Roger C. Alperin

Editor

Arboreal Group Theory

Proceedings of a Workshop Held
September 13-16, 1988

With 62 Illustrations

Roger C. Alperin
Department of Mathematics
 and Computer Science
San Jose State University
San Jose, California 95192
USA

Mathematical Sciences Research Institute
1000 Centennial Drive
Berkeley, CA 94720
USA

Mathematical Subject Classifications: 20E06, 20E99, 57M05, 57M99

Printed on acid-free paper.

The Boundary of Outer Space in Rank Two

MARC CULLER AND KAREN VOGTMANN

§1. Introduction

In [4] a space X_n was introduced on which the group $Out(F_n)$ of outer automorphisms of a free group of rank n acts virtually freely. Since then, this space has come to be known as "outer space." Outer space can be defined as a space of free actions of F_n on simplicial $\mathbb{R}$-trees; we require that all actions be minimal, and we identify two actions if they differ only by scaling the metric on the $\mathbb{R}$-tree. To describe the topology on outer space, we associate to each action $\alpha\colon F_n \times T \to T$ a length function $|\cdot|_\alpha\colon F_n \to \mathbb{R}$ defined by

$$|g|_\alpha = \inf_{x \in T} d(x, gx)$$

where d is the distance in the tree T. We have $|g|_\alpha = |h^{-1}gh|_\alpha$ and $|\cdot|_\alpha \equiv 0$ if and only if some point of T is fixed by all of F_n. Thus an action with no fixed point determines a point in $\mathbb{R}^{\mathcal{C}} - \{0\}$, where $\mathcal{C}$ is the set of conjugacy classes in F_n. Since actions differing by a scalar multiple define the same point of outer space, we have a map from X_n to the infinite dimensional projective space $\mathbf{P}^{\mathcal{C}} = \mathbb{R}^{\mathcal{C}} - \{0\}/\mathbb{R}^*$. It can be shown that this map is injective (see [3] or [1]). We topologize X_n as a subspace of $\mathbf{P}^{\mathcal{C}}$.

Part of the motivation for the definition of outer space was the idea of developing an analogy between the action of $Out(F_n)$ on outer space and the action of the mapping class group of a surface on the Teichmüller space of that surface. In particular, the action of the mapping class group on Teichmüller space was exploited by Thurston in his classification of automorphisms of surfaces. Thurston gives an embedding of Teichmüller space into an infinite dimensional projective space and shows that its closure in this projective space is a finite-dimensional ball; he then uses the fact that the ball has the fixed point property to analyze the action of a single automorphism on the closure of Teichmüller space.

We would like to know how much of Thurston's theory can be adapted to automorphisms of free groups. It is shown in [4] that outer space is contractible of dimension $3n - 4$, and in [3] that its closure $\overline{X}_n$ in $\mathbf{P}^{\mathcal{C}}$ is

compact. M. Steiner and R. Skora have recently announced proofs that $\overline{X}_n$ is contractible. It is not known whether the closure of outer space is finite dimensional or whether it is an ANR; if it is an ANR, this together with contractibility of $\overline{X}_n$ would imply that $\overline{X}_n$ has the fixed point property.

In this paper we restrict ourselves to the case $n = 2$. We give an explicit description of the closure $\overline{X} = \overline{X}_2$ of outer space in rank 2. In particular, we show that $\overline{X}$ is contractible, and give an imbedding of $\overline{X}$ as a two-dimensional subset of $\mathbb{R}^3$ which makes it clear that $\overline{X}$ is an ANR.

The paper is organized as follows. In section 2 we recall some basic definitions and properties of actions on $\mathbb{R}$-trees. In sections 3-5 we determine what lies on the "boundary" $\partial\overline{X} = \overline{X} - X$. As a starting point, we have by [3] that points in $\partial\overline{X}$ correspond to non-trivial actions of F_2 on $\mathbb{R}$-trees with cyclic arc stabilizers. In addition, we note in section 2 that the stabilizer of an arc in any limit of free actions is either trivial or is a maximal cyclic subgroup of the stabilizer of each of its endpoints. We remark that Cohen and Lustig [2] have given criteria for deciding when an action of F_n on a simplicial tree is a limit of free actions. For $n = 2$ these criteria consist of the above conditions on arc stabilizers. (For $n > 2$ an extra condition must be imposed.)

In section 6 we describe the imbedding of $\overline{X}$ into $\mathbb{R}^3$. One consequence of our analysis in sections 3-5 is that all actions in $\overline{X}$ are geometric, in the sense that they are isomorphic to actions of the fundamental group of a punctured torus or twice-punctured disk on trees which are dual to projective measured laminations on those surfaces. Our embedding demonstrates how the projective lamination spaces for various surfaces fit together to form outer space. A description of these projective lamination spaces is given in Hatcher [6].

We would like to thank Peter Shalen for helpful conversations and for inventing the name "outer space," Richard Skora for explaining the proof of his realization theorem, and Curt McMullen for supplying the computer programs used to generate the portrait of outer space.

§2. Group actions on trees.

In this section we establish some notation and conventions that we will use throughout the paper. For the general theory of group actions on $\mathbb{R}$-trees, we refer to [3] or [1].

By a *tree* we will always mean an $\mathbb{R}$-tree, with metric d. A *germ* at a point p of a tree is an equivalence class of rays from p, where two rays are equivalent if they agree in some neighborhood of p. An $\mathbb{R}$-tree is *simplicial* if it is homeomorphic to a simplicial tree, i.e. to a connected, 1-connected, 1-dimensional simplicial complex.

Our actions on trees will always be left actions. Two actions $\alpha_1 \colon G \times T_1 \to T_1$ and $\alpha_2 \colon G \times T_2 \to T_2$ are *isomorphic* if there is an equivariant isometry from T_1 to T_2. An action $\alpha \colon G \times T \to T$ is *simplicial* if T is a simplicial tree. The action is *minimal* if T has no proper invariant sub-tree. The *translation length* of an element $g \in G$ is given by

$$|g|_\alpha = \inf_{p \in T} d(p, gp).$$

We will often omit the subscript α, when no confusion will result. If $|g|_\alpha = 0$, then g fixes a subtree $\mathrm{Fix}(g)$ of T; in this case g is called *elliptic*. If g has no fixed point, g is called *hyperbolic* and has a translation axis $\mathrm{Axis}(g)$. By the *characteristic set* of g we will mean $\mathrm{Fix}(g)$ if g is elliptic and $\mathrm{Axis}(g)$ if g is hyperbolic.

Definition. An action of a group on an $\mathbb{R}$-tree will be said to have *maximal cyclic arc stabilizers* if any non-trivial stabilizer of an arc is a maximal cyclic subgroup of the stabilizers of each endpoint of the arc.

Lemma 2.1. *Let $\alpha \colon F_2 \times T \to T$ be an action which is a limit of free simplicial actions. Then α has maximal cyclic arc stabilizers.*

Proof. Let e be an arc of T with endpoints v and w. As noted in the introduction, [3] show that any limit of free actions has cyclic arc stabilizers, so the stabilizer of e is cyclic. If it is not a maximal cyclic subgroup of the stabilizer of v, there is an element g in the stabilizer of v such that $ge \neq e$ but $g^k e = e$ for some $k > 0$. Let $\{\alpha_n \colon F_2 \times T_n \to T_n\}$ be a sequence of free simplicial actions such that the associated (non-projective) length functions $|\cdot|_{\alpha_n}$ converge to the length function $|\cdot|_\alpha$.

By [11], the topology on X_2 as a subset of $\mathbf{P}^{\mathcal{C}}$ is the same as the Gromov topology. In particular, we can find points $w_n \in T_n$ such that $d(w_n, gw_n) \to$

$d(w, gw)$ and $d(w_n, g^k w_n) \to d(w, g^k w)$. Since g has an axis in each of th trees T_n, we have

$$d(w_n, g^k w_n) = d(w_n, gw_n) + (k-1)|g|_{\alpha_n}.$$

This gives a contradiction, since the left-hand side converges to $d(w, g^k w) = 0$, while the right-hand side converges to $d(w, gw) > 0$.   [

§3. Simplicial actions in the boundary.

In this section we describe all actions in $\overline{X}$ with the property that some primitive element of F_2 has a fixed point. When we complete our analysis of all actions in $\overline{X}$, we will see that these are exactly the simplicial actions in $\partial \overline{X} = \overline{X} - X$.

Definition. Let G be a group acting on an $\mathbb{R}$-tree T. A subtree H of T is a *fundamental domain* for the action if $GH = T$ and if for any $g \in G$, $gH \cap H$ is either empty, a single point, or equal to $H \cap \text{Fix}(g)$.

The following lemma gives a criterion for an action of F_2 on a tree to be simplicial, and produces a fundamental domain for the action. The lemma is easily generalized to free groups of higher rank; we give the statement for rank 2 to avoid complicated notation.

Lemma 3.1. *Let $\{a, b\}$ be a basis for F_2, let $F_2 \times T \to T$ be a minimal action, and let H be a non-empty closed subtree of T. Suppose that $T - H$ is the disjoint union of open sets $S_a, S_{a^{-1}}, S_b$, and $S_{b^{-1}}$ with finite boundary. Assume, for $x, y \in \{a, a^{-1}, b, b^{-1}\}$, that*

1. $xH \cap H \neq \emptyset$

2. $xH \subset \text{Fix}(x) \cup \overline{S_x}$;

3. $x(S_x \cup S_y \cup S_{y^{-1}}) \subset S_x$, for $x \neq y, y^{-1}$.

Then H is a fundamental domain for the action of F_2 on T, and T is a simplicial tree.

Proof. We may assume that $\text{Fix}(a) \cap \text{Fix}(b) = \emptyset$, since otherwise, by minimality, T is a point and the lemma is trivial.

In order to show that $F_2 H = T$, it suffices to show that $F_2 H$ is connected; $F_2 H$ is then an invariant subtree of T, which must be all of T since T is minimal. Let T_n denote the union of all translates gH, where $g \in F_2$ has word length less than or equal to n in the generators a and b. We show that T_n is connected for all n, by induction on n. If $n = 0$, this is just the statement that H is a subtree. If $n = 1$, this follows from hypothesis 1 of the lemma. If $n > 1$, $T_{n+1} = aT_n \cup a^{-1}T_n \cup bT_n \cup b^{-1}T_n$. Each of these four sets is connected by induction, and each contains H so their union is connected.

To show that $gH \cap H$ is of the right form for any $g \in F_2$, we use the following fact.

Claim. *Let $g \in F_2$, let w be the unique reduced word in a and b representing g, and let x be the first letter of w. Then*

(i) $gH \subset \overline{S_x} \cup (\mathrm{Fix}(x) \cap H)$, and

(ii) if gH is not contained in $\overline{S_x}$, then $g = x^k$ for some k.

Proof. The proof is by induction on the length n of w. If $n = 1$, statement (i) follows from hypothesis 2 and statement (ii) is trivial.

If $n > 1$, we have $w = xv$, where v is a word with first letter $y \neq x^{-1}$. Let p be a point of H. By induction, we know that $vp \in \overline{S_y} \cup (\mathrm{Fix}(y) \cap H)$. If $vp \in \overline{S_y}$, then $wp = xvp \in \overline{S_x}$ by hypothesis 3.

If vp is not in $\overline{S_y}$, then $vp \in \mathrm{Fix}(y) \cap H$, and we have two cases to consider. If $y \neq x$, then vp is not in $\mathrm{Fix}(x)$, since $\mathrm{Fix}(x)$ and $\mathrm{Fix}(y)$ are disjoint. Thus by hypothesis 2, $wp = xvp \in \overline{S_x}$. If $y = x$, then $wp = xvp = vp \in \mathrm{Fix}(x)$ since vp is not in $\overline{S_x}$, statement (ii) implies inductively that v, and hence w, is a power of x.   ⊏

We now show that $gH \cap H$ is either empty, a single point of H, or is contained in $\mathrm{Fix}(g)$. If g is not a power of x, then $gH \subset \overline{S_x}$ by part (ii) of the claim, so $gH \cap H \subset \partial S_x$. Since $gH \cap H$ is connected, and ∂S_x is discrete, $gH \cap H$ is either empty or a single point in this case.

If $g = x^k$ for some $k > 0$, the claim shows that $gH \cap H \subset (\mathrm{Fix}(x) \cap H) \cup \overline{S_x}$. We now claim that $\mathrm{Fix}(x) \cap H = \mathrm{Fix}(x^k) \cap H$, from which it follows that $gH \cap H \subset Fix(g)$. The inclusion $\mathrm{Fix}(x) \cap H \subset \mathrm{Fix}(x^k) \cap H$ is trivial. To show the opposite inclusion, let p be a point of H with $x^k p = p$. By hypothesis 2, $xp \in \mathrm{Fix}(x) \cup \overline{S_x}$. If $xp \in Fix(x)$, we are done. If $xp \in \overline{S_x}$, then by hypothesis 3, $x^i p$ is in $\overline{S_x}$ for all i. Thus $p = x^k p \in \partial S_x$. It follows that xp and $x^2 p$ are also in the ∂S_x; if they were in the interior of S_x, then $x^k p$ would be as well. The arc $[p, xp]$ is in H since p and xp are; therefore $x[p, xp] = [xp, x^2 p]$ is in $\overline{S_x} \cap H = \partial S_x$. But ∂S_x is a discrete set, so $[xp, x^2 p]$ is a single point, i.e. $xp = x^2 p = p$.

To see that T is a simplicial tree, consider the minimal subtree K of H which contains the boundary of H. This subtree K is a simplicial tree since the boundary of H is finite. We claim that $F_2 K$ is a simplicial tree; by minimality this will imply that $F_2 K = T$, and hence that T is simplicial. We have shown above that $H \cap gH$ is either empty, a single point, or equal to $\mathrm{Fix}(x) \cap H$ for $x = a$ or $x = b$; the last case occurs only if $g = x^k$. Thus, if $K \cap gK \subset H \cap gH$ is non-empty, then it is either a single point or one of the two trees $\mathrm{Fix}(x) \cap K$. This implies that $F_2 K$ is simplicial; it

remains only to show that F_2K is connected. For this it suffices to show that $K \cap xK$ is not empty for $x \in \{a, a^{-1}, b, b^{-1}\}$. If $\text{Fix}(x) \cap H$ is empty, then x maps a boundary point of H to a boundary point of H, so $K \cap xK$ is non-empty. Thus we need only show that if $\text{Fix}(x) \cap H$ is non-empty then so is $\text{Fix}(x) \cap K$. In fact, if x fixes a point of H then it fixes a boundary point. To see this consider a point p of ∂S_x and let $[p, q]$ be the bridge from p to $\text{Fix}(x) \cap H$. The interior points of $[p, q]$ are not fixed by x, so the image of the interior of $[p, q]$ under x must be contained in $\overline{S_x}$. Thus $q \in \overline{S_x}$, proving that a point of ∂H is fixed by x. $\square$

Proposition 3.2. *Let $F_2 \times T \to T$ be a minimal action with maximal cyclic arc stabilizers. If some primitive element $a \in F_2$ has a fixed point, then T is simplicial.*

Proof. Since a is primitive, there is an element $b \in F_2$ such that a and b generate F_2. We divide the proof into cases, applying Lemma 3.1 in each case to find a fundamental domain H for the action.

Case 1. b has a fixed point in T.

Since the action is non-trivial, we have $\text{Fix}(a) \cap \text{Fix}(b) = \emptyset$. Let $H = [p, q]$ be the bridge from $\text{Fix}(a)$ to $\text{Fix}(b)$. Let η_p and η_q be the germs at p and q determined by H. We take S_a (resp. $S_{a^{-1}}$) to be the union of all open rays in T emanating from p with germ $a^n \eta_p$ for some $n > 0$ (resp. $n < 0$). Similarly, let S_b (resp. $S_{b^{-1}}$) be the union of all open rays in T emanating from q with germ $b^n \eta_q$ for some $n > 0$ (resp. $n < 0$). (See Figure 1)
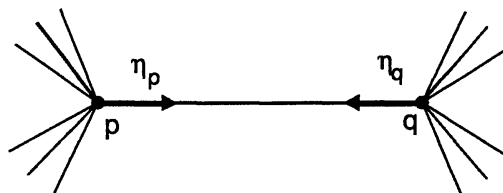

Figure 1

Note that $aH \cap H = \{p\}$, since otherwise a would fix a subarc of H, and H would not be the bridge from $\text{Fix}(a)$ to $\text{Fix}(b)$. Similarly, $bH \cap H = \{q\}$. Also, $a^n H \cap H = \{p\}$ and $b^n H \cap H = \{q\}$ for all $n \neq 0$, since otherwise the generator of an arc stabilizer would be included into a vertex stabilizer as a proper power, contradicting our hypothesis that the action have maximal cyclic arc stabilizers. This shows that the sets $S_a, S_{a^{-1}}, S_b, S_{b^{-1}}$, and H are disjoint. The other hypotheses of Lemma 3.1 are easily verified, showing that H is a fundamental domain for the action.

Case 2. b is hyperbolic and $\text{Fix}(a) \cap \text{Axis}(b) = \emptyset$.

Let $[p, q]$ be the bridge from $\text{Fix}(a)$ to $\text{Axis}(b)$, let $[r, br]$ be a segment of $\text{Axis}(b)$ containing q in its interior, and let $H = [p, q] \cup [r, br]$. Let η_p be the germ at p determined by $[p, q]$, let η_r be the germ at r determined by $b^{-1}[r, br]$, and let η_{br} be the germ at br determined by $b[r, br]$. We take S_a (resp. $S_{a^{-1}}$) to be the union of all open rays in T emanating from p with germ $a^n \eta_p$ for some $n > 0$ (resp. $n < 0$). We take S_b (resp. $S_{b^{-1}}$) to be the union of all open rays in T emanating from br (resp. r) with germ η_{br} (resp. η_r). (See Figure 2)

Figure 2

Case 3. b is hyperbolic and $\text{Fix}(a) \cap \text{Axis}(b)$ is a point.

Let $\{p\}$ be the intersection of $\text{Fix}(a)$ and $\text{Axis}(b)$, and let $[r, br]$ be a segment of $\text{Axis}(b)$ with p as its midpoint. If $a^n[p, br] = [p, r]$ for some n, then ba^n stabilizes br, and we are in Case 1 with b replaced by ba^n. If $a^n[p, br]$ is never equal to $[p, r]$, set $H = [r, br]$. Let $\eta_{p,1}$ be the germ at p determined by $[p, br]$, and $\eta_{p,2}$ be the germ at p determined by $[p, r]$. Let η_r be the germ at r determined by $b^{-1}[br, p]$, and η_{br} be the germ at br determined by $b[r, p]$. We take S_a (resp. $S_{a^{-1}}$) to be the union of all open rays in T emanating from p with germ $a^n \eta_{p,1}$ or $a^n \eta_{p,2}$ for some $n > 0$

(resp. $n < 0$). We take S_b (resp. S_{b-1}) to be the union of all open rays in T emanating from br (resp. r) with germ η_{br} (resp. η_r). (See Figure 3)

Figure 3

Case 4. b is hyperbolic and $\text{Fix}(a) \cap \text{Axis}(b)$ is an interval of length less than $|b|$.

Let $[p, q] = \text{Fix}(a) \cap \text{Axis}(b)$, and let $H = [r, br]$ be a segment of $\text{Axis}(b)$ containing $[p, q]$ in its interior. We may assume $q \in [p, br]$. Let η_p be the germ at p determined by $[p, r]$, and η_q be the germ at q determined by $[q, br]$. Let η_r be the germ at r determined by $b^{-1}[br, q]$, and η_{br} be the germ at br determined by $b[r, p]$. We take S_a (resp. S_{a-1}) to be the union of all open rays in T emanating from p with germ $a^n \eta_p$ or from q with germ $a^n \eta_q$ for some $n > 0$ (resp. $n < 0$). We take S_b (resp. S_{b-1}) to be the union of all open rays in T emanating from br (resp. r) with germ η_{br} (resp. η_r). (See Figure 4)
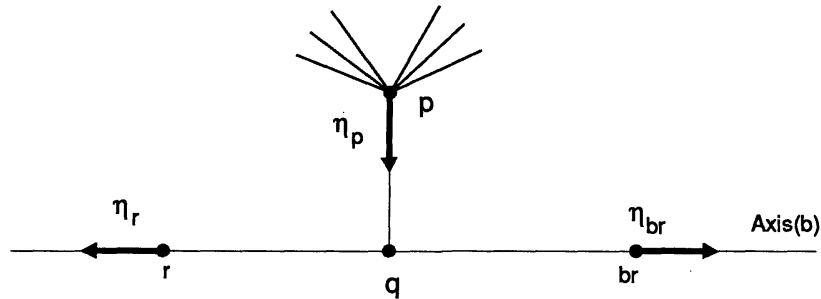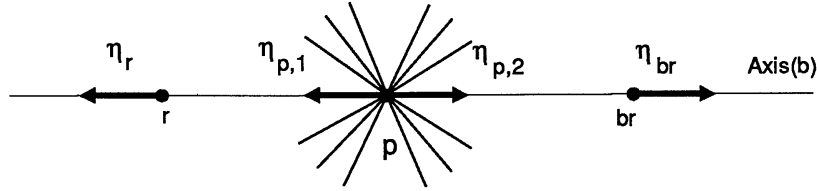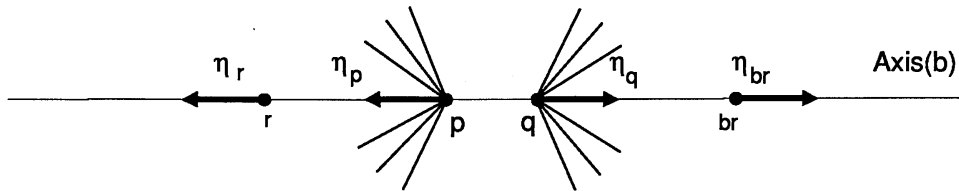

Figure 4

Case 5. b is hyperbolic and $\text{Fix}(a) \cap \text{Axis}(b)$ is an interval of length equal to $|b|$.

Let $H = [p, bp] = \text{Fix}(a) \cap \text{Axis}(b)$. Let η_p be the germ at p determined by $b^{-1}H$, and η_{bp} the germ at bp determined by bH. We take S_a (resp. S_{a-1}) to be the union of all open rays in T emanating from p with germ $a^n \eta_p$ or from bp with germ $a^n \eta_{bp}$ for some $n > 0$ (resp. $n < 0$). We take S_b

Figure 5

(resp. S_{b-1}) to be the union of all open rays in T emanating from p (resp. bp) with germ η_p (resp. η_{bp}). (See Figure 5)

Case 6. b is hyperbolic and $\mathrm{Fix}(a) \cap \mathrm{Axis}(b)$ is an interval of length greater than $|b|$.

Let $[p, q] = \mathrm{Axis}(b) \cap \mathrm{Fix}(a)$. Then a and $b^{-1}ab$ both fix an initial segment of $[p, q]$. But a and $b^{-1}ab$ do not commute, so the action has non-cyclic arc stabilizers, contradicting the fact that the action is a limit of free simplicial actions.   □

Definition. Let $F_2 \times T \to T$ be a simplicial action. The *quotient diagram* for the action is the quotient T/F_2 with vertices and edges labelled by the isomorphism type of their stabilizers. Two quotient diagrams are isomorphic if there is an isometry between the graphs so that corresponding edges and vertices have the same labels.

Given the fundamental domain for a simplicial action, one can easily construct the quotient diagram. Figure 6 shows the quotient diagrams, without specifying the lengths of the edges, in each case of the previous lemma.

We will need to be able to determine when two actions which satisfy the hypotheses of the previous proposition are actually isomorphic. Recall that two actions $\alpha_1 \colon G \times T_1 \to T_1$ and $\alpha_2 \colon G \times T_2 \to T_2$ on simplicial R-trees are isomorphic if there is an equivariant isometry $\phi \colon T_1 \to T_2$. Fundamental domains $H_1 \subset T_1$ and $H_2 \subset T_2$ are said to be isomorphic if there is an isometry $\psi \colon H_1 \to H_2$ sending vertices to vertices, such that the stabilizer of each vertex and edge in H_1 is the same as the stabilizer of its image under ψ. If α_1 and α_2 are isomorphic actions, and $H_1 \subset T_1$ is any fundamental domain for α_1, then $H_2 = \phi(H_1) \subset T_2$ is an isomorphic fundamental domain for α_2. Conversely, if $\psi \colon H_1 \to H_2$ is an isomorphism of fundamental
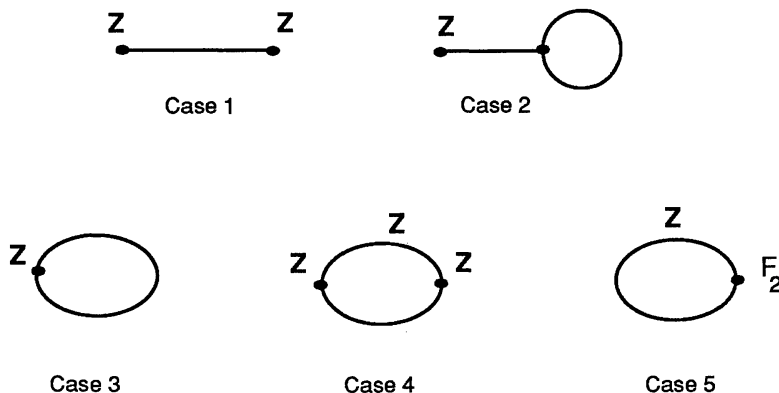
Figure 6

domains for α_1 and α_2, then ψ extends to an equivariant isometry of trees, so α_1 and α_2 are isomorphic.

The analysis in the proof of Proposition 3.2 includes the computation of a fundamental domain and of the stabilizers of the edges and vertices of the fundamental domain for any action satisfying the hypotheses of the proposition; thus a classification of all such actions is implicit in the proof. The following proposition makes this explicit.

Proposition 3.3. *Let a be a primitive element of F_2 which fixes a vertex in each of two actions $F_2 \times T_1 \to T_1$ and $F_2 \times T_2 \to T_2$ on simplicial $\mathbb{R}$-trees with maximal cyclic arc stabilizers. Assume that the quotient diagrams are isomorphic. Then*

1. *If T_i/F_2 is a single edge (Case 1 of Proposition 3.2), then there are elements b_1 and b_2 of F_2 such that $\{a, b_i\}$ is a basis of F_2 and b_i stabilizes a vertex of T_i. The actions are isomorphic if and only if b_1 is conjugate to b_2.*

2. *If T_i/F_2 has one free edge and one loop (Case 2 of Proposition 3.2), there are elements b_1 and b_2 of F_2 such that $\{a, b_i\}$ is a basis of F_2 and b_i is a stable letter for the HNN decomposition determined by the action of F_2 on T_i. The actions are isomorphic if and only if b_1 is conjugate to b_2.*

3. *If T_i/F_2 is homeomorphic to a circle (Cases 3,4 and 5 of Proposition 3.2), the actions are isomorphic.*

Proof. Let $p_i \in T_i$ be a vertex fixed by a. In the proof of Proposition 3.2 we describe a fundamental domain $H_i \subset T_i$ with $p_i \in D_i$. The hypothesis that

T_1/F_2 is isometric to T_2/F_2 implies that these fundamental domains are isometric. However, some choices were made in the construction of these fundamental domains; we must show in each case that the fundamental domains can be constructed so that they are isomorphic.

Note that if $\{a, b\}$ is a basis of F_2 then every basis which contains a has the form $\{a, a^m b a^n\}$ where $m, n \in \mathbb{Z}$. In particular if $\{a, b_1\}$ and $\{a, b_2\}$ are bases then b_1 is conjugate to b_2 if and only if $b_1 = a^n b_2 a^{-n}$ for some integer n.

If the actions are in Case 1 then any edge incident to p_i is a fundamental domain. For each integer n there is such an edge for which the stabilizer of the other endpoint is the cyclic group generated by $a^n b_i a^{-n}$. Thus there exist isomorphic fundamental domains for the two actions if and only if b_1 is conjugate to b_2.

If the actions are in Case 2 then any arc which consists of two edges and contains p_i as an endpoint is a fundamental domain. The edge which is disjoint from p_i is a fundamental domain for the action of an element of the form $a^n b_i a^{-n}$ by translation on its axis, and every integer n arises for some fundamental domain. Thus the two actions have isomorphic fundamental domains if and only if b_1 is conjugate to b_2.

In the other cases, let b be any primitive element such that that $\{a, b\}$ is a basis of F_2. Then in both trees the axis of b contains $\mathrm{Fix}(a)$, and a fundamental domain for the action of F_2 on T_i is constructed by taking a fundamental domain for the translation action of b along its axis which is symmetric about the midpoint of $\mathrm{Fix}(a)$. These are clearly isomorphic fundamental domains, so the actions are isomorphic.   ⊏

The following corollary will be used in the proof that all actions satisfying the hypotheses of Proposition 3.2 are limits of free simplicial actions.

Corollary 3.4. *Let $\{a, b\}$ be a basis of F_2. Let r be a non-negative real number or infinity. Up to scaling of the metric, there is a unique action of F_2 on an $\mathbb{R}$-tree with maximal cyclic arc stabilizers for which*

1. $|a| = 0$

2. $|aba^{-1}b^{-1}|/|b| = r$ and

3. $|b|$ is minimal among all primitive elements b such that $\{a, b\}$ is a basis of F_2.

This action is simplicial.

Proof. By Proposition 3.2, any action with maximal cyclic edge stabilizers and with $|a| = 0$ is simplicial. Thus to construct actions with properties (1)-(3), we need only give give fundamental domains; for this, we will refer to the cases of Proposition 3.2.

An action with $r = \infty$ (i.e. $|b| = 0$), is given by Case 1. The length of the edge in the quotient is $|aba^{-1}b^{-1}|/4$.

If $2 < r < \infty$, an action is given by Case 2, where the loop in the quotient has length $|b|$ and the other edge has length $(|aba^{-1}b^{-1}| - 2|b|)/4$.

If $r = 2$ an action is given by Case 3, where the length of the loop in the quotient is $|b|$.

If $0 < r < 2$ an action is given by Case 4, where the length of the edge fixed by a is $|b| - |aba^{-1}b^{-1}|/2$, and the length of the other edge is $|aba^{-1}b^{-1}|/2$.

If $r = 0$ an action is given by Case 5, where the length of the loop in the quotient is $|b|$.

By Proposition 3.3, the actions are uniquely determined by the primitive element a and the quotient diagram in Cases 3, 4 and 5. In cases 1 and 2, the additional fact that $|b|$ is minimal guarantees uniqueness.

$\square$

We finish this section by showing that the actions which we have been considering are actually points of $\overline{X}$.

Proposition 3.5. *Let $F \times T \to T$ be an action on a simplicial $\mathbb{R}$-tree with maximal cyclic arc stabilizers. If a primitive element a of F_2 fixes a vertex of T then this action is contained in $\overline{X}$.*

Proof. We will give an explicit sequence of free actions $F_2 \times T_n \to T_n$, where T_n is a simplicial $\mathbb{R}$-tree, which converge to the given action in the length function topology. Let $|\cdot|_n$ denote the length function for the action on T_n, and let $r = |aba^{-1}b^{-1}|/|b| \in [0, \infty]$. By Lemma 2.1, the limit of any convergent sequence of free actions is an action with maximal cyclic edge stabilizers. Thus by Corollary 3.4 we need only exhibit a convergent sequence of free actions so that $\lim_{n \to \infty} |a|_n = 0$ and $\lim_{n \to \infty} |aba^{-1}b^{-1}|_n/|b|_n = r$.

As in [4] we will describe the actions in this sequence by giving the quotient of the action together with a length for each edge and a labelling of the edges in the complement of a maximal tree by elements of F_2. (The labelling specifies an identification of F_2 with the fundamental group of the quotient.)

If $r = \infty$, the action on T_n is described by Figure 7, where the edges labelled a and b have length $1/n$ and the unlabelled edge has length 1.
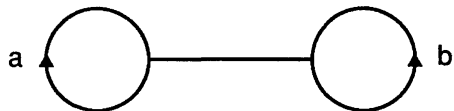
Figure 7

If $2 \leq r < \infty$, the action on T_n is again described by Figure 7, where the edge labelled a has length $1/n$, the edge labelled b has length 1 and the unlabelled edge has length $(r - 2)/4$.

If $0 < r < 2$, the action on T_n be described by Figure 8, where the edge labelled a has length $1/n$ and the edge labelled $a^n b$ has length $r/(2 - r)$.
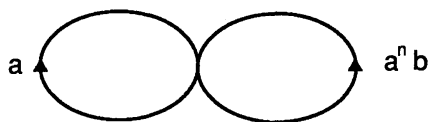

Figure 8

If $r = 0$, the action on T_n is again described by Figure 8, but both edges have length $1/n$.   □

§4. A division process.

Throughout this section we fix an action $F_2 \times T \to T$.

Definition. If $\{a, b\}$ is a basis of F_2 with a and b hyperbolic, we denote by $\Delta(a, b)$ the length of the segment $\text{Axis}(a) \cap \text{Axis}(b)$. If this intersection is empty, we set $\Delta(a, b) = 0$.

In the case where one of a or b has length less than or equal to $\Delta(a, b)$ it is possible to change basis by a Nielsen transformation to either reduce the length of the overlap of the axes or to produce a basis element with a fixed point:

Proposition 4.1. *Let* $\{a, b\}$ *be a basis for* F_2, *with* $0 < |b| \le |a|$, *such that the translation directions of* a *and* b *along their axes agree. Assume that* $|b| \le \Delta(a, b)$. *If* $|b^{-1}a| \ne 0$ *then* $\Delta(b, b^{-1}a) = \Delta(a, b) - |b|$. *If* $|b| < \Delta(a, b)$ *then* $|b^{-1}a| = |a| - |b|$.

Proof. Let p be the initial endpoint of the overlap $\text{Axis}(a) \cap \text{Axis}(b)$, where this segment is oriented in the common translation direction of a and b. We consider three cases.

Case 1. $|b| = |a|$

In this case $b^{-1}a$ fixes p, so $|b^{-1}a| = |a| - |b| = 0$.

Case 2. $|b| < \Delta(a, b)$ and $|b| < |a|$.

It is an easy exercise to check that the segment from p to $b^{-1}ap$ is a fundamental domain for the action of $b^{-1}a$ on its axis. This segment has length $|a| - |b|$ and meets $\text{Axis}(b)$ in its initial subsegment of length $\Delta(a, b) - |b|$ (see Figure 9).
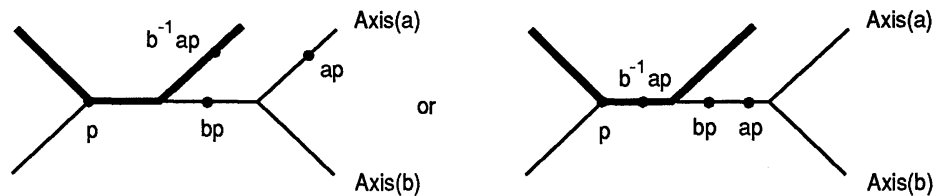

Figure 9

Case 3. $|b| = \Delta(a, b) < |a|.$

In this case we claim that $\mathrm{Axis}(b^{-1}a)$ is either disjoint from $\mathrm{Axis}(b)$ or meets it in a single point. In particular $\Delta(b^{-1}a, b) = \Delta(a, b) - |b| = 0.$ For this it suffices to check that the segment from p to $b^{-1}ap$ meets $\mathrm{Axis}(b)$ only at p, since this segment contains a fundamental domain for the action of $b^{-1}a$ on its axis. This is also an easy exercise; see Figure 10.
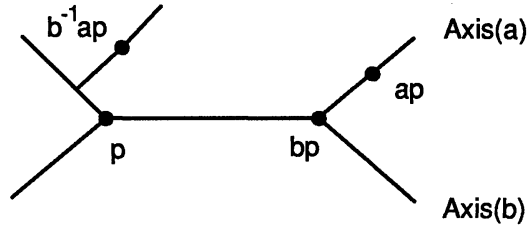
Figure 10

$\square$

The proposition above gives rise to a *division process*: beginning with a basis $\{a, b\}$ of F_2 consisting of hyperbolic elements, we first replace b by b^{-1} if necessary so that the axes for a and b are oriented in the same direction on their overlap; next, we interchange a and b if necessary so that $|b| \le |a|$. If $|b| \le \Delta(a, b)$, we replace the longer element a with $b^{-1}a$. If $|b^{-1}a| > 0$, we repeat the process from the beginning.

The process terminates if we obtain an elliptic generator or if we obtain a basis $\{a', b'\}$ whose shorter element is longer than $\Delta(a', b')$. If we obtain an elliptic generator, the action is simplicial and can be determined explicitly by the methods of section 3. If the shorter basis element has length greater than the overlap, we have the following proposition.

Proposition 4.2. *Let $\{a, b\}$ be a basis of F_2 with $|a| \ge |b| > \Delta(a, b)$. Then the action is free and simplicial.*

Proof. We consider separately the cases when the axes for a and b are disjoint and when they have non-empty intersection.

Case 1. $\mathrm{Axis}(a) \cap \mathrm{Axis}(b) = \emptyset$.

Let $[p, q]$ be the bridge from $\mathrm{Axis}(a)$ to $\mathrm{Axis}(b)$, and let $H = [p, ap] \cup [p, q] \cup [q, bq]$. Let η_p be the germ at p in the direction $a^{-1}[p, ap]$, let η_{ap} be the germ at ap in the direction $a[p, ap]$, let η_q be the germ at q in the direction $b^{-1}[q, bq]$, and let η_{bq} be the germ at bq in the direction $b[q, bq]$. We take S_a (resp. $S_{a^{-1}}$) to be the union of all open rays in T emanating from ap (resp. p) with germ η_{ap} (resp. η_p). We take S_b (resp. $S_{b^{-1}}$) to be the union of all open rays in T emanating from bq (resp. q) with germ η_{bq} (resp. η_q). We now apply Lemma 3.1 to see that H is a fundamental domain for the action of F_2 on T.

The action in this case is free and simplicial, with quotient a "barbell" as shown in Figure 11.
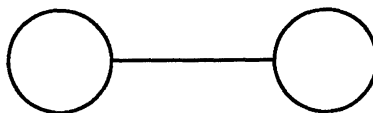
Figure 11

Case 2. $\mathrm{Axis}(a) \cap \mathrm{Axis}(b) \neq \emptyset$.

Let p be the initial point of the overlap and set $H = [p, ap] \cup [p, bp]$. Let $\eta_{p,1}$ be the germ at p in the direction $a^{-1}[p, ap]$, let η_{ap} be the germ at ap in the direction $a[p, ap]$, let $\eta_{p,2}$ be the germ at p in the direction $b^{-1}[p, bp]$, and let η_{bp} be the germ at bp in the direction $b[p, bp]$. We take S_a (resp. $S_{a^{-1}}$) to be the union of all open rays in T emanating from ap (resp. p) with germ η_{ap} (resp. $\eta_{p,1}$). We take S_b (resp. $S_{b^{-1}}$) to be the union of all open rays in T emanating from bp (resp. p) with germ η_{bp} (resp. $\eta_{p,2}$). Lemma 3.1 again applies to show that H is a fundamental domain for the action of F_2 on T; the action is free and simplicial, with quotient a "theta graph" as shown in Figure 12. If $\Delta(a, b) = 0$, the theta graph degenerates to a "rose".

$\square$

Remark. If the division process does not terminate, the lengths of the basis elements approach zero, while the difference $\Delta(a, b) - |a| - |b|$ remains constant.
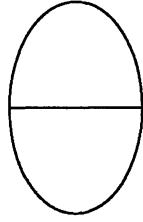
Figure 12

For the rest of the section we fix a basis $\{a, b\}$ of F_2. Let α be any action of F_2 on an $\mathbb{R}$-tree T.

Definition. If a and b are hyperbolic and $\mathrm{Axis}(a) \cap \mathrm{Axis}(b)$ is a nontrivial arc, we define the *slope* of α to be $\pm|a|/|b|$, where the sign is positive if and only if the translation directions of a and b agree on the overlap. If the overlap is empty or a point, or if a or b is not hyperbolic, then the slope is zero if $|a| < |b|$, infinite if $|a| > |b|$ and undefined if $|a| = |b|$.

Note that the division process always terminates if the slope is rational.

Proposition 4.3. *Let $F_2 \times T_1 \to T_1$ and $F_2 \times T_2 \to T_2$ be actions on $\mathbb{R}$-trees. Assume that for each action there is a primitive element which has a fixed point and that the quotient diagrams are isomorphic and homeomorphic to a circle (Cases 3, 4, and 5 of Proposition 3.2). The actions are isomorphic if and only if they have the same slope.*

Proof.

By Proposition 3.3, actions of this type are determined up to isomorphism by the conjugacy class of the primitive element of length zero and the quotient diagram. Thus we must show that the slope determines which primitive element has a fixed point. This is clear if the slope is 0 or ∞. Otherwise a and b are hyperbolic in each action, so the division process can be started. The slope is rational because both $|a|$ and $|b|$ are integer multiples of the total length of the quotient circle. Thus the division process must terminate. Since the actions are not free, it terminates by producing a primitive element with a fixed point.

We claim that for both actions, the division process produce the same primitive element. This follows from the observation that the division process, and the associated sequence of basis changes (Nielsen transformations) are completely determined by the slope. The sign of the slope determines

whether b should be replaced by b^{-1} at the beginning of the process. The generators a and b must be interchanged if and only if the absolute value of the slope is less than 1. From that point onward the process is determined by the ratio $|a|/|b|$. $\qquad\square$

Next we use the division process to classify actions with overlap which is either shorter or longer than $|a| + |b|$.

Proposition 4.4. *If $\Delta(a,b) < |a| + |b|$, the action is simplicial.*

Proof. By Proposition 4.2, it suffices to show that the division process always terminates if $\Delta(a,b) < |a| + |b|$. Assume that the division process does not terminate. Then we have a sequence $\{a_i, b_i\}$ of bases for F_2 such that the lengths of a_i and b_i approach zero. But $\Delta(a_i, b_i) - |a_i| - |b_i| = \Delta(a,b) - |a| - |b|$ is constant and less than zero for all i. This is a contradiction, since $\Delta(a_i, b_i)$ is positive for all i. $\qquad\square$

Proposition 4.5. *If $\Delta(a,b) > |a|+|b|$, there is an arc stabilizer containing a free subgroup of rank two in F_2.*

Proof. Since $|a|+|b| < \Delta(a,b)$, we can apply the division process to obtain a sequence $\{a_i, b_i\}$ of bases; since $\Delta(a_i, b_i) - |a_i| - |b_i|$ remains (constant and) positive, the only way the process can terminate is if b_n is elliptic for some n. In this case, the basis $\{a_{n-1}, b_{n-1}\}$ has $|a_{n-1}| = |b_{n-1}| < \frac{1}{2}\Delta(a_{n-1}, b_{n-1})$. Now note that $b_{n-1}^{-1}a_{n-1}$ and the commutator $[a_{n-1}^{-1}, b_{n-1}^{-1}]$ each fix an initial segment of the overlap $\mathrm{Axis}(a_{n-1}) \cap \mathrm{Axis}(b_{n-1})$. Since these two elements generate a free group of rank two, we have a non-cyclic arc stabilizer.

If $|a|/|b|$ is not rational, the division process does not terminate. In this case the lengths of a_i and b_i approach zero, while $\Delta_i = \Delta(a_i, b_i)$ approaches a positive constant. Thus, for n sufficiently large, we have $|a_n|+2|b_n| < \Delta_n$. Since $|a_n| + |b_n| < \Delta_n$, the commutator $[a_n, b_n]$ fixes an initial segment of the overlap $\mathrm{Axis}(a_n) \cap \mathrm{Axis}(b_n)$, and since $|a_n|+2|b_n| < \Delta_n$, the commutator $[a_n, b_n^2]$ fixes a (smaller) initial segment of the same overlap. Since these two commutators generate a free group of rank two, we again have a non-cyclic arc stabilizer. $\qquad\square$

§5. Critical overlap: $\Delta(a, b) = |a| + |b|$

Throughout this section we let $F_2 \times T \to T$ be a minimal action and we assume that F_2 is generated by hyperbolic elements a and b, where the axes of a and b meet in a segment of length $\Delta = |a| + |b|$. Note that $|aba^{-1}b^{-1}| = 0$, so this action is not in X. As in section 4, we define the *slope* of the action to be the extended real number $m = \pm|a|/|b|$, where the sign is positive if and only if the translation directions of a and b agree on the overlap of their axes. For convenience we will assume that for the action on T the translation directions agree; there is no loss of generality since this condition holds after replacing b by b^{-1}. We begin by giving a simple construction, for each m, of a tree GT_m with a natural F_2-action. We then show that T is equivariantly isometric to GT_m.

The tree GT_m is the dual tree to the measured foliation of the punctured torus by lines of slope m. Specifically, it is constructed as the space of leaves of a PL measured foliation of a simplicial complex Σ, which is equivariantly homeomorphic to the universal cover of the punctured torus with a countable number of points added at infinity. The general theory of such measured foliations and their relation to actions on $\mathbb{R}$-trees is developed in [5]. While our construction is self-contained, we use results from [5] to show that the space which we construct is an $\mathbb{R}$-tree.

If m is rational then the division process described in the last section, when applied to $\{a, b\}$, terminates at a basis containing one elliptic generator which fixes a fundamental domain for the action of the other generator on its axis. Such actions are described in Case 5 of Proposition 3.2: the action is simplicial, the quotient graph is a circle with one vertex, the stabilizer of the edge is infinite cyclic, and the vertex stabilizer is free of rank 2. The tree GT_m, for m rational, is the dual tree to a foliation of the punctured torus by parallel simple closed curves, which is a simplicial tree. One can check that T is equivariantly isometric to GT_m. However, the details of the argument in the rational case differ from those in the irrational case; since the rational case can be handled by the methods of section 3, we assume for the rest of this section that m is irrational.

Let S be the unit square $[0, 1] \times [0, 1]$ foliated by line segments of slope m. To construct the simplicial complex Σ, we take a copy wS of S for each reduced word w in the generators a and b of F_2, and glue them along their edges according to:

$$w_1(1, y) \sim w_2(0, y) \text{ if } w_1^{-1}w_2 = a$$

$$w_1(x, 0) \sim w_2(x, 1) \text{ if } w_1^{-1} w_2 = b$$

There is a natural left action of F_2 on Σ. The foliations on the squares wS patch together to give a PL foliation $\mathcal{F}_m$ of Σ. There is a transverse measure on $\mathcal{F}_m$ which is inherited from the transverse measure on the square, in which the length of a monotone arc transverse to the leaves is the length of its projection onto the direction of slope -1.

Deleting the corner points of S and identifying opposite sides gives a punctured torus. The complement in Σ of the set of corner points of the squares wS can be identified in a natural way with the universal cover of this punctured torus. We also have a natural identification of the universal abelian cover of this torus with $\mathbb{R}^2 - \mathbb{Z}^2$, where b acts by shifting downward and a acts by shifting to the right. The covering projection from the universal cover to the universal abelian cover extends to a map $\pi \colon \Sigma \to \mathbb{R}^2$ which sends S to $[0, 1] \times [0, 1]$.

A leaf of $\mathcal{F}_m$ meets a square wS in a line segment. This gives each leaf a natural simplicial structure in which these line segments are the edges. Following Gillet and Shalen, a leaf of $\mathcal{F}_m$ is called *classical* if does not contain a corner of any translate of S, and *singular* if it does contain a corner. A classical leaf is homeomorphic to the real line, while a singular leaf is homeomorphic to a simplicial tree, with a vertex of infinite valence coming from each corner which it contains. Note that the map $\pi \colon \Sigma \to \mathbb{R}^2$ takes leaves of $\mathcal{F}_m$ to lines of slope m. It follows that a singular leaf has only one vertex which is not bivalent, since a line of irrational slope passes through at most one point of the integer lattice in $\mathbb{R}^2$.

In the language of [5], Σ is a *simply-connected uniform $\mathbb{R}$-foliated surface with points at infinity*. It is proved in [5, Theorem 5.20 and Proposition 5.25] that the *leaf space* of such a foliation is an $\mathbb{R}$-tree. We define GT_m to be the leaf space of $\mathcal{F}_m$. Thus the underlying topological space of GT_m is the quotient $\Sigma / \sim$, where $p \sim q$ if and only if p and q lie on the same leaf of the foliation. The metric on GT_m is induced by the transverse measure on $\mathcal{F}_m$; the distance between two leaves is the measure of an arc joining them which meets each leaf of $\mathcal{F}_m$ in a connected set.

We will call a point of GT_m a *vertex point* if it corresponds to a singular leaf of the foliation, and an *edge point* if it corresponds to a classical leaf. This agrees with the usual notion of a vertex point as a point from which more than two germs of arcs emanate. An edge point has exactly two germs of arcs emanating from it. This implies that if an edge point is

contained in the interior of two arcs then it is contained in the interior of their intersection.

The next step is to define an equivariant map $\phi\colon GT_m \to T$. We will do this by defining a map ϕ_0 on Σ which is constant on leaves of $\mathcal{F}_m$; this map factors through GT_m to give ϕ. To define ϕ_0, let $[p, q] \subset T$ be the intersection of the axes of a and b, where $abp = q$. By scaling the metric on T we may assume that $[p, q]$ has length $\sqrt{2}$. Define ϕ_0 on the diagonal of slope -1 of S to be an isometry with $\phi_0(0, 1) = p$ and $\phi_0(1, 0) = q$. The diagonal contains a unique point of each foliating line segment of S; extend ϕ_0 over S by making it constant on these line segments.

Next extend ϕ_0 to all of Σ equivariantly: if $x \in wS$, define $\phi_0(x) = w\phi_0(w^{-1}x)$. To check that ϕ_0 is well-defined, we must show both that $\phi_0(w(x, 0)) = \phi_0(wb(x, 1))$ and $\phi_0(w(1, y)) = \phi_0(wa(0, y))$ for all $x, y \in [0, 1]$ and all words w. This reduces to showing that $\phi_0(x, 0) = b\phi_0((x, 1))$ and $\phi_0(1, y) = a\phi_0((0, y))$. This can be checked using analytic geometry and remembering that $|a| + |b| = \sqrt{2}$ and $m = |a|/|b|$ (see Figure 13).
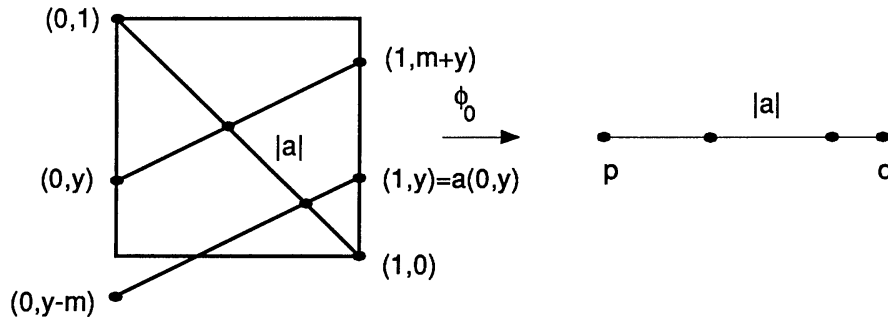
Figure 13

Finally, since ϕ_0 is constant on leaves of $\mathcal{F}_m$, it descends to an equivariant map ϕ on the quotient tree GT_m.

Definition. An arc α of GT_m is *straight* if α is the image of a non-degenerate line segment which is contained in some square wS and is transverse to the foliation on wS.

Note that the image of a square wS in GT_m is a straight arc which, by definition, is mapped isometrically by ϕ to the arc $\operatorname{Axis}(waw^{-1}) \cap \operatorname{Axis}(wbw^{-1})$ in T.

The following lemma shows that ϕ is a morphism of $\mathbb{R}$-trees in the sense of [8], i.e. that each point of GT_m is in the interior of an arc which is mapped isometrically by ϕ.

Lemma 5.1. *There is a straight arc through each point of GT_m.*

Proof. The preimage of a point in GT_m is a leaf of $\mathcal{F}_m$, which contains a point in the interior of some square. □

The map ϕ is also a morphism of $\mathbb{R}$-trees in the sense of [5], i.e. each arc of GT_m can be subdivided into a finite number of subarcs which are mapped isometrically by ϕ. This fact is contained in the following lemma.

Lemma 5.2. *The unique arc joining two points of GT_m is a finite union of straight arcs.*

Proof. Every pair of points in Σ can be joined by a path which meets only finitely many fundamental domains. Thus the unique arc joining the images of those points in GT_m is contained in the subtree of GT_m which is the image of this path. This subtree is a finite union of straight arcs. Since a nondegenerate subarc of a straight arc is straight, any arc contained in this subtree is also a finite union of straight arcs. □

We now apply the lemmas above to show that ϕ can only fail to be injective in a very special way: *If ϕ is not injective then there exist two straight arcs in GT_m, emanating from a vertex point v and meeting only at v, which are identified under ϕ.* We describe this by saying that ϕ has a *fold* at v. Assume that p and q are distinct points with $\phi(p) = \phi(q)$. Then the segment $[p, q]$ maps under ϕ to a closed loop in the tree T. By Lemma 5.2, $[p, q]$ can be subdivided into a finite number of straight arcs; since straight arcs are mapped isometrically by ϕ, a fold must occur at one of these subdivision points. We claim that this subdivision point must be a vertex point. There are exactly two germs of arcs emanating from an edge point and these are represented by a straight arc through that point. Thus these germs cannot be folded by ϕ.

Skora [14] has recently shown that if a closed surface group acts on an $\mathbb{R}$-tree with cyclic arc stabilizers then the tree is the leaf space of a measured foliation on the surface. His result also applies to surfaces with boundary if one assumes that the peripheral elements of the fundamental group have fixed points in the tree. In our situation this condition is satisfied since

$|a| + |b| = \Delta(a, b)$ implies that the commutator $aba^{-1}b^{-1}$ has a fixed point. For completeness we will prove that if T has cyclic arc stabilizers then ϕ is an isomorphism; the proof we give is essentially due to Skora.

Remark. It is not clear *a priori* whether the condition $|a| + |b| = \Delta(a, b)$ implies that the action is $\mathbb{Q}$-rank 2; if it were clear, then one could deduce from [5] that T is dual to a measured foliation on a punctured torus.

Proposition 5.3. *If the stabilizer of each arc in T is a cyclic subgroup of F_2 then $\phi: GT_m \to T$ is an equivariant isomorphism. In particular, if T is a limit of free simplicial actions then ϕ is an isomorphism.*

The main ingredient in the proof of Proposition 5.3 is an argument going back to Plante [12], which is also used in [9], [7] and [14]. We describe this argument before beginning the proof of 5.3.

By a *partial translation* on an interval I we mean an orientation preserving isometry whose domain and range are open subintervals of I. If f and g are partial translations of I the composition $f \circ g$ is a partial translation with domain $g^{-1}(\text{Domain}(f) \cap \text{Range}(g))$ if this set is non-empty. Otherwise we say that $f \circ g$ is undefined. Recall that a collection of partial translations on I is a pseudogroup if it is closed under the composition operation; the group laws hold pointwise on the subset where all of the relevant compositions are defined.

We consider a finitely generated pseudogroup Ψ of partial translations on an interval I, with generators $\{t_1, t_2, \dots, t_n\}$. For each point p of I we can construct the *tree of definition* D_p as follows: Let G_n be the Cayley graph of the free group on $\{t_1, t_2, \dots, t_n\}$, i.e. the graph whose vertices are the elements of the free group with an edge joining w to $t_i w$ for each group element w and generator t_i. A vertex of G_n determines a composition of the generators of Ψ. In G_n consider all the vertices such that the corresponding composition is defined in a neighborhood of p. This is the set of vertices of a subtree of G_n, which we define to be D_p. If w is a vertex of D_p then $\psi(w)$ will denote the corresponding element of Ψ.

Suppose that B is a subtree of D_p containing 1. We will say that B *grows exponentially* if the number of vertices of B contained in a ball of radius n about 1 grows exponentially as a function of n.

Observe that if I is embedded as an interval in the real line then each partial translation in Ψ extends uniquely to a translation of the line. This

defines a homomorphism from Ψ to a finitely generated group of translations of the line, generated by the extensions of $t_1, \ldots, t_n$. Since a finitely generated abelian group has polynomial growth, this has the following consequence.

If a subtree B of D_p has exponential growth for some $p \in I$ then, for large r, there is a family $W_r \subset B$ of words of length r in $t_1, t_2, \ldots, t_n$ such that

1. if v and w are both contained in W_r then the partial translations $\psi(v)$ and $\psi(w)$ agree on an open interval containing p; and
2. the cardinality of W_r grows exponentially in r.

In our situation the interval I will arise as an arc in the tree T and the generators of Ψ will be restrictions to subintervals of I of elements of F_2. Suppose that $g_1, g_2, \ldots, g_n$ are elements of F_2 whose restrictions to appropriate subintervals of I equal $t_1, t_2, \ldots, t_n$ respectively. Let ρ denote the homomorphism from the free group on $\{t_1, \ldots, t_n\}$ to F_2 with $\rho(t_i) = g_i$, for $i = 1, \ldots, n$.

Claim 5.4. Assume that there is a point p of I such that the tree of definition D_p contains an exponentially growing subtree B and that ρ is injective on B. Then T must have a non-cyclic arc stabilizer.

Proof. Consider the exponentially growing family W_r constructed above. Note that for $w \in W_r$ we have that the length in the generators a and b of $\rho(w)$ is bounded by Mr where M is an upper bound for the lengths of $g_1, g_2, \ldots, g_n$. Also, if we fix $w_0 \in W_r$ then $w_0^{-1} W_r$ is a set of words which fixes a subarc of I containing p. We now observe that a collection of commuting elements of F_2 whose lengths are bounded by Mr must have cardinality less than $2Mr + 1$, since they are all powers of some fixed element. Since the size of W_r grows exponentially, for large r there must exist $v, w \in W_r$ such that $\rho(w_0^{-1} w)$ and $\rho(w_0^{-1} v)$ do not commute. Thus T has a noncyclic arc stabilizer, and the claim is proved. □

Proof of Proposition 5.3. Note that $\phi \colon GT_m \to T$ is necessarily surjective since T is assumed to be minimal. Thus we need only show that if T has cyclic arc stabilizers then ϕ is injective. If ϕ is not injective then we have seen that there must be a fold at a vertex point of GT_m. Following [14], we will use this fold to construct a pseudogroup to which we can apply Plante's argument and Claim 5.4 to show that there is an arc in T with noncyclic stabilizer.

We begin by recording some important features of the action of F_2 on GT_m which follow easily from the construction of GT_m. We remind the reader that m is assumed to be irrational.

1. The action of F_2 is transitive on the vertex points of GT_m, which form a dense subset of GT_m.

2. The stabilizer of a vertex point is a (maximal) cyclic subgroup of F_2, generated by a conjugate of $aba^{-1}b^{-1}$. Under the action of the stabilizer of a vertex point v of GT_m there are exactly two orbits of germs of arcs emanating from v.

Let v be a vertex point at which ϕ folds, and let σ_1 and σ_2 be straight arcs emanating from v which are folded under ϕ. (Of course if there is a fold at one vertex then there is a fold at every vertex by equivariance.) Let $\sigma = \phi(\sigma_1) = \phi(\sigma_2)$ be the image of these arcs in the tree T.

We consider two cases.

Case 1. The germs of σ_1 and σ_2 are in the same orbit under the action of the stabilizer of v.

Since σ_1 and σ_2 are straight arcs, we have $\sigma_1 = g\sigma_2$ for some g in the stabilizer of v. This means that g stabilizes $\sigma \subset T$.

By the first fact above, the vertex points of GT_m are dense in GT_m. Choose any vertex point $w \neq v$ on σ_1 and a free group element h with $hv = w$ such that $h\sigma_1$ meets σ_1 in a positively oriented arc τ_1 emanating from w. Then $h^{-1}\tau_1$ is a subarc of σ_1 emanating from v whose image under ϕ is fixed by $h^{-1}gh$ as well as by g. Recall that g is in the stabilizer of v, which is a cyclic group. If h and g commute, then the subgroup of F_2 generated by h and the stabilizer of v is abelian, and hence cyclic, and properly contains the stabilizer of v, contradicting the second fact above. Thus g and $h^{-1}gh$ must not commute, i.e. they generate a free group of rank 2 which stabilizes a subarc of σ. This contradicts the fact that T has cyclic arc stabilizers.

Case 2. The germs of σ_1 and σ_2 are in distinct orbits under the action of the stabilizer of v.

Since σ_1 and σ_2 are straight arcs, we can take them to be images of line segments δ_1 and δ_2 contained in squares $w_1 S$ and $w_2 S$ of Σ. In fact, we may assume that δ_1 and δ_2 are contained in the diagonal of slope -1 of their respective squares, and begin at a corner. Since σ_1 and σ_2 are in distinct orbits under the action of the stabilizer of v, δ_1 and δ_2 begin at opposite

corners, i.e. we may assume that δ_1 begins at $w_1(0,1)$ and δ_2 begins at $w_2(1,0)$. We may further assume that δ_i is short enough so that its interior is disjoint from any line of slope m through a corner of the square w_iS. See Figure 14.
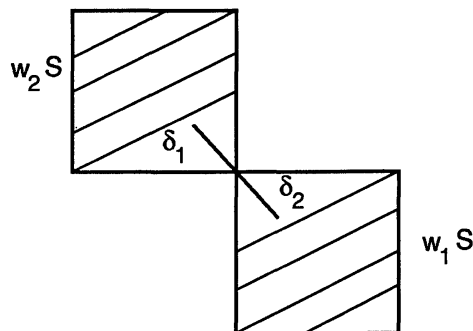

Figure 14

Our argument makes use of the Poincaré map for a measured foliation on a closed orientable surface. Consider an arc α which is transverse to the leaves of a measured foliation (and does not pass through any of the singularities of the foliation). Choose a normal direction for α. For each point p on α, follow the leaf of the foliation leaving in the chosen normal direction. This leaf will eventually either return to a point $P(p)$ on the arc or encounter a singularity. (Note that, since there are only finitely many singularities, only a finite number of leaves encounter singularities before returning to α.) This defines the Poincaré map P, the domain of which is the complement of a finite set in α. This map is an interval exchange map, i.e. the arc α can be divided into finitely many subintervals such that the restriction of P to the interior of each subinterval is a translation along α. Thus we obtain a finitely generated pseudogroup of partial translations on α such that the domains of the generators cover all but a finite subset of α.

The quotient of Σ by F_2 is a torus with a single distinguished point, the image of the corners of S. The foliation $\mathcal{F}_m$ descends to a foliation of the torus by lines of slope m with a single 2-pronged singularity at the distinguished point.

The Poincaré maps on the images of δ_1 and δ_2 in the torus determine pseudogroups of partial translations as above. These pseudogroups lift to δ_1 and δ_2, which are identified with σ_1 and σ_2 under the quotient map from

Σ to GT_m. Since σ_1 and σ_2 are identified under the equivariant map ϕ, we obtain two finitely generated pseudogroups Ψ_1 and Ψ_2 of partial translations on the arc σ in T. Let Ψ be the pseudogroup of partial translations on σ generated by Ψ_1 and Ψ_2. For p in the complement of a finite set of points in σ, there is a generator of Ψ_1 and a generator of Ψ_2 defined at p. Thus for p in the complement of a countable set of points, the tree of definition D_p for Ψ contains a rooted binary tree B_p of infinite depth consisting of the vertices of D_p corresponding to positive words in the generators for Ψ_1 and Ψ_2. Such a tree has exponential growth. Therefore, by Claim 5.4, Proposition 5.3 follows once we show that the map ρ from the free group on the generators of Ψ_1 and Ψ_2 to F_2 is injective on B_p.

To prove that ρ is injective on B_p we need to describe ρ in terms of the foliation $\mathcal{F}_m$ on Σ. Let x be a point in σ. Let x_1 and x_2 be the points in σ_1 and σ_2 which map to x under $\phi\colon GT_m \to T$. Let $\tilde{x}_1$ and $\tilde{x}_2$ be points in δ_1 and δ_2 which map to x_1 and x_2 under the quotient map from Σ to GT_m. Starting at the point $\tilde{x}_i$ of δ_i we can follow a leaf of $\mathcal{F}_m$ until it first meets a translate $g_i^{-1}\delta_i$ at a point $\tilde{y}_i$. The Poincaré map on the image of δ_i in the torus takes the image of $\tilde{x}_i$ to the image of $g_i\tilde{y}_i$. Let y_i be the point of σ_i which is the image of $\tilde{y}_i$. Thus g_i acts on the arc σ_i by a partial translation which sends x_i to the image of y_i. Thus if $x \in \sigma$ is contained in the domain of a generator t of the pseudogroup Ψ_i on σ then $\rho(t) = g_i$.

We will give a simple prescription, in terms of the map $\pi\colon \Sigma \to \mathbb{R}^2$, for how to find the word in the generators a and b of F_2 which represents g. Since the leaves of $\mathcal{F}_m$ project to lines of slope m, π maps the leaf joining $\tilde{x}_i$ and $\tilde{y}_i$ to a line of slope m from $\pi(\delta_i)$ to $\pi(g^{-1}\delta_i)$. Label each intersection of this line with the integral grid so that a transition from the square $\pi(gS)$ to the square $\pi(gxS)$ is labelled x, for $x \in \{a, a^{-1}, b, b^{-1}\}$. (Intersections with horizontal lines are labelled b or b^{-1} and intersections with vertical lines are labelled a or a^{-1}). Reading the labels from $\pi(\tilde{x}_i)$ to $\pi(\tilde{y}_i)$ then spells the word representing g_i.

The important thing to notice here is that, because we chose δ_i to be short, the word representing g_1 begins with b^{-1} while that representing g_2 begins with a. Also, notice that b^{-1} and a are the only letters which occur in g_i. We will think of the vertices of B_p as positive words in the generators of Ψ. Since ρ sends each generator of Ψ to a word which only involves the letters b^{-1} and a, no cancellation occurs when multiplying images of generators.

Let v and w be two vertices of B_p. We can write $v = uv_1$, $w = uw_1$, where the first letters of v_1 and w_1 are distinct. There is at most one generator of each of the pseudogroups Ψ_1 and Ψ_2 defined at a given point of σ. Assume without loss of generality that the first letter of v_1 is contained in Ψ_1; then the first letter of w_1 must be contained in Ψ_2. Therefore the first letter of $\rho(v_1)$ is b^{-1} and the first letter of $\rho(w_1)$ is a, showing that $\rho(v)$ is not equal to $\rho(w)$. $\qquad\square$

We close this section by noting that the actions considered here, with $\Delta(a, b) = |a| + |b|$ are in the closure of X, i.e. they are limits of free simplicial actions. Fix one such action α. If $m = |a|/|b|$ is rational, the division algorithm of Section 4 terminates at a basis $\{u, v\}$ for F_2 such that u has a fixed point. These actions are in the closure of X by Proposition 3.5, with quotient diagram given by Case 5 of Proposition 3.2. The primitive element u is uniquely determined up to conjugacy by the rational number m, and the action is uniquely determined by u, by Corollary 3.4. If m is irrational, approximate m by a sequence $\{m_i\}$ of rational numbers and let α_i be the unique action of slope m_i with $\Delta(a, b) = |a| + |b|$. We can then approximate each α_i by a free simplicial action β_i to obtain a sequence converging to α.

§6. A finite-dimensional embedding of $\overline{X}$.

Recall that the space X is topologized via an embedding $X \to \mathbf{P}^{\mathcal{C}}$ where $\mathbf{P}^{\mathcal{C}}$ is the infinite dimensional projective space whose coordinates are indexed by the set $\mathcal{C}$ of conjugacy classes in F_2. The image in $\mathbf{P}^{\mathcal{C}}$ of an action ϵ is the projective class of the translation length function $|\cdot|$ of the action The compactification $\overline{X}$ is defined to be the closure of X in $\mathbf{P}^{\mathcal{C}}$, so $\overline{X}$ is, by construction, embedded in an infinite dimensional projective space. In this section we will exhibit an embedding of $\overline{X}$ in a 3-dimensional projective space. We remark that the definition of $\overline{X}$ is patterned after Thurston's compactification of Teichmüller space, and that finite dimensional embeddings of Thurston's compactification can be obtained by projecting onto a subspace spanned by finitely many of the standard coordinates. It is not known, however, whether such a projection can ever be an embedding in the case of $\overline{X}_n$; our construction is not of this type. In an effort to make this section readable, several routine computational proofs have been left to the reader.

Every point in the closure $\overline{X}$ is the length function of a unique minimal action of F_2 on an $\mathbb{R}$-tree, by [3] and [10]. Let $\hat{X}$ denote the subset of $\overline{X}$ consisting of actions on simplicial trees. The space $\hat{X}$ is the union of a family of simplices which, abstractly, form a simplicial complex; however the weak topology on this complex does not agree with the topology on $\hat{X}$ Linear coordinates on the open simplices can be described as follows: Given a minimal action of F_2 on a simplicial $\mathbb{R}$-tree T, consider the quotient graph T/F_2. Each edge of T/F_2 inherits a metric from T in which it is isometric to an interval in the real line. We may vary the metric on T/F_2 so as to change the lengths of these edges and consider the lifted metrics on T This defines a k-parameter family of actions in $\hat{X}$, where k is the number of edges of T/F_2. Since the points of $\hat{X}$ are equivalence classes of actions under scaling of the metric on the tree, this family is a $(k-1)$-simplex embedded in $\mathbf{P}^{\mathcal{C}}$. Note that this is a linear embedding since the translation length of any element of F_2 is given by an integral linear combination of the lengths of edges in T/F_2.

Two actions have the same *type* if there is a homeomorphism between their quotient diagrams which preserves the labels. Any two actions in the same open simplex of $\hat{X}$ have the same type, which we will call the type of the simplex. The type of a closed simplex is the type of its interior.

Let α be an action in $\overline{X}$, and fix a basis $\{a, b\}$ of F_2. By applying the division process of Section 4 to $\{a, b\}$, we obtain either a basis element with a fixed point, a basis whose axes overlap less than the length of either basis element, or a basis whose lengths are not rationally related, and whose axes overlap in $|a| + |b|$ (Propositions 4.3-4.5). If we arrive at a basis element with a fixed point, the possible quotient diagrams are given in Proposition 3.2. If we obtain a basis with short overlap, the possible quotient diagrams are given in Proposition 4.2. If the overlap is $|a| + |b|$ and the slope is irrational, the action is determined in Section 5; it is not simplicial, so is not in $\hat{X}$. We summarize this data below. There are two types of 2-simplices in $\hat{X}$ (which are, in fact, in X), shown in Figure 15. These 2-simplices will be called *tiles* and *fins* respectively.
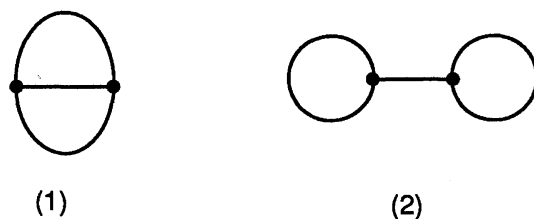

(1)   (2)

Figure 15

There are three types of 1-simplices, listed in Figure 16. These 1-simplices will be called *edges, spikes* and *free edges* respectively.

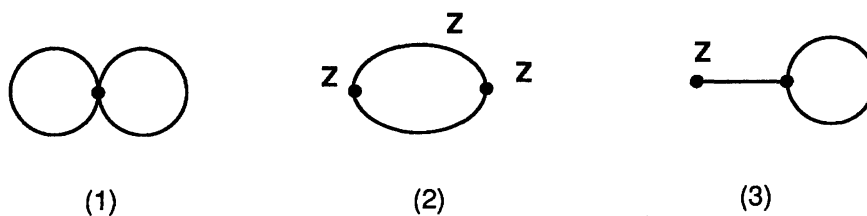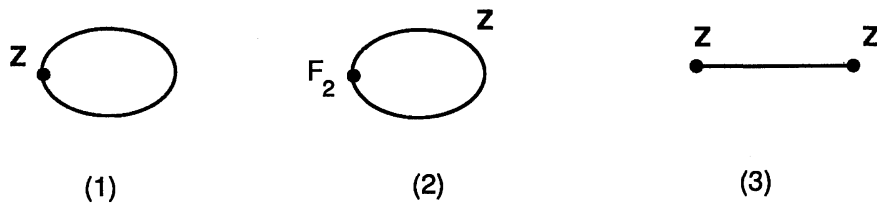(1)   (2)   (3)

Figure 16

Figure 17

There are three types of 0-simplices, shown in Figure 17. The first two occur as endpoints of spikes and will be called *heads* and *tips* respectively Heads are also vertices of tiles and fins. The third type of 0-simplex is a vertex of a fin but not of a tile; it will be called a *peak*.

We will refer to either open or closed simplices by these names; it will be clear from the context which is meant.

There is a natural right action of $Out(F_2)$ on $\mathbf{P}^C$ by permuting the co ordinates. The action on length functions is given as follows. For a lengtl function $l: C \to \mathbb{R}$ and $\phi \in Out(F_2)$ we have $l\phi([g]) = l([\phi(g)])$. This pre serves homothety classes and hence descends to an action on the projectiv length functions. This action permutes the simplices of $\hat{X}$, preserving type and maps each simplex to its image by a projective linear transformation

We identify the abelianization of $F_2 = \langle a, b \rangle$ with the free abelian grou $\mathbb{Z} \oplus \mathbb{Z}$, using the basis of the abelianization consisting of the images of a and b under the canonical epimorphism. We will identify $Out(F_2)$ with $GL_2(\mathbb{Z}$ by identifying an outer automorphism with the matrix of its induced iso morphism of $\mathbb{Z} \oplus \mathbb{Z}$. There is a right action of $GL_2(\mathbb{Z})$ on the projectiv line by right multiplication on row vectors. Thus our identifications defin a right action of $Out(F_2)$ on the projective line.

We are now ready to describe our finite dimensional embedding of $\overline{X}$ We begin by considering the map $f_1: \overline{X} \to \mathbf{P}^2$ given by

$$f_1(\alpha) = (|ab| - |ab^{-1}| : |a| - |b| : |a| + |b| + |aba^{-1}b^{-1}|),$$

where $| \cdot |$ denotes the translation length function for α. The map f_1 i continuous since the coordinates are linear combinations of the standar coordinates on $\mathbf{P}^C$.

Lemma 6.1. $|a| + |b| + |aba^{-1}b^{-1}| > 0$ *for all actions in* $\overline{X}$.

Proof. If $|a| = |b| = 0$, then by Proposition 3.2, Case 1, the action is simplicial, with quotient an interval whose endpoints are stabilized by a and b respectively. For this action, $|aba^{-1}b^{-1}|$ is equal to four times the length of the interval, which is positive since the action is non-trivial.   □

Thus the image of f_1 lies in the affine plane $\{(x : y : 1)\}$ in $\mathbf{P}^2$.

We define an action in $\overline{X}$ to be *special* if the characteristic sets for a and b are disjoint or meet in a point, or if a and b are hyperbolic with $\Delta(a, b) < \min(|a|, |b|)$. If an action is not special it is said to be *generic*.

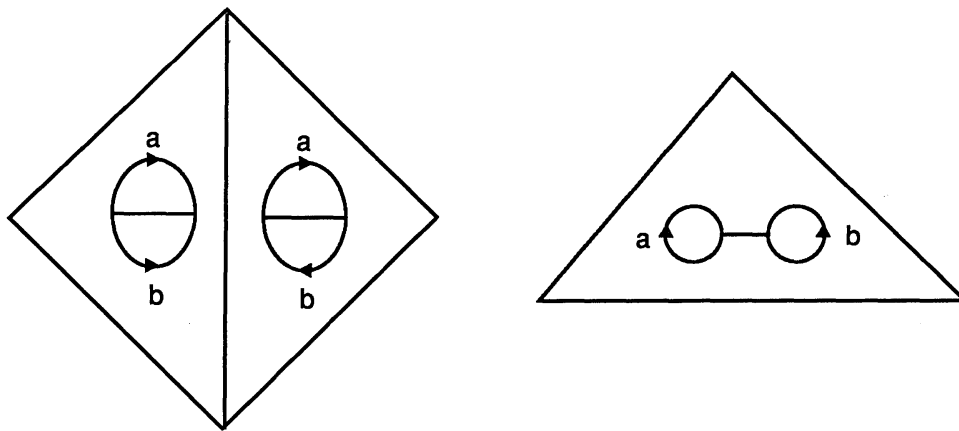**Lemma 6.2.** *The special actions are the two open tiles and the closure of the fin shown in Figure 18:*

Figure 18

Proof. This follows from the classification of simplicial actions in sections 3 and 4.   □

A straightforward computation gives the image of the special actions under the map f_1 . For example, consider the points α_1 and α_2 shown in
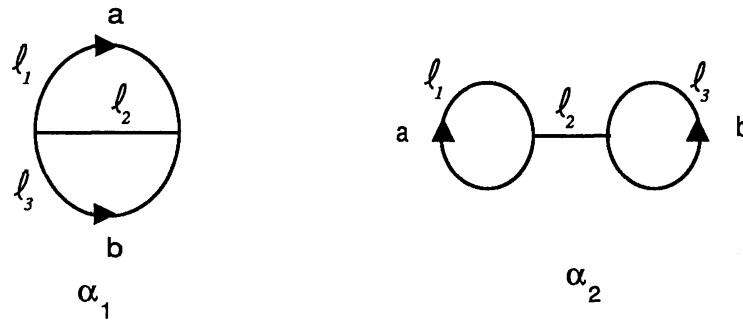
Figure 19

Figure 19. The lengths of the relevant group elements are given in the table below:

	$|a|$	$|b|$	$|ab|$	$|ab^{-1}|$	$|aba^{-1}b^{-1}|$
α_1	$\ell_1 + \ell_2$	$\ell_2 + \ell_3$	$\ell_1 + 2\ell_2 + \ell_3$	$\ell_1 + \ell_3$	$2\ell_1 + 2\ell_2 + 2\ell_3$
α_2	ℓ_1	ℓ_3	$\ell_1 + 2\ell_2 + \ell_3$	$\ell_1 + 2\ell_2 + \ell_3$	$2\ell_1 + 4\ell_2 + 2\ell_3$

Thus,

$$f_1(\alpha_1) = (2\ell_2 : \ell_1 - \ell_3 : 3\ell_1 + 4\ell_2 + 3\ell_3) =$$

$$\left(\frac{2\ell_2}{3\ell_1 + 4\ell_2 + 3\ell_3} : \frac{\ell_1 - \ell_3}{3\ell_1 + 4\ell_2 + 3\ell_3} : 1 \right)$$

and

$$f_1(\alpha_2) = (0 : \ell_1 - \ell_3 : 3\ell_1 + 4\ell_2 + 3\ell_3) = \left(0 : \frac{\ell_1 - \ell_3}{3\ell_1 + 4\ell_2 + 3\ell_3} : 1 \right).$$

From these computations, we obtain the following lemma.

Lemma 6.3. *The map f_1 gives a homeomorphism of the union of the two special tiles and the open face they have in common onto the interior of the quadrilateral in $\mathbf{P}^2$ with vertices $(-1/2 : 0 : 1), (0 : 1/3 : 1), (1/2 : 0 : 1)$ and $(0 : -1/3 : 1)$.*   □

Consider the subset of $\overline{X}$ consisting of generic actions for which a and b are hyperbolic. This has a partition into four subsets X_I, X_{II}, X_{III}, and X_{IV} according to the following scheme: $X_I \cup X_{IV}$ is the of the set of generic actions for which a and b are hyperbolic and their translation

directions agree on the overlap of their axes, and $X_I \cup X_{II}$ is the set of generic actions for which $|a| > |b| > 0$. The set of generic actions is the union of the closures of these four sets, which we denote by $\overline{X}_I, \overline{X}_{II}, \overline{X}_{III}$ and $\overline{X}_{IV}$.

Lemma 6.4. *The sets $\overline{X}_I, \overline{X}_{II}, \overline{X}_{III}$ and $\overline{X}_{IV}$ map into quadrants $I, II,$ III and IV respectively under f_1.*   □

Let τ_1 and τ_2 be the two involutions in $Out(F_2)$ given by

$$\tau_1(a) = b \qquad \tau_1(b) = a$$

and

$$\tau_2(a) = a \qquad \tau_2(b) = b^{-1}.$$

Let σ_1 and σ_2 be the reflections of the affine plane $\{(x : y : 1)\}$ about the x- and y-axes respectively. Then for any action α,

$$f_1(\alpha \cdot \tau_1) = \sigma_1 \circ f_1(\alpha) \text{ and } f_1(\alpha \cdot \tau_2) = \sigma_2 \circ f_1(\alpha).$$

Thus in order to describe the map f_1 it suffices to consider its restriction to the closed set $\overline{X}_I$; the behavior on the other sets can be determined by symmetry.

For generic actions, the first coordinate of f_1 can be expressed in terms of the lengths of the generators a and b:

Lemma 6.5. *For a generic action, $|ab| - |ab^{-1}| = \pm 2\min(|a|, |b|)$, where the sign is positive if and only if the action is in $\overline{X}_I \cup \overline{X}_{IV}$.*

Proof. As an example, consider the case where $|a| > |b| > 0$ and the translation directions of a and b agree on the overlap of their axes (i.e. the action is in X_I). Since the action is generic, $|b| < \Delta(a, b)$; this is illustrated in Figure 9 of Proposition 4.1. One checks that $|ab^{-1}| = |a| - |b|$ and $|ab| = |a| + |b|$, giving the lemma in this case. The other cases are similar. □

Lemma 6.6. *For any action α, $f_1(\alpha) = (x : y : 1)$ with $|x| + |y| \leq 1$. Equality holds if and only if $|aba^{-1}b^{-1}| = 0$.*

Proof. This is an easy calculation for actions in $\overline{X}_I$. The lemma follows by symmetry together with Lemma 6.3.   □

For any generic action α, recall that the *slope* of α is defined to be $\pm |a|/|b|$, where the sign is positive if and only if α is in $\overline{X}_I \cup \overline{X}_{IV}$.

Lemma 6.7. *For a generic action α of slope m, we have $f_1(\alpha) = (x : y : z)$, where*

$$\frac{y}{x} = \begin{cases} \frac{1}{2}(m-1) & \text{if } \alpha \in \overline{X}_I \\ \frac{1}{2}(1-\frac{1}{m}) & \text{if } \alpha \in \overline{X}_{IV} \\ \frac{1}{2}(-m-1) & \text{if } \alpha \in \overline{X}_{II} \\ \frac{1}{2}(1+\frac{1}{m}) & \text{if } \alpha \in \overline{X}_{III} \end{cases}$$

Proof. If $\alpha \in \overline{X}_I$, then $|a| \geq |b| \geq 0$, and if $|b| > 0$, the translation directions for a and b agree on their overlap (which is an arc since α is generic). Thus by Lemma 6.5, $|ab| - |ab^{-1}| = 2|b|$, so

$$\frac{y}{x} = \frac{|a| - |b|}{|ab| - |ab^{-1}|} = \frac{|a| - |b|}{2|b|} = \frac{1}{2}(m-1).$$

Computations in the other quadrants are similar.   □

By Propositions 4.3 and 5.3, the actions with $|aba^{-1}b^{-1}| = 0$ are determined by the slope m. Thus Lemma 6.7 shows that f_1 is injective on this set and maps it to the boundary of the diamond

$$D = \{(x : y : 1) : |x| + |y| = 1\}$$

By Proposition 4.3, every action on a given spike has the same (rational) slope and distinct spikes have distinct slopes. Therefore, Lemma 6.7 also shows that distinct spikes have disjoint images under f_1. We calculate the image explicitly below.

Lemma 6.8. *The spike of slope $m = p/q$, $(p,q) = 1$, is mapped by f_1 to the segment from $(2q : p - q : p + q + 2)$ to $(2q : p - q : p + q)$.*

Proof. Again, we give the computation for a spike in $\overline{X}_I$, where $|a| \geq |b| \geq 0$ and the translation directions of a and b agree if b is hyperbolic.

Normalize so that the quotient graph of an action on the spike has length one (which doesn't change the point in projective space). Then $|a|$ and $|b|$ are integers, and are relatively prime since $\{a, b\}$ is a basis for F_2. By Lemma 6.7, we have $f_1(\alpha) = (x : y : z)$, with

$$\frac{y}{x} = \frac{|a| - |b|}{2|b|} = \frac{1}{2}(m-1) = \frac{(p-q)}{2q}.$$

Since p and q are also relatively prime, this implies $|b| = q$ and $|a| = p$.

Now recall that, in any action, the commutator of any basis is conjugate to $aba^{-1}b^{-1}$ or its inverse, so has the same length. At the tip of the spike, the commutator has length zero, and at the head, the commutator has length 2. Thus $|a| = p, |b| = q, |ab| - |ab|^{-1} = 2q$, and $0 \le |aba^{-1}b^{-1}| \le 2$. This gives the lemma for spikes in $\overline{X}_I$, and the other cases are similar. $\square$

The vertices of a closed tile are heads of spikes with slopes p/q, r/s, and $(p + q)/(r + s)$ where $ps - qr = \pm 1$. This allows us to compute the image of each of these 2-simplices, giving the following simple geometric construction.

Lemma 6.9. *Suppose that the spikes of slopes p/q and r/s, with $ps - qr = 1$, are both contained in one of the sets $\overline{X}_I$, $\overline{X}_{II}$, $\overline{X}_{III}$, or $\overline{X}_{IV}$. Consider the two diagonal lines joining the image under f_1 of the head of one spike to the image of the tip of the other. Then f_1 maps the head of the spike of slope $(p + q)/(r + s)$ to the point of intersection of these lines.* $\square$
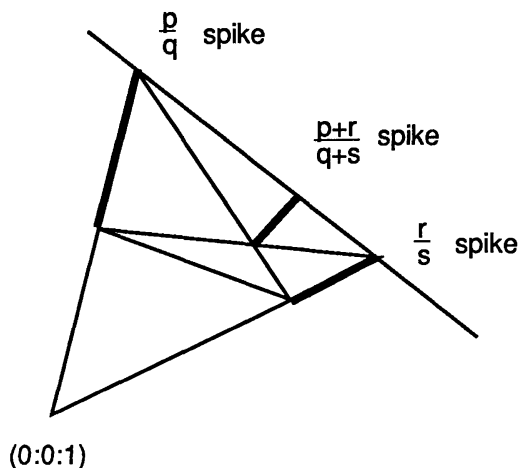
This construction is illustrated in Figure 20.

Figure 20

Let $\overline{Y}$ denote the subspace of $\overline{X}$ consisting of actions for which no pair of basis elements has disjoint characteristic sets. This is the complement of the open fins, free edges and peaks.

Proposition 6.10. *The map f_1 restricted to $\overline{Y}$ is a homeomorphism onto the diamond $D = \{(x : y : 1): |x| + |y| = 1\}$.*

Proof. We first prove injectivity. Inductively define the *level* of a closed tile as follows. The two special tiles have level 0. If a tile has an edge in common with a tile of level n and is not of level less than or equal to n then we define it to have level $n + 1$. Define the level of a closed spike to be the minimum level of the tiles with which it has a common vertex. Let L_n denote the union of the tiles and spikes of level less than or equal to n. We show by induction on n that the restriction of f_1 to L_n is a homeomorphism onto its image. For $n = 0$ this follows from Lemma 6.2. The induction step follows from the geometric construction given in Lemma 6.9; the image of each tile of level n meets $f_1(L_n)$ in one edge and the complement of this edge is disjoint from the images of the other tiles of level n. Each tile of level $n > 0$ meets exactly one spike of level n in one vertex and the construction in Lemma 6.9 shows that the image of this spike is disjoint from the images of the other tiles. Since we have seen that f_1 is injective on spikes, this shows that f_1 is injective on $\overline{Y}$.

To prove that f_1 maps $\overline{Y}$ onto the diamond, consider a ray emanating from the origin in the plane $\{((x : y : 1)\}$. Let P be the last point on this ray which is contained in $f_1(\overline{Y})$. Suppose that P is not on the boundary of D. Since the image of a spike is contained in a ray from the origin, P cannot be in the image of a spike. Thus it is the image of a non-vertex point in the boundary of a closed tile. But this tile adjoins another tile whose image, by the construction in Lemma 6.9, covers points further out along the ray. This contradiction shows that P is on the boundary of D and that $f_1(\overline{Y}) = D$.   □

A picture of $f_1(\overline{Y})$ appears as Figure 21 at the end of the paper. It is not difficult to give a complete description of the induced action of $Out(F_2)$ on $f_1(\overline{Y})$. The group $Out(F_2)$ is generated by the involutions τ_1 and τ_2 together with the outer automorphism π given by

$$\pi(a) = ab \qquad \pi(b) = b.$$

We have already described the action of the involutions. For a generic tile t, the action of π on $f_1(t)$ is by the unique projective linear transformation which fixes the origin and maps $f_1(t)$ to $f_1(\pi(t))$. The action on the images of the special tiles can be computed directly.

Next we will extend f_1 to an embedding $f\colon \overline{X} \to \mathbf{P}^3$. Since this involves extending f_1 over the closed fins, it will be convenient to have a way of indexing the fins. Since the actions in the interior of a fin have quotients which are barbells (see Figure 11), the two circles of the barbell determine a conjugacy class of unordered bases of F_2 up to an equivalence relation in which the conjugacy class of $\{u, v\}$ is equivalent to that of $\{u, v^{-1}\}$. Denote the equivalence class of the basis $\{u, v\}$ by (u, v). The fins are indexed by the classes (u, v) where u and v are a basis of F_2.

Given a basis $\{u, v\}$ of F_2 and an action α of F_2 on an $\mathbb{R}$-tree T, we define

$$b_{(u,v)}(\alpha) = \max(|uv|, |uv^{-1}|) - |u| - |v|.$$

Lemma 6.11. *Let α, T and $b_{(u,v)}$ be as above. Then $b_{(u,v)}(\alpha)$ is equal to the distance between the characteristic sets for u and for v.* $\square$

It is an easy consequence of the lemma that $b_{(u,v)}$ is non-zero only on the fin which corresponds to actions described by the barbell with labels "u" and "v", and is zero on the actions corresponding to the edge along which this fin meets $\overline{Y}$. In particular, a fin minus its intersection with $\overline{Y}$ is an open set in $\overline{X}$.

Define, for any action α corresponding to a point of $\overline{X}$,

$$w(\alpha) = \sum_{(u,v)} \frac{b_{(u,v)}(\alpha)}{\ell(u) + \ell(v)},$$

where for $g \in F_2$, $\ell(g)$ denotes the length of a cyclic reduced word in a and b representing the conjugacy class of g. We now define a map $f_2\colon \overline{X} \to \mathbf{P}^3$ by

$$f_2(\alpha) = (|ab| - |ab^{-1}| : |a| - |b| : |a| + |b| + |aba^{-1}b^{-1}| : w(\alpha)).$$

Since each coordinate of f_2 transforms linearly under scaling of the metric this gives a well-defined map into projective space. This will be a first approximation to our embedding of $\overline{X}$ in $\mathbf{P}^3$, but it is not quite injective and must be perturbed slightly to construct the embedding. Before constructing the embedding we pause to observe:

Lemma 6.12. *The map $f_2 \colon \overline{X} \to \mathbf{P}^3$ is continuous.*

Proof. It suffices to check that

$$\frac{w(\alpha)}{|a| + |b| + |aba^{-1}b^{-1}|}$$

is continuous on $\overline{X}$. This function is an infinite sum of functions such that at most one is non-zero at any point of $\overline{X}$. Thus in order to force the partial sums to converge uniformly it suffices to make the maximum values of these terms converge to zero. The denominators $\ell(u) + \ell(v)$ were chosen, more or less arbitrarily, to guarantee this.   □

The image of f_2 is contained in the affine space $\{(x : y : 1 : w)\} \subset \mathbf{P}^3$. An easy computation gives

Lemma 6.13. *Consider the fin which contains the heads of slopes p/q and r/s with $ps - qr = \pm 1$. Assume that p/q and r/s are contained in $[1, \infty]$, $p + q < r + s$, and $\{p/q, r/s\} \neq \{1, \infty\}$. Then the fin maps under f_2 to the triangle with vertices $(2q : p - q : p + q + 2 : 0)$, $(2s : r - s : r + s + 2 : 0)$, and $(2q : p - q : p + q + 2 : \frac{1}{p+q+r+s})$.*   □

Note that the fourth coordinate of f_2 is invariant under the involutions τ_1 and τ_2. Thus by symmetry the lemma above describes the images of all of the fins which do not meet the special tiles.

The lemma shows that the image of each fin which does not meet a special tile is orthogonal to the plane $\{(x : y : 1 : 0)\}$ which contains the image of $\overline{Y}$ under f_2. However, it also shows that one edge of such a fin is mapped to a vertical line segment, and hence that there are infinitely many fins whose images all meet a vertical line segment emanating from the image of a head. Thus f_2 is not injective on these edges.

In order to construct our embedding we have to correct this defect in f_2. The map f_2 also has a second defect: the four fins corresponding to the classes $(a, ab), (a, ab^{-1}), (b, ab)$, and (b, ab^{-1}) do not map to vertical triangles. This is easily corrected by redefining the first coordinate on these four fins to be $2\min(|a|, |b|)$. This agrees with the original coordinate on the edges where these fins meet $\overline{Y}$, so the modified map is continuous. Let f_3 be the map obtained from f_2 by this modification.

The map f_3 still has the first defect. To correct this, we perturb the map slightly on the fins as follows. For our embedding f we will use a map

which agrees with f_3 on $\overline{Y}$ and maps each fin to a vertical triangle by a projective linear map. However, if P is the peak of a fin which is mapped by f_3 to a point above a head, then we define f so that $f(P)$ is displaced horizontally from $f_3(P)$ by an amount ϵ_P and lies above an interior point of the image of the edge where the fin meets $\overline{Y}$. This makes f injective. As long as the displacements ϵ_P converge to 0, the map f will be continuous.

Pictures of the image of f, drawn using $Mathematica^{TM}$ and some computer programs written by Curt McMullen, are shown in Figure 22.
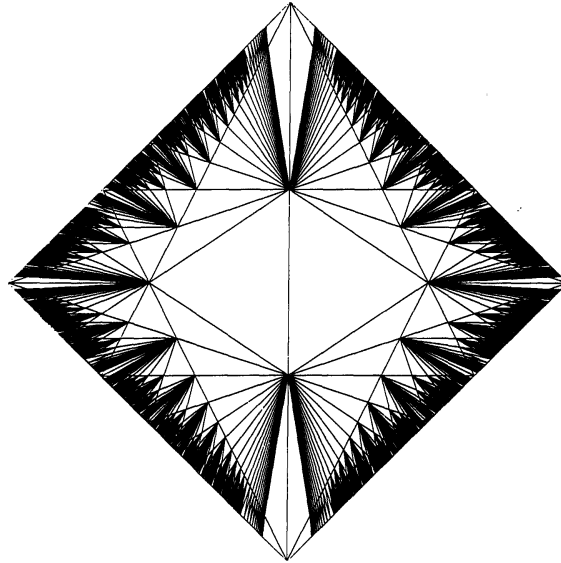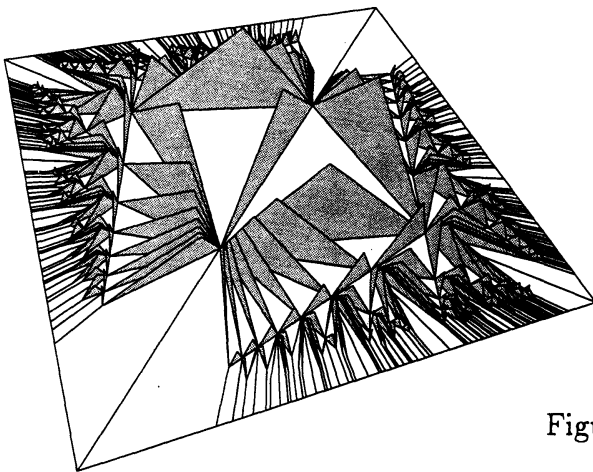

Figure 21. $f_1(\overline{Y})$

Figure 22. $f(\overline{X})$

REFERENCES

1. R. Alperin and H. Bass, *Length Functions of Group Actions on Λ- trees*, in "Combinatorial Group Theory and Topology," S.M. Gersten et al., eds., Princeton University Press, Princeton, 1987, pp. 265–378.

2. M. Cohen and M. Lustig, *Very Small Group Actions on $\mathbb{R}$-trees and Dehn-twist Automorphisms*.

3. M. Culler and J.W. Morgan, *Group Actions on $\mathbb{R}$-trees*, Proc. Lond. Math. Soc. **55** (1987), 571–604.

4. M. Culler and K. Vogtmann, *Moduli of Graphs and Automorphisms of Free Groups*, Invent. Math. **84** (1986), 91–119.

5. H. Gillet and P.B. Shalen, *Dendrology of Groups of Low $\mathbb{Q}$-ranks*, preprint, J. Diff. Geom. (1988) (to appear).

6. A.E. Hatcher, *Measured Lamination Spaces for Surfaces, from the Topological Viewpoint*, Topology and its Applications **30** (1988), 63–88.

7. J.W. Morgan, *Ergodic Theory and Free Actions on Trees*, Invent. Math. **94** (1988), 605–622.

8. J.W. Morgan and J.P. Otal, *Relative Growth Rates of Closed Geodesics on a Surface under Varying Hyperbolic Structures*, Comment. Math. Helv. (to appear).

9. J.W. Morgan and P.B. Shalen, *Degenerations of Hyperbolic Structures, II: Measured Laminations in 3-Manifolds*, Ann. of Math. **127** (1988), 403–456.

10. W. Parry. these proceedings.

11. F. Paulin, *Topologie de Gromov equivariante, Structures Hyperboliques et Arbres Reels*, Invent. Math. **94** (1988), 53–80.

12. J. Plante, *Foliations with Measure-Preserving Holonomy*, Ann. of Math. **102** (1975), 327–361.

13. J.P. Serre, "Trees," Springer-Verlag, Berlin, 1980. Translation of "Arbres, Amalgames, SL_2," Astérisque **46**, 1977.

14. R. Skora, *Splittings of surfaces*, preprint.