ISSN 1472-2739 (on-line) 1472-2747 (printed)

Algebraic & Geometric Topology
Volume 5 (2005) 1481–1503
Published: 1 November 2005

Contractibility of deformation spaces of G–trees

Matt Clay

Abstract Forester has defined spaces of simplicial tree actions for a finitely generated group, called deformation spaces. Culler and Vogtmann's Outer space is an example of a deformation space. Using ideas from Skora's proof of the contractibility of Outer space, we show that under some mild hypotheses deformation spaces are contractible.

AMS Classification 20E08; 20F65, 20F28

Keywords G–tree, deformation space, Outer space

Culler and Vogtmann's Outer space is a good geometric model for $\mathrm{Out}(F_n)$, the outer automorphism group of a finitely generated free group of rank $n \geq 2$, for three reasons:

(1) Outer space is contractible;
(2) point stabilizers are finite; and
(3) there is a equivariant deformation retract on which the action is cocompact [4].

Outer space is the analog of Teichmüller space for the mapping class group of a closed negatively curved surface or of the symmetric space for an arithmetic group. See [1] and [12] for a survey of some results about $\mathrm{Out}(F_n)$ obtained from using this connection between the three classes of groups. Also see [2] for some open questions about the similarities and differences.

Recall that Outer space is the moduli space of free actions of a free group on a simplicial tree. Forester has defined a generalization of Outer space for an arbitrary finitely generated group G [5]. The generalization allows actions which are not free but requires the subgroups with fixed points to be the same among all actions in the moduli space. Unfortunately these spaces are not $\mathrm{Out}(G)$–invariant in general. Nevertheless, in the cases when the space is invariant under $\mathrm{Out}(G)$ these spaces have the potential to provide information about the structure of $\mathrm{Out}(G)$. The purpose of this paper is to show that these spaces share the first of the above mentioned properties with Outer space, i.e. they are contractible.

For a finitely generated group G, a *G–tree* is a metric simplicial tree on which G acts by isometries. Two G–trees T and T' are equivalent if there is a G–equivariant isometry between then. When we speak of a G–tree we will always mean the equivalence class of the G–tree. A subgroup is called an *elliptic* subgroup for T if it has a fixed point in T. Given a G–tree there are two moves one can perform to the tree that do not change whether or not subgroups of G are elliptic. These moves correspond to the isomorphism $A \cong A *_C C$ and are called *collapse* and *expansion*. For a detailed description of the moves see [5]. In [5] Forester proves the converse, namely if two cocompact G–trees have the same elliptic subgroups, then there is a finite sequence of collapses and expansions (called an *elementary deformation*) transforming one G–tree to the other. A G–tree T is *cocompact* if the quotient T/G is a finite graph.

We let $\mathcal{X}$ denote a maximal set of cocompact G–trees which are related by an elementary deformation. By the theorem of Forester mentioned above, an equivalent definition is as the set of all cocompact G–trees that have the same elliptic subgroups as some fixed G–tree. Both of these interpretations are utilized in the following. This set $\mathcal{X}$ is called a *unnormalized deformation space.* We will always assume that the G–trees are minimal, irreducible and that G acts without inversions. See section 1 for these definitions.

As is common practice in spaces of this nature, we projectivize by taking the quotient of $\mathcal{X}$ under the action of $\mathbb{R}^+$ by homothety. The quotient $\mathcal{X}/\mathbb{R}^+$ is called a *deformation space* and is denoted $\mathcal{D}$. Outer space is an example of a deformation space for a finitely generated free group where the only elliptic subgroup is the trivial group. Culler and Vogtmann described a contraction of the spine of Outer space using combinatorial methods and a "Morse-like" function [4]. Skora showed in a different manner that Outer space is contractible [11]. The method of Skora is to homotope the unnormalized deformation space projecting to Outer space to a set homeomorphic to a simplex$\times\mathbb{R}^+$ by continuously unfolding G–trees in the unnormalized deformation space. This homotopy descends to Outer space, proving its contractibility. It is this idea which we extend to show:

Theorem 6.7 *For a finitely generated group G, any irreducible deformation space which contains a G–tree with finitely generated vertex groups is contractible.*

The outline of the proof is as follows: starting with an unnormalized deformation space $\mathcal{X}$, we look at the space $\mathcal{M}(\mathcal{X})$ of morphisms between elements of $\mathcal{X}$. A *morphism* is a G–equivariant map between G–trees which on each segment either folds or is an isometry. Given a morphism $\phi\colon T \to Y$ we show that we can continuously interpolate between the two G–trees. We then fix some reduced G–tree $T \in \mathcal{X}$. For another G–tree $Y \in \mathcal{X}$ we define a map $B(Y)\colon T \to Y$, which is not a morphism but is nice in certain respects. The assignment $Y \mapsto B(Y)$ is a continuous function between the appropriate spaces. We redefine the metric on T to obtain another G–tree T_Y (equivariantly homeomorphic to T) such that $B(T)\colon T_Y \to Y$ is a morphism. Thus we can homotope $\mathcal{X}$ to the space of trees equivariantly homeomorphic to T. We show this space is homeomorphic to a simplex$\times\mathbb{R}^+$, thus $\mathcal{X}$ is contractible. This homotopy descends to a contraction of the deformation space $\mathcal{D}$.

Originally, the following proof was only for finitely generated generalized Baumslag–Solitar groups, for which there is a natural $\mathrm{Out}(G)$–invariant deformation space. A *generalized Baumslag–Solitar group* is a group which admits an action on a simplicial tree where the stabilizer of any point is isomorphic to $\mathbb{Z}$. However after a research announcement by Guirardel and Levitt [7], which contains Theorem 6.7, we noticed that our proof for generalized Baumslag–Solitar groups went through in the general case after modifying case (ii) in Lemma 6.4. We are grateful for their announcement. They have proven Theorem 6.7 in the case of a free product and have given several consequences [8].

The majority of material presented within is in Skora's preprint [11]. As this preprint was never published, we present the full details here. The main difference from [11] is section 6.

Acknowledgements This work was done under the supervision of my advisor Mladen Bestvina. In addition to thanking him for the helpful discussions, I am also grateful for discussions with Lars Louder and for the research announcement of Vincent Guirardel and Gilbert Levitt. Thanks are also due to the referee for suggestions improving the exposition.

1 Preliminaries

For a G–tree T, the length function $l_T\colon G \to [0,\infty)$ is defined by $l_T(g) = \min_{x\in T} d(x,gx)$. The *characteristic set* T_g, of a element $g \in G$ is where this minimum is realized, i.e. $T_g = \{x \in T \mid d(x,gx) = l_T(g)\}$. An element is *elliptic* if $l_T(g) = 0$ and *hyperbolic* otherwise. For $g \in G$ hyperbolic, the characteristic set is isometric to $\mathbb{R}$ and g acts on T_g by translation by $l_T(g)$. In this case the characteristic set of g is often called the *axis* of g. Note that $d(x,gx) = 2d(x,T_g) + l_T(g)$ for both g elliptic or g hyperbolic. If a subgroup

$H \subseteq G$ is elliptic, we define the characteristic set of H as $T_H = \{x \in T \mid hx = x \; \forall h \in H\}$. For a closed set $A \subseteq T$ we let $p_A \colon T \to A$ denote the nearest point projection. A map between metric simplicial trees $\phi \colon T \to T'$ is morphism if for any segment $[x,y] \subseteq T$ there is a subsegment $[x,x'] \subseteq [x,y]$ on which ϕ is an isometry. If T and T' are G–trees, we also require that ϕ is G–equivariant.

We have following dictionary of group actions on trees [3]. A G–tree is *trivial* if there is a fixed point and *minimal* if there is no proper invariant subtree. A G–tree T is *reducible* if:

(1) every element fixes a point (equivalent to being trivial for finitely generated groups); or

(2) G fixes exactly one end of T; or

(3) G leaves a set of two ends of T invariant.

If T is not reducible, it is *irreducible*. A G–tree is irreducible if and only if there are two hyperbolic elements whose axes are either disjoint or intersect in a compact set [3]. This feature is preserved by elementary deformations [5], hence any G–tree obtained via an elementary deformation from an irreducible G–tree is also irreducible.

Unless otherwise stated, we will always assume G–trees are minimal and irreducible. We say a deformation space is *irreducible* if every G–tree in the space is irreducible. By the above statement, a deformation space is irreducible if any G–tree in the space is irreducible.

2 Topology on deformation spaces

We endow an unnormalized deformation space $\mathcal{X}$ with the *Hausdorff–Gromov* topology. Gromov introduced this topology as a way to compare two distinct metric spaces [6]. This topology generalizes the Hausdorff distance between two closed sets in a metric space. The deformation space $\mathcal{D}$ is then topologized as the quotient $\mathcal{X}/\mathbb{R}^+$.

The Hausdorff–Gromov topology is defined as follows. Let X, Y be metric G–spaces, i.e. metric spaces equipped with isometric G–actions. For any $\epsilon > 0$, an *ϵ–approximation* is a set $R \subseteq X \times Y$ that surjects onto each factor such that if $x, x' \in X$ and $y, y' \in Y$ with xRy (i.e. $(x,y) \in R$) and $x'Ry'$ then $|d(x,x') - d(y,y')| < \epsilon$. We say that R is a *closed* ϵ–approximation if R is closed in $X \times Y$. For a finite subset $P \subseteq G$ and subspaces $K \subseteq X, L \subseteq Y$ the

ϵ–approximation in $K \times L$ is P*–equivariant* if whenever $g \in P$, $x, gx \in K$ and $y \in L$ with xRy then $gy \in L$ and $gxRgy$.

Given an ϵ–approximation $R \subseteq X \times Y$, we let R_δ denote the closed δ–neighborhood of R using the L^1 metric. In other words $R_\delta = \{(x,y) \in X \times Y \mid \exists (x',y') \in R \text{ with } d(x,x') + d(y,y') \leq \delta\}$. One can show that R_δ is a $(\epsilon + 2\delta)$–approximation. If R is P–equivariant, then R_δ is P–equivariant.

These ϵ–approximations can topologize any set of metric G–spaces. In particular, they can topologize any unnormalized deformation space $\mathcal{X}$. Let $\mathcal{S}$ be such a set of metric G–spaces. Then for $X \in \mathcal{S}$, $K \subseteq X$ compact, $P \subseteq G$ finite and $\epsilon > 0$ define a basic open set $U(X,K,P,\epsilon)$ to be the set of all $Y \in \mathcal{S}$ such that there is a compact set $L \subseteq Y$ and a P–equivariant closed ϵ–approximation $R \subseteq K \times L$. If $K \subseteq K'$ and $P \subseteq P'$ then $U(X,K',P',\epsilon) \subseteq U(X,K,P,\epsilon)$. This will allow us to assume that certain subsets of X and G are contained in K and P respectively by shrinking our basic open set.

Given an ϵ–approximation $R \subseteq X \times Y$, we will assume it is *full*: i.e. if xRy and $x'Ry'$ then every point in $[x,x']$ is related by R to some point in $[y,y']$ and vice versa. This is not necessary but it cleans up some of the proofs in sections 5 and 6. When the set $\mathcal{S}$ contains only trees the two topologies generated are the same. For X, P, ϵ as above let $U_f(X,K,P,\epsilon)$ be the set of all $Y \in \mathcal{S}$ such that there is a finite subtree $L \subseteq Y$ and a P–equivariant closed full ϵ–approximation $R \subseteq K \times L$. Clearly we have $U_f(X,K,P,\epsilon) \subseteq U(X,K,P,\epsilon)$. We now show the opposite inclusion of bases.

For two trees X, Y, subsets $K \subseteq X, L \subseteq Y$ related by an ϵ–approximation $R \subseteq K \times L$ and a finite segment $[x_1,x_2] \subseteq K$, let $R([x_1,x_2]) = \{z \in L \mid \exists x \in [x_1,x_2] \text{ with } xRz\}$. For $z \in R([x_1,x_2])$ with x_iRy_i for some $y_i \in Y, i = 1,2$ we have $d(z,[y_1,y_2]) < 2\epsilon$. We have the following statement about the density of $R([x_1,x_2])$.

Lemma 2.1 *If $z_0 \in [y_1,y_2] \subseteq Y$ where x_iRy_i for $i = 1,2$, then there is a $z \in R([x_1,x_2])$ such that $d(z_0,z) < 2\epsilon$.*

Proof We assume this is not the case. Let $d(y_1,z_0) = d_1, d(y_2,z_0) = d_2$. As $y_1, y_2 \in R([x_1,x_2])$ we can assume both d_1 and d_2 are larger than ϵ. Take $x \in [x_1,x_2]$ such that $d(x_1,x) = d_1$. Therefore $d(x_2,x) < d_2 + \epsilon$. There is a $z \in R([x_1,x_2])$ such that xRz. For this z, $d(y_1,z) < d_1 + 2\epsilon$ and $d(y_2,z) < d_2 + 2\epsilon$. Now we let $z' = p_{[y_1,y_2]}(z)$. Hence by our initial assumption $d(z_0,z') + d(z,z') \geq d(z_0,z) \geq 2\epsilon$. Assume without loss of generality that z' is closer to y_1 than z_0 is. Then $d(y_2,z) = d(y_2,z_0) + d(z_0,z') + d(z',z) \geq d_2 + 2\epsilon$, a contradiction. □

To finish up the claim that the two above mentioned topologies are the same we let $\delta = \frac{\epsilon}{5}$. Then for $Y \in U(X, K, P, \delta)$ we have a P–equivariant δ–approximation between K and some finite subtree $L \subseteq Y$. By the above $R_{2\delta}$ is a full P–equivariant ϵ–approximation between K and L. Therefore $U(X, K, P, \delta) \subseteq U_f(X, K, P, \epsilon)$ and the two topologies are indeed the same.

We will also topologize the space of morphisms between elements in a deformation space. Let $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ be G–equivariant maps, a closed ϵ–approximation between these two maps is a pair (R, R') such that:

(1) $R \subseteq X \times Y$ and $R' \subseteq X' \times Y'$ are closed ϵ–approximations; and

(2) for $x \in X, y \in Y$ if xRy then $\phi(x)R'\psi(y)$.

Let $P \subseteq G$ be finite and $K \subseteq X, K' \subseteq X', L \subseteq Y, L' \subseteq Y'$ be subspaces. The ϵ–approximation (R, R') is P–equivariant if R and R' are P–equivariant on the appropriate subspaces. Note that if $z \in \text{graph}(\phi)$ then by definition there is a $w \in \text{graph}(\psi)$ with $z(R, R')w$.

As above this allows us to topologize a set of G–equivariant maps between G–spaces. In particular we can topologize $\mathcal{M}(\mathcal{X})$, the set of morphisms between elements of $\mathcal{X}$. Let $\mathcal{S}'$ a set of G–equivariant maps between G–spaces. For $\phi\colon X \to X'$ in $\mathcal{S}'$, and $K \subseteq X, K' \subseteq X'$ both compact with $\phi(K) \subseteq K'$, $P \subseteq G$ finite and $\epsilon > 0$ define the basic open set $U(\phi, K \times K', P, \epsilon)$ to be the set of all maps $\psi\colon Y \to Y'$ in $\mathcal{S}'$ such that there are compact sets $L \subseteq Y$, $L' \subseteq Y'$ with $\psi(L) \subseteq L'$ and a P–equivariant closed ϵ–approximation (R, R') bewtween $\phi\colon K \to K'$ and $\psi\colon L \to L'$.

For a space $\mathcal{S}$ of metric G–spaces and a space $\mathcal{S}'$ of G–equivariant maps between the elements of $\mathcal{S}$ we have the two continuous maps $\mathcal{D}o$ and $\mathcal{R}a$ defined from $\mathcal{S}'$ to $\mathcal{S}$ which send a map to its domain and range respectively. In other words, for $\phi\colon X \to X'$ an element of S' we have $\mathcal{D}o(\phi) = X$ and $\mathcal{R}a(\phi) = X'$.

There are two other topologies one might use to topologize a deformation space. Let $\mathcal{C}$ be the set of conjugacy classes for G. Then we have a function $l\colon \mathcal{X} \to \mathbb{R}^{\mathcal{C}}$ where the coordinates are given by the length functions $l_T(c)$ where $c \in \mathcal{C}$. Culler and Morgan showed that for minimal irreducible actions on $\mathbb{R}$-trees this function is injective [3]. This defines a topology on $\mathcal{X}$ (and hence on $\mathcal{D}$) called the *axes* topology. Paulin proved that for spaces of minimal irreducible actions on $\mathbb{R}$-trees, the Hausdorff–Gromov topology is the same as the axes topology [10].

We can define the *weak* topology directly on $\mathcal{D}$. The *volume* of a G–tree T, denoted $\text{vol}(T)$, is the sum of the lengths of the unoriented edges of T/G. We identify $\mathcal{D}$ with the G–trees in $\mathcal{X}$ that have volume one. By reassigning the lengths of the edges of T/G in a manner to hold the volume constant we can define a simplex in $\mathcal{D}$. The weak topology is defined by considering $\mathcal{D}$ as the union of such simplicies. In general, the weak topology is different from the axes and Hausdorff–Gromov topology, see [9] for an example.

3 Deforming trees

A morphism $\phi\colon T \to T'$ between trees in an unnormalized deformation space $\mathcal{X}$ can be decomposed into elementary deformations [5]. We will define trees T_t which continuously interpolate between T and T'.

For the morphism $\phi\colon T \to T'$, a nontrivial segment $[x,x'] \subseteq T$ is *folded* if $\phi(x) = \phi(x')$. A folded segment is *maximally folded* if it cannot be locally extended to a segment which is folded. On a maximally folded segment $[x,x']$ the function $d(\phi(z),\phi(x))$ attains a local maximum at possibly several points. Such points are called *fold points* of the morphism ϕ. The points at where the global maxima are obtained are called *maximal fold points.* We remark that every fold point is a maximal fold point for some maximally folded segment. A fold point z is *d-deep* if $d(\phi(x),\phi(z)) > d$ for some maximally folded segment $[x,x']$ of which z is a maximal fold point.

We let $m(\phi) = \sup\{d(\phi(z),\phi(x)) \mid z \in [x,x'] \text{ where } \phi(x) = \phi(x')\}$. Then $m(\phi)$ is finite as elementary deformations are quasi-isometries [5]. Notice that $m(\phi) = 0$ if and only if ϕ is an isometry and hence $T = T'$ as G–trees. For $0 \le t \le 1$ we define $V_t = \{(x,y) \in T \times T' \mid d(\phi(x),y) \le m(\phi)t\}$. For $(x,y) \in V_t$ let $C_t(x,y)$ denote the path component of $V_t \cap (T \times \{y\})$ which contains (x,y). Finally, we define:

$$W_t = \{(x,y) \in V_t \mid C_t(x,y) \cap \operatorname{graph}(\phi) \neq \emptyset\}.$$

Thus W_t is a thickening of $\operatorname{graph}(\phi) \subseteq T \times T'$. We will write $W_t(\phi)$ when we need to specify the morphism. Let $\mathcal{F}_t$ be a partition of W_t into sets which are the path components of $W_t \cap (T \times \{y\})$ for $y \in T$ and $T_t = W_t/\mathcal{F}_t$. We denote points in T_t by $[z]_t$ for $z \in W_t$. As $\mathcal{F}_t$ is G–equivariant, T_t is a G–tree. For $0 \le s \le t \le 1$ the inclusions $W_s \to W_t$ induce G–equivariant maps $\phi_{st}\colon T_s \to T_t$, Figure 1.

A path $\gamma : [0,1] \to W_t$ is *taut* if for components A in $\mathcal{F}_t$, $\gamma^{-1}(A)$ is connected. For $t > 0$, a non-backtracking path γ in T_t lifts to a path $\tilde{\gamma}$ in W_t with endpoints in $\operatorname{graph}(\phi)$. This lift $\tilde{\gamma}$ is homotopic relative to these endpoints to

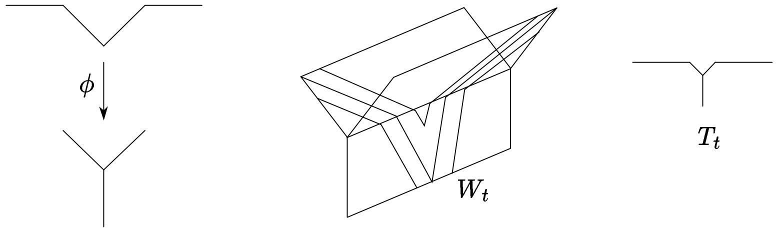

Figure 1: W_t and T_t for the morphism on the left

a taut product of paths $\gamma_1 \cdots \gamma_k$ where each γ_i lies either in a component of $\mathcal{F}_t$ or is a non-backtracking path in graph(ϕ). For $t = 0$ a non-backtracking path γ in T_0 lifts to a path $\tilde{\gamma}$ which is homotopic relative to its endpoints to a taut product of paths $\gamma_1 \cdots \gamma_k$ where ϕ is an isometry on each γ_i. We call these decompositions *taut corner paths*, the pieces lying in graph(ϕ) are called *essential*, the pieces lying in some component of $\mathcal{F}_t$ are called *nonessential*, see Figure 2. Metrize T_t by setting length(γ) equal to the sum of the lengths of the essential pieces measured in T' (or equivalently measured in T). With this metric the maps ϕ_{st} are morphisms.

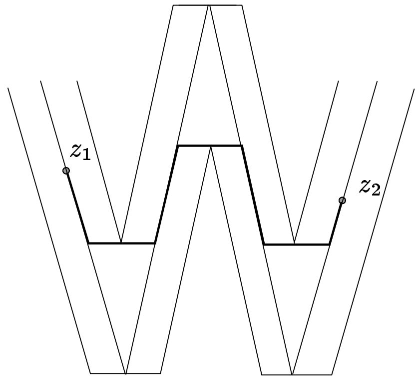

Figure 2: A taut corner path in the subset W_t between the points z_1 and z_2. The central line is the graph of the morphism. The essential pieces are the segments which lie in the graph; the nonessential pieces are the horizontal segments.

Lemma 3.1 *For the above definitions: $T_0 = T$, $T_1 = T'$ as G–trees, $\phi_{00} = Id_T$ and $\phi_{01} = \phi$.*

Proof The only nonobvious claim here is $T_1 = T'$. This is equivalent to saying that the sets $W_1 \cap (T \times \{y\})$ are connected. Let $(x_1, y), (x_2, y) \in W_1 \cap (T \times$

$\{y\})$. We will show that these two points lie in the same component. Choose $(z_1,y),(z_2,y)\in W_1\cap(T\times\{y\})$ such that $\phi(z_i)=y$ and $(x_i,y),(z_i,y)$ are in the same component of $W_1\cap(T\times\{y\})$ for $i=1,2$. For $z\in[z_1,z_2]$ we have $d(\phi(z),y)\leq m(\phi)$. Thus the pairs of points $(z_1,y),(z_2,y)$ are in the same component of $W_1\cap(T\times\{y\})$. Then as (x_1,y) is in the same component as (z_1,y) and (x_2,y) is in the same component as (z_2,y), the points (x_1,y) and (x_2,y) are in the same component. Thus $W_1\cap(T\times\{y\})$ is connected. □

Lemma 3.2 *If T' is irreducible, then so is T_t for $0\leq t\leq 1$.*

Proof As the G–tree T' is irreducible, there are $g,h\in G$ which act hyperbolically on T' such that $T_g\cap T_h$ is empty or compact [3]. As equivariant maps cannot make elliptic elements act hyperbolically, g,h act hyperbolically in T_t. The maps ϕ_{st} are quasi-isometries, hence the axes of g and h have empty or compact intersection. This implies that the G–tree T_t is irreducible. □

The following lemma is obvious.

Lemma 3.3 *T_t is in the same unnormalized deformation space as T and T' for $0\leq t\leq 1$.*

Remark 3.4 For future reference we remark that the above construction is invariant under the $\mathbb{R}^+$–action. In other words if we scale both T and T' by a nonzero positive number k, then the trees T_t are scaled by k.

4 Continuity of deformation

Fix an unnormalized deformation space $\mathcal{X}$. Recall that $\mathcal{M}(\mathcal{X})$ is the space of all morphisms between G–trees in $\mathcal{X}$. Define $\Phi\colon \mathcal{M}(\mathcal{X})\times\{(s,t)\mid 0\leq s\leq t\leq 1\}\to\mathcal{M}(\mathcal{X})$ by $\Phi(\phi,(s,t))=\phi_{st}$. The goal of this section is the following theorem:

Theorem 4.1 *Φ is continuous.*

We have some work before we can prove this. The approach is the same as in Skora's preprint [11], with the addition of Lemmas 4.2 and 4.4. We will consider a fixed morphism $\phi\colon X\to X'$ between finite simplicial trees and prove some results about morphisms $\psi\colon Y\to Y'$ which are close to ϕ. The main step is to show Lemma 4.7: if the two morphisms $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ are close and we fold both X and Y for a similar amount of time, then the two folded trees have comparable lengths. To prove this, we show that for a taut corner path in $W_t(\phi)$, the individual pieces are related to taut corner paths of comparable length in $W_s(\psi)$ when both s and t are close and ϕ and ψ are close.

Our first step is to show that maps close to ϕ have similar folding data.

Lemma 4.2 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Then for all $\epsilon > 0$ there is a $\delta > 0$ such if (R, R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ then $|m(\phi) - m(\psi)| < \epsilon$.*

Proof Let $\epsilon > 0$ be arbitrary. There are two cases to deal with. Notice that the lemma is symmetric with respect to ϕ and ψ.

Case 1 $m(\phi), m(\psi) > 0$

Set $\delta = \min\{\frac{\epsilon}{2}, \frac{m(\phi)}{3}, \frac{m(\psi)}{3}\}$. Let $z \in [x, x']$ be such that $d(\phi(z), \phi(x)) = m(\phi)$ and $\phi(x) = \phi(x')$. There are corresponding points $y, y', w \in Y$ such that $xRy, x'Ry'$ and zRw. As we can assume that R is full, we may assume that $w \in [y, y']$. Then $d(\psi(y), \psi(y')) < \delta$ and $d(\psi(w), \psi(y)), d(\psi(w), \psi(y')) > m(\phi) - \delta > \delta$. Hence there is a subsegment contained in $[y, y']$ and containing w which is folded. Thus $m(\psi) > m(\phi) - 2\delta$. Repeating the argument for Y we see that $m(\phi) > m(\psi) - 2\delta$. Hence we see that $|m(\phi) - m(\psi)| < \epsilon$.

Case 2 $m(\phi) = 0$ and $m(\psi) > 0$

Set $\delta = \frac{\epsilon}{2}$. Let $[y, y'] \subseteq Y$ be a folded segment where $w \in [y, y']$ attains $d(\psi(w), \psi(y)) = m(\psi)$. For corresponding points $x, x', z \in X$, we have that $[x, x']$ is embedded and $d(\phi(x), \phi(x')) < \delta$. Hence for all $z \in [x, x']$, we have $d(\phi(z), \phi(x)) < \delta$. Thus $m(\psi) < \epsilon$. □

Let $F(\phi)$ denote the number of fold points for the morphism $\phi\colon X \to X'$. Thus for $N(\phi) = 3(F(\phi) + 1)$ we have that any taut corner path γ in $W_t(\phi)$ can be written as $\gamma = \gamma_1 \cdots \gamma_n$ with $n \leq N(\phi)$ where each γ_i is either essential or nonessential.

We need a similar statement about morphisms close to ϕ. It is easy to see that we cannot expect a universal bound, but we can bound the number of large folds, which is sufficient. For $d > 0$, we introduce an equivalence relation on the set of fold points defined by $z \sim_d z'$ if there is a sequence of fold points: $z = z_0, \ldots, z_n = z'$ such that $d(z_i, z_{i+1}) < 2d$. Let

$$F_d(\psi) = \{\{\text{fold points for } \psi\}/\sim_d\} \setminus \{\text{classes without a } d\text{-deep point}\}.$$

Notice that $|F_{m(\psi)s}(\psi)|$ is the number of fold points for the map $\psi_{s1}\colon Y_s \to Y'$. Suppose (R, R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ where $\delta \leq d$. Then $|F_d(\psi)|$ is bounded independent of ψ as for each class in $F_d(\psi)$ we have a $\frac{d}{2}$ neighborhood in X, and the neighborhoods for different classes are disjoint. Set $F_d(\phi)$ to be the maximum of $|F_d(\psi)|$ over all morphisms $\psi\colon Y \to Y'$ for which there is a d–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$. As above we let $N_d(\phi) = 3(F_d(\phi)+1)$. Thus if ζ is a taut corner path in $W_s(\psi)$ then we can write $\zeta = \zeta_1 \cdots \zeta_n$ with $n \leq N_d(\phi)$ where each ζ_i is either nonessential or has length equal to the length of its image in Y'.

We now show that for a taut corner path in $W_t(\phi)$, the individual pieces are related to a taut corner path in $W_s(\psi)$ of comparable length. This is proven for the essential pieces first. As a convention when taking several points in X and points related to them in Y, if some of the points in X are the same we require that the related points in Y are the same.

Lemma 4.3 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Let z_1, z_2 be points in $\mathrm{graph}(\phi) \subseteq W_t(\phi)$ such that the taut corner path γ between them lies entirely in $\mathrm{graph}(\phi)$. Then for all $\epsilon > 0$ there is a $\delta > 0$ such that if (R, R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ and $w_i \in \mathrm{graph}(\psi)$ where $z_i(R, R')w_i$ for $i = 1, 2$, then $|\,\mathrm{length}(\gamma) - \mathrm{length}(\zeta)| < \epsilon$ where ζ is the taut corner path ζ in $W_s(\psi)$ from w_1 to w_2,*

Proof Let $\epsilon > 0$ be arbitrary. Let $\delta = \epsilon$ and assume the data in the hypothesis. Let $z_i = (x_i, x_i'), w_i = (y_i, y_i')$ for $i = 1, 2$ and let ζ be the taut corner path in $W_s(\psi)$ connecting w_1 to w_2. By hypothesis $\mathrm{length}(\gamma) = d(x_1, x_2) = d(x_1', x_2')$. As $\mathrm{length}(\zeta) \leq d(y_1, y_2) < d(x_1, x_2) + \delta$ and $\mathrm{length}(\zeta) \geq d(y_1', y_2') > d(x_1', x_2') - \delta$, we have the conclusion of the lemma. □

Next we have a similar statement for the nonessential pieces:

Lemma 4.4 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Let z_1, z_2 be points in $\mathrm{graph}(\phi) \subseteq W_t(\phi)$ such that the taut corner path γ between them lies entirely in a component of $\mathcal{F}_t$. Then for all $\epsilon > 0$ there is a $\delta > 0$ such that if (R, R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$, $|m(\psi)s - m(\phi)t| < \delta$, and $w_i \in \mathrm{graph}(\psi)$ where $z_i(R, R')w_i$ $i = 1, 2$, then $\mathrm{length}(\zeta) < \epsilon$ where ζ is the taut corner path in $W_s(\psi)$ from w_1 to w_2,*

Proof Let $\epsilon > 0$ be arbitrary. We have two cases depending on $m(\phi)$ and t. Let (R, R') be a δ–approximation with z_i, w_i as in the statement above where δ is chosen in the individual cases. Say $z_i = (x_i, x_i'), w_i = (y_i, y_i')$ for $i = 1, 2$. Then from the definitions we have $x_1' = x_2'$ and $\phi([x_1, x_2])$ stays within $m(\phi)t$ of x_1'.

Case 1 $t = 0$ or $m(\phi) = 0$

Let $\delta = 1$. Then as $z_1 = z_2$, we have $w_1 = w_2$ by the above convention. Hence $\mathrm{length}(\zeta) = 0$.

Case 2 $t > 0$ and $m(\phi) > 0$

Let $N = N_d(\phi)$ as above where $d = \frac{2m(\phi)t}{3}$ and set $\delta = \min\{\frac{\epsilon}{2N}, \frac{m(\phi)t}{3}\}$. As $m(\psi)s > m(\phi)t - \delta \geq d$, the number of fold points for ψ_{s1} is less than $F_d(\phi)$. Therefore we can write $\zeta = \zeta_1 \cdots \zeta_n$ where $n \leq N$ and each ζ_i is nonessential or has length equal to the length of its image in Y'.

If $\psi([y_1, y_2])$ is contained within a $m(\psi)s$ neighborhood about y_1', then ζ is nonessential. This might not be the case, but the length of an essential piece of ζ is bounded by how far $\psi([y_1, y_2])$ travels away from y_1': $\mathrm{length}(\zeta_i) \leq \max\{\{d(\psi(y), y_1') - m(\psi)s \mid y \in [y_1, y_2]\}, 0\}$. Now we use fullness of the approximations to see: $\mathrm{length}(\zeta_i) \leq \max\{\{d(\phi(x), x_1') - m(\psi)s + \delta \mid x \in [x_1, x_2]\}, 0\} \leq \max\{\{d(\phi(x), x_1') - m(\phi)t + 2\delta \mid x \in [x_1, x_2]\}, 0\} \leq 2\delta$.

Thus we have $\mathrm{length}(\zeta) \leq \sum \mathrm{length}(\zeta_i) \leq 2\delta n < \epsilon$. □

Putting together the previous two lemmas we have:

Lemma 4.5 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Let z_1, z_2 be points in $\mathrm{graph}(\phi \subseteq W_t(\phi)$ and $\gamma = \gamma_1 \cdots \gamma_n$ the taut corner path between them. Then for all $\epsilon > 0$ there is a $\delta > 0$ such if (R, R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$, $|m(\psi)s - m(\phi)t| < \delta$ and $w_i \in \mathrm{graph}(\psi)$ where $z_i(R, R')w_i$ for $i = 1, 2$, then there is a path $\zeta = \zeta_1 \cdots \zeta_n$ in $W_s(\psi)$ from w_1 to w_2 with each ζ_i a taut corner path which satisfies $|\,\mathrm{length}(\gamma_i) - \mathrm{length}(\zeta_i)| < \epsilon$ for $i = 1, \ldots, n$.*

The next lemma is a converse to Lemma 4.4 and the proof is simpler as we know how many fold points ϕ has. Recall that the image of $z \in W_t$ in the quotient tree is denoted $[z]_t$.

Lemma 4.6 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Then for all $\epsilon > 0$ there is a $\delta > 0$ such that if (R, R') is a δ–approximation between*

$\phi\colon X \to X'$ and $\psi : Y \to Y'$ and $|m(\psi)s - m(\phi)t| < \delta$ then $d([z_1]_t, [z_2]_t) < \epsilon$ where $z_i \in \mathrm{graph}(\phi)$, $w_i \in \mathrm{graph}(\psi)$ with $z_i(R,R')w_i$ for $i = 1,2$ and $[w_1]_s = [w_2]_s$.

Proof Let $\epsilon > 0$ be arbitrary and $\delta = \frac{\epsilon}{2N}$, where $N = N(\phi)$. Let $\gamma = \gamma_1 \cdots \gamma_n$ be the taut corner path from z_1 to z_2 where each piece is either essential or nonessential and $n \leq N$. Using the same argument as in case 2 for 4.4, we can bound the lengths of the γ_i by 2δ. Thus $\mathrm{length}(\gamma) \leq 2\delta n < \epsilon$. □

Now using the previous two lemmas, we are able to show that close morphisms which are folded for a similar amount of time have comparable lengths. We will also remove the dependence on the folding data using Lemma 4.2.

Lemma 4.7 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. Then for all $\epsilon > 0$ there is a $\delta > 0$ such that if (R,R') is a δ-approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ and $|s - t| < \delta$ then $|d([z_1]_t, [z_2]_t) - d([w_1]_s, [w_2]_s)| < \epsilon$ where $z_i \in \mathrm{graph}(\phi)$, $w_i \in \mathrm{graph}(\psi)$ with $z_i(R,R')w_i$ for $i = 1,2$.*

Proof Let ϵ be arbitrary. Set $\epsilon_1 = \epsilon_2 = \frac{\epsilon}{4N}$, where $N = N(\phi)$. Use these to find δ_1, δ_2 from Lemma 4.5 and Lemma 4.6 respectively. Let $\epsilon_3 = \frac{1}{2}\min\{\delta_1, \delta_2\}$ and take δ_3 from Lemma 4.2 using ϵ_3. Finally set $\delta = \min\{\epsilon_2, \frac{\epsilon_3}{m(\phi)}, \delta_3\}$.

The choice of these parameters implies that if (R,R') is a δ-approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$, and $|s-t| < \delta$, then $|m(\psi) - m(\phi)| < \epsilon_3$. Thus $m(\psi)s - m(\phi)t < (m(\phi) + \epsilon_3)s - m(\phi)t < m(\phi)(s-t) + \epsilon_3 < \delta_1, \delta_2$ and similarly $m(\phi)t - m(\psi)s < \delta_1, \delta_2$. Therefore we can use Lemma 4.5 and Lemma 4.6.

We can write the taut corner path connecting z_1 and z_2 as $\gamma = \gamma_1 \cdots \gamma_n$ where $n \leq N$ and each γ_i is either essential or nonessential. Hence by Lemma 4.5 we have a path $\zeta = \zeta_1 \cdots \zeta_n$ connecting w_1 to w_2 where each piece is a taut corner path and $|\,\mathrm{length}(\gamma_i) - \mathrm{length}(\zeta_i)| < \epsilon_1$. Hence $d([w_1]_s, [w_2]_s) \leq \sum \mathrm{length}(\zeta_i) < \sum(\mathrm{length}(\gamma_i) + \epsilon_1) < d([z_1]_t, [z_2]_t) + \epsilon$.

If $d([w_1]_s, [w_2]_s) \leq d([z_1]_t, [z_2]_t) - \epsilon$, then as $d([z_1]_t, [z_2]_t) = \sum \mathrm{length}(\gamma_i) < (\sum \mathrm{length}(\zeta_i)) + \frac{\epsilon}{2}$ we get that $d([w_1]_s, [w_2]_s) < \sum \mathrm{length}(\zeta_i) - \frac{\epsilon}{2}$. Since the only folds of $[\zeta]$ in Y_s are at the intersection points of $[\zeta_i]$ with $[\zeta_{i+1}]$, there are two points q_1 and q_2 on ζ such that the length along ζ between these two points is greater than $\frac{\epsilon}{2N}$ but these are the same point in Y_s. Thus for points $p_1, p_2 \in W_t(\phi)$ with $p_i(R,R')q_i$ for $i = 1,2$ we have $d([p_1]_t, [p_2]_t) > \frac{\epsilon}{2N} - \delta \geq \epsilon_2$. However the choice of δ_2 implies that $d([p_1]_t, [p_2]_t) < \epsilon_2$ by Lemma 4.6. Hence we have a contradiction. Therefore $|d([z_1]_t, z_2]_t) - d([w_1]_s, [w_2]_s)| < \epsilon$. □

Thus the folded trees have comparable lengths. We can use this to build an ϵ–approximation between these trees. For morphisms $\phi\colon X \to X'$, $\psi\colon Y \to Y'$ which are related by an ϵ–approximation (R, R') we define a new relation $[R,R]_{ts}$ from X_t to Y_s by $[z]_t[R,R']_{ts}[w]_s$ whenever $z(R,R')w$ for $z \in \mathrm{graph}(\phi)$ and $w \in \mathrm{graph}(\psi)$. We now prove a lemma about this relation when s and t are close.

Lemma 4.8 *Let $\phi\colon X \to X'$ be a morphism of finite simplicial trees. For all $\epsilon > 0$ there is a $\delta > 0$ such that if (R,R') is a δ–approximation between $\phi\colon X \to X'$ and $\psi\colon Y \to Y'$ and $|s-t| < \delta$ then $[R,R']_{ts}$ is an ϵ–approximation from X_t to Y_s. If R and R' are P–equivariant, then so is $[R,R']_{ts}$.*

Proof Let $\epsilon > 0$ be arbitrary and choose δ from Lemma 4.7. Given data as in the hypothesis, $[R,R']_{ts}$ is an ϵ–approximation. It also follows that if R and R' are P–equivariant, then so is $[R,R]_{ts}$. □

Given an arbitrary morphism $\phi\colon T \to T'$ between G–trees in the unnormalized deformation space $\mathcal{X}$, for subtrees $X \subseteq T, X' \subseteq T'$ such that $\phi(X) \subseteq X'$ we can define X_t as $W_t(\phi|_X)/(\mathcal{F}_t \cap (X \times X'))$. We can now prove that Φ is continuous.

Proof Let $\epsilon > 0$ be arbitrary. Let $\phi\colon T \to T'$ and $0 \le s \le t \le 1$ be given. Assume U is the basic open set around ϕ_{st} given by $U = U(\phi_{st}, X \times X', P, \epsilon)$ where $X \subseteq T, X' \subseteq T'$ are finite subtrees and P is a finite subset of G. Let δ be given by Lemma 4.8, and $V = U(\phi, X \times X', P, \delta)$.

Suppose $\psi\colon \widetilde{T} \to \widetilde{T}'$ with $\psi \in V$ and $|p-s| < \delta, |q-t| < \delta$. We will show that $\psi_{pq} \in U$.

For some finite subtrees $Y \subseteq \widetilde{T}, Y' \subseteq \widetilde{T}'$ there is a δ–approximation (R,R') from $\phi : X \to X'$ to $\psi : Y \to Y'$. The claim is that $([R,R']_{sp}, [R,R]_{tq})$ is a closed ϵ–approximation from $\phi_{st} : X_s \to X_t$ to $\psi_{pq} : Y_p \to Y_q$. The choice of δ implies that both $[R,R']_{sp}$ and $[R,R']_{qt}$ are ϵ–approximations by Lemma 4.8. If $[z]_s[R,R']_{sp}[w]_p$ then we have that $[z]_t[R,R']_{tq}[w]_q$. Therefore $\psi_{pq} \in U$. □

5 Continuity of base point

For a G–tree $T \in \mathcal{X}$ define $l_T(S) = \min_{x\in T} \max_{g\in S} d(x, gx)$, where S is some finite subset of G. The characteristic set of S is $T_S = \{x \in T \mid l_T(S) =$

$\max_{g\in S} d(x,gx)\}$. This agrees with the earlier notion for characteristic set when the subgroup generated by S is elliptic. Clearly for $g \in S$ we have $l_T(g) \leq l_T(S)$. We let S' be the subset of S where this is an equality, i.e. $S' = \{g \in S \mid l_T(g) = l_T(S)\}$. Finally we define $Z_S = \bigcap_{g\in S'} T_g$.

Lemma 5.1 *Let T be a G–tree and let S be a finite subset of G. Then T_S is contained in the union of a finite simplicial tree and Z_S. In particular, if Z_S is a finite simplicial tree, then T_S is a finite simplicial tree.*

Proof Let $x \in T$ and X be the union of all arcs from x to T_g for $g \in S$, then X is a finite simplicial tree. If $y \in T_S$ is not in X, let z be the closest point in X to y. Then $d(y,gy) \geq d(z,gz)$ for all $g \in S$ as $d(y,T_g) \geq d(z,T_g)$ with equality only if $y \in T_g$. If $g \in S'$ then $d(y,gy) \geq d(z,gz) \geq l_T(S)$. As $y \in T_S$ we have $l_T(S) \geq d(y,gy)$. Hence we have equality $d(y,gy) = d(z,gz)$ for $g \in S'$. Thus $y \in T_g$ for all $g \in S'$ and hence $y \in Z_S$. □

Let S generate G. Then for irreducible G–trees T, Z_S is finite, hence so is T_S. We have some simple lemmas on the shape and position of T_S based on $l_T(S)$ and $l_T(g)$ that will be used in Proposition 5.4.

Lemma 5.2 *Suppose that T_S is finite. Then T_S is either a point or a segment. Moreover, the latter only occurs when there is a $g \in S$ such that $l_T(g) = l_T(S)$. In both cases, there are distinct $g_1, g_2 \in S$ such that $d(x,g_1x) = d(x,g_2x) = l_T(S)$ for all $x \in T_S$.*

Proof Suppose $l_T(S) > \max_{g\in S} l_T(g)$ and there are distinct points $x_1, x_2 \in T_S$. Let $g_1, g_2 \in S$ be such that $\max_{g\in S} d(x_i,gx_i) = d(x_i,g_ix_i) = l_T(S)$ for $i = 1,2$. Thus $x_i \notin T_{g_i}$. Consider the segment $[x_1,x_2]$. Let $y \in [x_1,x_2]$ and $y \neq x_1,x_2$. Then for any $g \in S$, $d(y,T_g) < d(x_i,T_g)$ for either $i = 1$ or 2, hence $\max_{g\in S} d(y,gy) < l_T(S)$. This is a contradiction, therefore $T_S = \{x\}$. Now notice that there are $g_1, g_2 \in S$ such that $d(x,g_ix) = l_T(S)$. For if there was only one such g, then for some point y near x on the arc from x to T_g, $\max_{g\in S} d(y,gy) < d(x,gx) = l_T(S)$, which is a contradiction.

If $l_T(S) = l_T(g)$ for $g \in S$ then $T_S \subset T_g$. Therefore T_S is either a point or a segment. If there were only one such $g \in S$ such that $l_T(S) = l_T(g)$, then T_S is open by a similar argument as above. This is a contradiction. □

Recall that for $A \subseteq T$ closed, we let $p_A\colon T \to A$ denote the nearest point projection.

Lemma 5.3 *Let $z \in T \setminus T_S$ and $x = p_{T_S}(z)$. Then for some $g \in S$ such that $d(x, gx) = l_T(S)$, we have that x is on the arc from z to T_g.*

Proof Suppose not. Then for points $x' \in [x, z]$ near x, $d(x', T_g) \leq d(x, T_g)$ for all $g \in S$ such that $d(x, gx) = l_T(S)$. This is a contradiction. □

For an irreducible G–tree T, let x_* denote the midpoint of T_S. This is called the *basepoint* of the action. Define a map $b(T)\colon G \to T$ by $g \mapsto gx_*$. This defines a map $b : \mathcal{X} \to \mathcal{E}(G, \mathcal{X})$ where $\mathcal{E}(G, \mathcal{X})$ is the space of equivariant maps from G to G–trees in $\mathcal{X}$. The topology for $\mathcal{E}(G, \mathcal{X})$ is the Gromov-Hausdorff topology defined in section 2 where we consider G as a metric G-space. The actual metric we place on G does not matter as the domain is fixed in $\mathcal{E}(G, \mathcal{X})$. The remainder of this section is used to prove that $b\colon \mathcal{X} \to \mathcal{E}(G, \mathcal{X})$ is a continuous function.

Proposition 5.4 *b is continuous.*

Proof This amounts to showing that close G–trees in $\mathcal{X}$ have close basepoints. Let $T \in \mathcal{X}$, there are two cases depending on $l_T(S)$.

Case 1 $l_T(S) > \max_{g \in S} l_T(g)$

By Lemma 5.2 we have that $T_S = \{x_*\}$. Within the set of $g \in S$ such that $d(x_*, gx_*) = l_T(S)$, there are two elements g_1, g_2 such that x_* is on the spanning arc from T_{g_1} to T_{g_2}. Let x_i be the point on T_{g_i} nearest to x_*. Thus $x_* \in [x_1, x_2]$ and $d(x_1, x_2) = d(x_1, x_*) + d(x_*, x_2)$.

Let U be the basic open set $U = U(b(T), P \times K, P, \epsilon)$, where $S \subseteq P$ and $P(\{x_*, x_1, x_2\}) \subseteq K$. By the remark in section 2, we can assume that P and K contain these subsets by shrinking U. Also let $V = U(T, K, P, \delta)$, where $\delta = \frac{1}{4}\min\{\epsilon, d(x_*, x_1), d(x_*, x_2)\}$. Suppose that $Y \in V$, we will show that $b(Y) \in U$. By definition, there is a P–equivariant closed δ–approximation $R \subseteq K \times L$ for some finite subtree $L \subseteq Y$.

Fix related points in L: $x_* R y_*, x_i R y_i$ for $i = 1, 2$. By fullness of R, we may assume that $y_* \in [y_1, y_2]$. Our object now is to show that y_* is close to every point in Y_S, in particular, the midpoint of Y_S. This involves some inequalities. As $|d(x_*, gx_*) - d(y_*, gy_*)| < \delta$ for all $g \in S$ we have $\max_{g \in S} d(y_*, gy_*) < \max_{g \in S} d(x_*, gx_*) + \delta = d(x_*, g_i x_*) + \delta < d(y_*, g_i y_*) + 2\delta$ for $i = 1, 2$. Therefore, if $y \in Y_S$, then $d(y, g_i y) \leq \max_{g \in S} d(y, gy) \leq \max_{g \in S} d(y_*, gy_*) < d(y_*, g_i y_*) +$

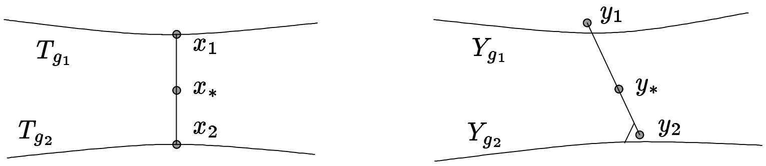

Figure 3: The characteristic sets in T and the related points in Y for case 1 in Proposition 5.4

2δ hence $d(y, Y_{g_i}) < d(y_*, Y_{g_i}) + \delta$ for $i = 1, 2$. We will now show that y, y_* are close to the spanning arc α, from Y_{g_1} to Y_{g_2}.

$$\begin{aligned} |d(y_*, Y_{g_i}) - d(y_*, y_i) - d(y_i, Y_{g_i})| = \; & |\tfrac{1}{2}(d(y_*, g_i y_*) - l_Y(g_i)) - d(y_*, y_i) - \\ & \tfrac{1}{2}(d(y_i, g_i y_i) - l_Y(g_i))| \\ = \; & |\tfrac{1}{2}(d(y_*, g_i y_*) - d(x_*, g_i x_*)) + \\ & \tfrac{1}{2}(d(x_i, g_i x_i) - d(y_i, g_i y_i)) + \\ & (d(x_*, x_i) - d(y_*, y_i))| \\ < \; & 2\delta. \end{aligned}$$

Hence $d(y_*, Y_{g_i}) - d(y_i, Y_{g_i}) > d(y_*, y_i) - 2\delta > d(x_*, x_i) - 3\delta > 0$. As $y_* \in [y_1, y_2]$ we have that y_* is on α. Thus for $y \in Y_S$, $d(y_*, y) < \delta$. Let $y_0 \in Y_S$ be the basepoint.

We claim that (Id_G, R_δ) is a P–equivariant closed ϵ–approximation between $b(T)\colon P \to K$ and $b(Y)\colon P \to L$. The only nontrivial check is that for $g \in P$, $b(T)(g) R_\delta b(Y)(g)$. This follows from the following calculation as for $g \in P$ we have $g x_* R g y_*$:

$$d(gx_*, b(T)(g)) + d(gy_*, b(Y)(g)) = d(gy_*, gy_0) = d(y_*, y_0) < \delta.$$

This implies $b(Y) \in U$.

Case 2 : $l_T(S) = \max_{g \in S} l_T(g)$

Let $h \in S$ be such that $l_T(h) = l_T(S)$, then $T_S \subset T_h$ as in Lemma 5.2. If $x_1 \neq x_2$ assume that h translates from x_1 to x_2.

Let U be the basic open set $U = U(b(T), P \times K, P, \epsilon)$ where $S, S^{-1} \subseteq P$ and $P([h^{-1}x_1, hx_2]) \subseteq K$. As in case 1, this is possible by shrinking U. Let $V = U(T, K, P, \delta)$ where $\delta = \frac{1}{9}\min\{\epsilon, l_T(S)\}$. Suppose that $Y \in V$, we will show that $b(Y) \in U$. By definition, there is a P–equivariant closed δ–approximation $R \subseteq K \times L$ for some finite subtree $L \subseteq Y$.

Fix related points in L: $x_* R y_*, x_i R y_i$ for $i = 1, 2$. Again, by the fullness of R we may assume that $y_* \in [y_1, y_2]$. Our object now is to show that the

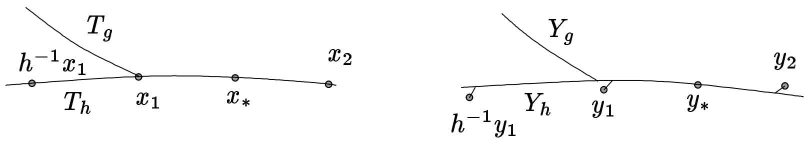

Figure 4: The characteristic sets in T and the related points in Y for case 2 in Proposition 5.4

Hausdorff distance between $[y_1, y_2]$ and Y_S is small. As before, this involves some inequalities. Our first step is to show that points in $[y_1, y_2]$ are close to some point in Y_S.

Let $z \in [y_1, y_2]$ and $x \in [x_1, x_2] = T_S$ where xRz. Then $\max_{g\in S} d(z, gz) < \max_{g\in S} d(x, gx) + \delta = d(x, hx) + \delta < d(z, hz) + 2\delta$. Since R is full, this is true for any $z \in [y_1, y_2]$. Note that the above inequality implies $l_Y(S) < l_T(S) + \delta$.

We now show that the segment $[h^{-1}y_1, hy_2]$ is close to the axis Y_h. Let $z, z' \in [h^{-1}y_1, hy_2]$ and $x, x' \in [h^{-1}x_1, hx_2]$ where $xRz, x'Rz'$.

$$\begin{aligned} |d(z, Y_h) - d(z', Y_h)| &= |\tfrac{1}{2}(d(z, hz) - l_Y(h)) - \tfrac{1}{2}(d(z', hz') - l_Y(h))| \\ &= \tfrac{1}{2}|(d(z, hz) - d(x, hx)) + (d(x', hx') - d(z, hz'))| \\ &< \delta. \end{aligned}$$

In particular $|d(h^{-1}y_1, Y_h) - d(hy_2, Y_h)| < \delta$, as $d(h^{-1}y_1, hy_2) > 2l_T(h) - \delta > 2\delta$ this implies that there is a $z_0 \in [h^{-1}y_1, hy_2] \cap Y_h$. Hence for any $z \in [h^{-1}y_1, hy_2]$ we have $d(z, Y_h) < \delta$. Likewise the same is true for $z \in [y_1, y_2]$.

Now for $z \in [y_1, y_2]$, $l_Y(S) - 2\delta \leq \max_{g\in S} d(z, gz) - 2\delta < d(z, hz) < l_Y(h) + 2\delta < l_Y(S) + 2\delta$. For $z \in [y_1, y_2]$ that are not in Y_S, let $y = p_{Y_S}(z)$ and let $g' \in S$ be given by Lemma 5.3. Then $d(z, y) = \frac{1}{2}(d(z, g'z) - d(y, g'y)) \leq \frac{1}{2}(\max_{g\in S} d(z, gz) - l_Y(S)) < \delta$. Hence for $z \in [y_1, y_2]$, we have $d(z, Y_S) < \delta$.

For the opposite inequality we show that points in Y_S are close to some point in $[y_1, y_2]$. We do so by showing that points far enough away from $[y_1, y_2]$ cannot lie in Y_S. First note that the above inequality implies: $l_Y(S) - l_Y(h) < 4\delta$. Hence if $y' \in Y_S$, then $2d(y', Y_h) = d(y', hy') - l_Y(h) < l_Y(S) - (l_Y(S) - 4\delta)$. Thus $d(y', Y_h) < 2\delta$. Recall that we have shown $l_Y(S) < l_T(S) + \delta$.

The idea now is to use Lemma 5.3 on points far from $[y_1, y_2]$. Assume that $y' \in Y_S$ and $d(y', [y_1, y_2]) \geq 4\delta$. Then there is some point $y \in Y_h \cap L$ with $d(y, [y_1, y_2]) \geq 2\delta$. Without loss of generality, we assume that y is closer to y_1 than to y_2. Let $x \in T_h \cap K$ be such that xRy. Then $d(x, x_1) \geq \delta$. Hence by Lemma 5.3 there is a $g \in S$ such that $d(x, gx) \geq l_T(S) + 2\delta$. Therefore $l_Y(S) \geq d(y, gy) \geq l_T(S) + \delta > l_Y(S)$, which is a contradiction. Therefore the Hausdorff distance between Y_S and $[y_1, y_2]$ is less than 4δ. Let $y_0 \in Y_S$ be the basepoint, then $d(y_0, y_*) < 4\delta$. Now proceed as in case 1 using the P–equivariant closed ϵ–approximation $(Id_G, R_{4\delta})$.

This completes the proof. □

Remark 5.5 The technical statement proved in the above which is used later on in Lemma 6.4 is that if two trees Y and Z have subtrees, $L \subseteq Y$, $M \subseteq Z$ with $P\{b(Y)(1)\} \subseteq L$, $S \subseteq P$ and a P–equivariant ϵ–approximation $R \subseteq L \times M$, then if $z \in Z$ with $b(Y)(1)Rz$, we have $d(z, b(Z)(1)) < 4\epsilon$. In other words, any point related to the basepoint of Y is within 4ϵ of the basepoint of Z.

6 Contractibility of deformation space

To prove the contractibility of the unnormalized deformation space $\mathcal{X}$, we construct a homotopy onto a contractible subset. To define the homotopy, for any G–tree $T' \in \mathcal{X}$ we need to build a nice map from some fixed G–tree $T \in \mathcal{X}$ to T'. To ensure that the map $T \to T'$ is nice, we will need T to be reduced.

Definition 6.1 A G–tree T is *reduced* if for all edges $e = [u, v]$, u is G–equivalent to v if $G_e = G_u$.

This is equivalent to Forester's definition in [5] where a tree is said to be reduced if it admits no collapse moves. We will use this notion via the next lemma.

Lemma 6.2 *Let T be a reduced G–tree and $u, v \in T$ vertices such that there is an edge $e = [u, v]$ and x a vertex with $G_u, G_v \subseteq G_x$. Then u is G–equivalent to v.*

Proof Without loss of generality, assume that v is closer to x than u is. Then $[u, x] = e \cup [v, x]$ and as G_u stabilizes $[u, x]$ this implies that $G_u = G_e$. Hence as T is reduced, the two endpoints of e are G–equivalent. □

We now require that our unnormalized deformation space $\mathcal{X}$ contains a G–tree with finitely generated vertex groups. In particular as all G–trees in $\mathcal{X}$ are cocompact, there is a reduced tree $T \in \mathcal{X}$ with finitely generated vertex groups. Define $\mathcal{T}(T, \mathcal{X})$ as the space of all continuous maps from T to G–trees in $\mathcal{X}$ that take vertices to vertices and are injective on the edges of T. We call such maps *transverse*. This has a different meaning than in [11], where transverse only implies cellular. We topologize $\mathcal{T}(T,\mathcal{X})$ using the Gromov-Hausdorff topology from section 2.

Our aim now is to build a section $B\colon \mathcal{X} \to \mathcal{T}(T,\mathcal{X})$. Let G be finitely generated by S and fix $X \subseteq T$ a subtree whose edges map bijectively to T/G. We follow Forester's construction from Proposition 4.16 in [5]. Order the vertices of X as $\{v_1,\dots,v_k\}$ where vertices in the same orbit are consecutive. For the ith orbit $v_{i_0},\dots,v_{i_0+d}$ let $g_{i_0}=1$ and fix $g_{i_0+q}\in G$ such that $g_{i_0+q}v_{i_0}=v_{i_0+q}$ for $1\le q\le d$. As the path $[v_{i_0}, g_{i_0+q}^{-1}g_{i_0+p}v_{i_0}]$ for $1\le p,q\le d$ is contained in X, it maps bijectively to T/G. Therefore the products $g_{i_0+q}^{-1}g_{i_0+p}$ are hyperbolic for $p\neq q$ (Lemma 2.7(b) [5]).

Given $Y\in\mathcal{X}$, we define the map $B(Y)\colon T\to Y$ first on the vertices of X. Let $y_*=b(Y)(1)$, where b is the basepoint map of Proposition 5.4. Recall that p_A is projection onto the closed subset A. Consider an orbit $\{v_{i_0},\dots,v_{i_0+d}\}$. If $G_{v_{i_0}}\neq\{1\}$ then let $Y_{i_0}\subseteq Y$ be the characteristic set of $G_{v_{i_0}}$. Otherwise, let $Y_{i_0}=y_*$. As T is reduced, $G_{v_{i_0}}=\{1\}$ can only happen if G is a finitely generated free group of rank at least 2. In which case T/G is a rose and there is only one orbit of vertices in X. Define $B(Y)$ on the orbit by: $v_{i_0+d}\mapsto g_{i_0+d}p_{Y_{i_0}}(y_*)$.

We now show that $B(Y)$ can be extended to a transverse map. If there is an edge $e\subseteq X$ where $e=[u,v]$ with $B(Y)(u)=B(Y)(v)=x'\in Y$, then $G_u,G_v\subseteq G_{x'}$. This subgroup must fix a vertex $x\in T$, hence $G_u,G_v\subseteq G_x$ and by Lemma 6.2, u and v must be in the same orbit. But if v_i and v_j are in the same orbit then as ${g_i}^{-1}g_j$ is hyperbolic for $i\neq j$ necessarily $B(Y)(v_i)\neq B(Y)(v_j)$. Thus we can linearly map each edge of X injectively into Y. Now extend $B(Y)$ to all of T equivariantly. As $B(Y)$ is injective on each edge this defines $B\colon\mathcal{X}\to\mathcal{T}(T,\mathcal{X})$.

For the ith orbit, let G_i be the vertex stabilizer of the first vertex in this orbit and denote the characteristic set for G_i by the subscript i, i.e. $Y_{G_i}=Y_i$. If $G_i=1$ then as before, set $Y_i=y_*$. Let G_i be finitely generated by S_i, then for any G–tree $Y\in\mathcal{X}$ we have $Y_i=\cap_{s\in S_i}Y_s$. Let $Q\subseteq G$ be the union of the S_i's and S, a finite generating set for G.

Lemma 6.3 *B is continuous and $\mathcal{R}a(B(Y))=Y$ for all $Y\in\mathcal{X}$.*

If G is finitely generated free group of rank at least 2, then this follows from Proposition 5.4. Thus we assume that G is not free. Before we prove this lemma in general, we prove a statement about the position of the basepoint relative the fixed point sets.

Lemma 6.4 *Let $Y, Z \in \mathcal{X}$, and Y_i, Z_i be the characteristic sets as described above. Let $L \subseteq Y, M \subseteq Z$ be subtrees with $R \subseteq L \times M$ a P–equivariant ϵ–approximation where $Q \subseteq P$ and $P\{y_*, p_{Y_i}(y_*)\} \subseteq L$. Then for $z \in Z$ such that $p_{Y_i}(y_*)Rz$ we have $d(z, p_{Z_i}(z_*)) < 10\epsilon$.*

Proof Fix $\hat{z} \in Z$ where $y_* R \hat{z}$, then by Remark 5.5, $d(\hat{z}, z_*) < 4\epsilon$. The lemma follows from 3 observations:

(i) $|d(z_*, z) - d(y_*, p_{Y_i}(y_*)| < |d(\hat{z}, z) - d(y_*, p_{Y_i}(y_*))| + 4\epsilon < 5\epsilon$.

(ii) Let $g \in S_i$ be such that $p_{Z_i}(z_*) = p_{Z_g}(z_*)$. Then:

$$\begin{aligned} 2d(z_*, p_{Z_i}(z_*)) = d(z_*, gz_*) &< d(\hat{z}, g\hat{z}) + 8\epsilon \\ &< d(y_*, gy_*) + 9\epsilon = 2d(y_*, p_{Y_g}(y_*)) + 9\epsilon \\ &\leq 2d(y_*, p_{Y_i}(y_*)) + 9\epsilon. \end{aligned}$$

Likewise, running this argument with $h \in S_i$ such that $p_{Y_i}(y_*) = p_{Y_h}(y_*)$, we see that $|d(z_*, p_{Z_i}(z_*)) - d(y_*, p_{Y_i}(y_*))| < 5\epsilon$.

(iii) Let $g \in S_i$ be such that $p_{Z_i}(z) = p_{Z_g}(z)$. Then: $2d(z, p_{Z_i}(z)) = d(z, gz) < d(p_{Y_i}(y_*), gp_{Y_i}(y_*)) + \epsilon = \epsilon$.

Putting (i) and (ii) together: $|d(z_*, z) - d(z_*, p_{Z_i}(z_*))| < 10\epsilon$. If $[z_*, z]$ passes through Z_i then $d(z, p_{Z_i}(z_*)) < 10\epsilon$. If $[z_*, z]$ doesn't pass through Z_i, then $p_{Z_i}(z) = p_{Z_i}(z_*)$ and hence $d(z, p_{Z_i}(z_*)) = d(z, p_{Z_i}(z)) < \epsilon$ by (iii). □

Now we can prove Lemma 6.3.

Proof Let $Y \in \mathcal{X}$ and let U be a basic open set of $B(Y), U = U(B(Y), K \times L, P, \epsilon)$ where $Q \subseteq P$ and $P\{y_*, p_{Y_i}(y_*)\} \subseteq L$. Enlarge P such that $K \subseteq PX$. Also let V be a basic open set for $Y \in \mathcal{X}$, $V = U(Y, L, P, \delta)$ where $\delta = \frac{\epsilon}{21}$. If $Z \in V$, we have by definition a δ–approximation $R \subseteq L \times M$ for some $M \subseteq Z$. This is the set-up in Lemma 6.4. As before, we will show that $B(Z) \in U$.

We claim that $(Id, R_{10\delta})$ is an ϵ–approximation from $B(Y)$ to $B(Z)$, hence $B(Z) \in U$. As in the proof of Proposition 5.4, the only nontrivial check is that $B(Y)(x) R_{10\delta} B(Z)(x)$ for $x \in K$. Without loss of generality, we can assume that x is a vertex as the maps are linear on the edges. Now $x = gv$ for some $v \in X$, ordered first in its orbit and some $g \in P$. Let $A = Y_{G_v}, B = Z_{G_v}$, then $B(Y)(x) = gp_A(y_*)$ and $B(Z)(x) = gp_B(z_*)$. Let $z \in Z$ be such that $p_A(y_*)Rz$. Thus by Lemma 6.4:

$$d(gp_A(y_*), B(Y)(x)) + d(gz, B(Z)(x)) = d(gz, gp_B(z_*)) = d(z, p_B(z_*)) < 10\delta.$$

This completes the proof. □

The map $B(Y)$ is not a morphism in the sense used within this paper. However we can redefine the metric on T to get a new G–tree T_Y such that the map $B(Y)$ when regarded as a map $B(Y)\colon T_Y \to Y$ is a morphism. As each edge of T is mapped injectively via $B(Y)$ to Y we can remetrize each edge by pulling back the metric on Y. Thus we can remetrize T by setting the distance between two points to be the length of the geodesic path between them. Call this new G–tree T_Y. Then T_Y is equivariantly homeomorphic to the G–tree T. Let $\mathcal{G}(T)$ denote the set of G–trees in $\mathcal{X}$ which are equivariantly homeomorphic to the G–tree T. Recall that the volume of T is defined as $\mathrm{vol}(T) = \sum \mathrm{length}(e)$ where the sum is over the unoriented edges of T/G.

Proposition 6.5 *$\mathcal{G}(T)$ is homeomorphic to $\sigma \times \mathbb{R}^+$ where σ is an open simplex of dimension one less than the number of edges of T/G.*

Proof Fix an ordering $e_1, \ldots, e_n$ of the edges of T. This in turn gives an ordering of the edges of $T' \in \mathcal{G}(T)$. Let $h : \mathcal{G}(T) \to \sigma \times \mathbb{R}^+$ be defined by:

$$h(T') = (\frac{1}{\mathrm{vol}(T')}(\mathrm{length}(e_1), \ldots, \mathrm{length}(e_n)), \mathrm{vol}(T')). \tag{1}$$

It is clear that this map gives a bijection between the sets. As we are working with irreducible G–trees, as mentioned in section 2 the Gromov–Hausdorff topology is the same as the axes topology. A small change in $\mathcal{G}(T)$ of the length functions results in a small change in the lengths of the edges. And conversely, a small change in the length of the edges of T/G results in a small change of the length functions only for the hyperbolic conjugacy classes whose axis project down to paths which cross the rescaled edges. Therefore, h is a homeomorphism. □

Denote by $\beta(Y)\colon T_Y \to Y$ the morphism induced by the transverse map $B(Y)\colon T \to Y$. Hence β defines a map $\beta\colon \mathcal{X} \to \mathcal{M}(\mathcal{G}(T), \mathcal{X})$. As $B\colon \mathcal{X} \to \mathcal{T}(T, \mathcal{X})$ is continuous and the newly defined metric on T_Y depends continuously on the metric on Y, β is continuous. We can now define a homotopy equivalence from $\mathcal{X}$ to $\mathcal{G}(T)$.

Theorem 6.6 *For a finitely generated group G, any irreducible unnormalized deformation space which contains a G–tree with finitely generated vertex groups is contractible.*

Proof Let $\beta\colon \mathcal{X} \to \mathcal{M}(\mathcal{G}(T), \mathcal{X}), \Phi\colon \mathcal{M}(\mathcal{X}) \times \{(s,t) \mid 0 \le s \le t \le 1\} \to \mathcal{M}(\mathcal{X})$ and $\mathcal{R}a\colon \mathcal{M}(\mathcal{X}) \to \mathcal{X}$ be the continuous functions defined above.

Define a homotopy $H\colon \mathcal{X}\times[0,1]\to\mathcal{X}$ by $H_{(1-t)}(Y)=\mathcal{R}a(\Phi(\beta(Y),0,t))$. Then $H_0(Y)=\mathcal{R}a(\Phi(\beta(Y)),0,1))=\mathcal{R}a(\beta(Y))=Y$ and $H_1(\mathcal{X})=\mathcal{G}(T)$, which is contractible by 6.5. □

Recall that $\mathcal{D}=\mathcal{X}/\mathbb{R}^+$. As $\mathcal{R}a\circ\Phi$ is $\mathbb{R}^+$–invariant (Remark 3.4) and β clearly is also, H descends to a homotopy of $\mathcal{D}$. Therefore we have the following theorem as stated in the introduction:

Theorem 6.7 *For a finitely generated group G, any irreducible deformation space which contains a G–tree with finitely generated vertex groups is contractible.*

References

[1] **M Bestvina**, *The topology of* $\mathrm{Out}(F_n)$, Proc. of the ICM, Beijing II, Higher Ed. Press, Beijing (2002) 373–384

[2] **M Bridson**, **K Vogtmann**, *Automorphism groups of free groups, surface groups and free abelian groups*, `arXiv:math.GR/0507612`

[3] **M Culler**, **J W Morgan**, *Group actions on* **R***-trees*, Proc. London Math. Soc. (3) 55 (1987) 571–604

[4] **M Culler**, **K Vogtmann**, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986) 91–119

[5] **M Forester**, *Deformation and rigidity of simplicial group actions on trees*, Geom. Topol. 6 (2002) 219–267

[6] **M Gromov**, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. (1981) 53–73

[7] **V Guirardel**, **G Levitt**, *A general construction of JSJ splittings*, research announcement

[8] **V Guirardel**, **G Levitt**, *The outer space of a free product*, e-print (2005) `arXiv:math.GR/0501288`

[9] **D McCullough**, **A Miller**, *Symmetric automorphisms of free products*, Mem. Amer. Math. Soc. 122 (1996) viii+97pp

[10] **F Paulin**, *The Gromov topology on* **R***-trees*, Topology Appl. 32 (1989) 197–221

[11] **R Skora**, *Deformation of length functions in groups*, preprint

[12] **K Vogtmann**, *Automorphisms of free groups and outer space*, Geom. Dedicata 94 (2002) 1–31

Department of Mathematics, University of Utah
Salt Lake City, UT 84112-0090, USA

Email: `clay@math.utah.edu`

Received: 19 November 2004