agree on Σ , and which pass through an arbitrarily large number of roses. \square

We now pause to prove Theorem 2. The proof of Theorem 1 for rank greater than 3 will follow this proof.

Proof of Theorem 2. In order to prove this theorem, we note that the homeomorphism type of the graph plays a relatively minor role in the basic construction. If we fix $n \geq 3$, $\Sigma \subset F_n$, and any graph G of genus n with no separating edges, then we can find two elements of Outer Space homeomorphic to G which have the same lengths for all words in Σ . We proceed as before. We find: an embedded loop (playing the role of b), a single edge e_a with both endpoints on that loop, and another loop in the graph which does not involve e_a (playing the role of c). The following proposition shows that this is always possible.

Recall that a graph G is *minimal* if it is connected and has no free edges or bivalent vertices.

PROPOSITION. *Let G be a minimal graph of genus greater than or equal to 2 with no separating edges. Then there is an embedded loop σ in G with the property that some component of $G - \sigma$ is a single edge with both endpoints on σ .*

Proof. Note that G has an embedded loop, since G is not a tree. If τ is an embedded loop of G , we denote by $c(\tau)$ the smallest number of edges in any single connected component of $G - \tau$. Let $c(G)$ denote the minimum value of $c(\tau)$ over all embedded loops τ in G . Since G has no free edges, the conclusion of the theorem is equivalent to the statement that $c(G) = 1$.

Let σ be an embedded loop with $c(\sigma) = c(G)$, and let C be a connected component of $G - \sigma$ with $c(G)$ edges. Suppose that the closure \bar{C} of C in G intersects σ in at least two points. Choose two such points p and q which are adjacent on σ ; that is, there is a subarc α of σ joining p to q which doesn't contain any other points of $\bar{C} \cap \sigma$. Choose an embedded path β in \bar{C} joining p to q . By replacing the subarc α of σ by β , we obtain a new embedded loop σ' ; the component C of $G - \sigma$ breaks up into a union of components of $G - \sigma'$, and none of these components contains any of the edges of β . If C consists of more than one edge (i.e. if $c(G) > 1$) then one of these components is nonempty and contains fewer edges than C , contradicting our minimality assumptions.

Now suppose that C intersects σ in a single point v . Then v is a separating vertex. If the closure \bar{C} is a graph of genus 2 or more, we can use induction on the genus to find an embedded loop and single edge with endpoints on that edge in $\bar{C} \subset G$, contradicting the minimality of C . Thus \bar{C} has genus 1. Since G has no separating edges, \bar{C} must be a single loop, that is, $c(G) = 1$. \square

To finish the proof of Theorem 2, we need to apply the observations made in the proof of Theorem 1; if k is large enough, we can move the endpoints of e_a all the way around the embedded loop representing b until they return to their original positions. The end result is a graph G' isometric to G ; we have changed the marking, and therefore the length function, but not the lengths of elements in Σ . \square

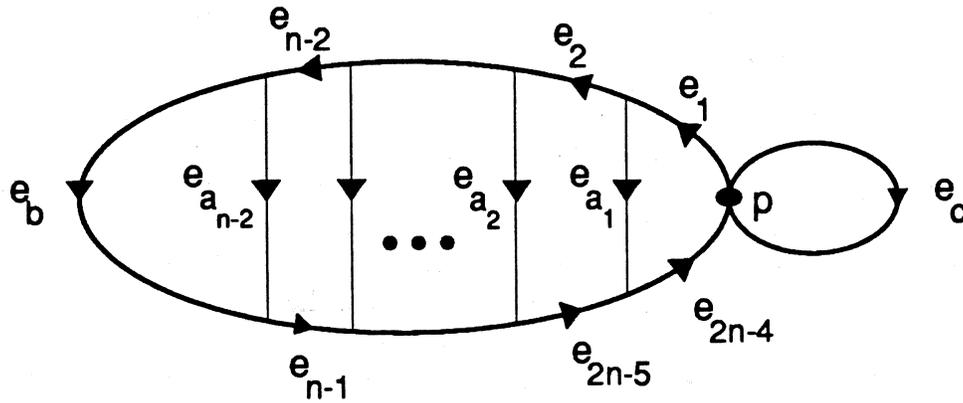


Figure 2

To complete the proof of Theorem 1 for $n > 3$, and to prove Theorems 3 and 4, we refer to Figure 2.

Proof of Theorem 1 for $n > 3$. If $n > 3$, consider the graph shown in Figure 2. Replace the automorphism ϕ in the basic construction by an automorphism sending a_i to $c^{m_i}a_i c^{m_i}$ for some $m_1 \gg m_2 \gg \dots \gg m_{n-2} \gg 0$, and replace γ by an automorphism sending a_i to $b^k a_i b^{-k}$ for some $k > 0$. By varying the lengths of the edges e_i , we obtain an $(n - 2)$ -parameter family of length functions which all agree on Σ . The proof of Theorem 1 in the case $n = 3$ generalizes to show that this family can be extended to include arbitrarily many roses by choosing k large enough. \square

Recall that Theorem 3 asserts the existence of a large-dimensional family of normalized marked graphs so that the length of each element of Σ is the same in every graph in the family.

Proof of Theorem 3. In order to obtain a $(2n - 5)$ -dimensional family, perform only the first automorphism, which sends a_i to $c^{m_i}a_i c^{m_i}$ for some $m_1 \gg m_2 \gg \dots \gg m_{n-2} \gg 0$. The reduced paths representing the words in Σ have the property that they never make the turns $\bar{e}_{i+1}e_{a_i}$ or $e_{a_i}e_{2n-3-i}$ or their inverses. For $i = 1, \dots, n - 2$, we may now identify a small initial segment of e_{a_i} with an initial segment of e_{i+1} , and a small terminal segment of e_{a_i} with a terminal segment of e_{2n-4-i} without changing the length of any word in Σ . Note that the total length of the graph is changed by each identification (in contrast with the previous construction). If we want to preserve the total length of the graph, the amount of the last identification is determined; thus we have a $(2n - 4) - 1 = (2n - 5)$ -parameter family of length functions which agree on Σ . \square

We now turn our attention to the boundary of Outer Space. The universal cover of an \mathbf{R} -graph is a (simplicial) \mathbf{R} -tree. A marked graph gives an

action of F_n on the universal cover of the graph. In fact, the space of marked graphs can be identified with the space of free actions of F_n on discrete \mathbf{R} -trees. The boundary of Outer Space consists of actions of F_n on \mathbf{R} -trees which are limits of free actions on discrete \mathbf{R} -trees in the appropriate sense. (See [CM] for more information on \mathbf{R} -trees and on the boundary of Outer Space.)

Proof of Theorem 4. To construct a family S_∞ of length functions in the boundary of Outer Space which all agree on Σ , we need only shrink the edge e_c in the graph in Figure 2 to a point. This produces a family of actions of F_n on \mathbf{R} -trees which is not free: p has the group generated by c as stabilizer. This completes the proof. \square

References

- [AB] R. Alperin and H. Bass, *Length functions of group actions on Λ -trees*, Combinatorial Group Theory and Topology (S. M. Gersten et al., eds.), pp. 265–378, Princeton Univ. Press, Princeton, NJ, 1987.
- [CLS] M. Cohen, M. Lustig, and R. Steiner, *\mathbf{R} -tree actions are not determined by the translation lengths of finitely many elements*, Arboreal Group Theory (R. Alperin, ed.), MSRI Publications, 19, pp. 183–187, Springer, New York, 1991.
- [CM] M. Culler and J. W. Morgan, *Group actions on \mathbf{R} -trees*, Proc. London Math. Soc. 55 (1987), 571–604.
- [CV1] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986), 91–119.
- [CV2] ———, *The boundary of Outer Space in rank 2*, Arboreal Group Theory (R. Alperin, ed.), MSRI Publications, 19, pp. 189–230, Springer, New York, 1991.
- [H] A. Hatcher, *Measured lamination spaces for surfaces, from the topological viewpoint*, Topology Appl. 30 (1988), 63–88.
- [T] W. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19 (1988), 417–431.

Department of Mathematics
Cornell University
Ithaca, NY 14853-7901