DENDROLOGY OF GROUPS IN LOW $\mathbb{Q}$-RANKS

HENRI GILLET & PETER B. SHALEN

Introduction

The work of Bass and Serre [15] established the study of actions of groups on simplicial trees as a basic theme in the theory of infinite groups. From the standpoint of abstract group theory, this "simplicial dendrology" provides an efficient and elegant context for the classical theory of free products with amalgamation. From the standpoint of linear group theory, simplicial trees arise naturally as Tits buildings for rank-1 algebraic groups (such as SL_2) over discretely valued fields.

Actions of groups (by isometries) on $\mathbb{R}$-trees, or more generally on Λ-trees, where Λ is an ordered abelian group, are natural generalizations of actions on simplicial trees. (The definition of a Λ-tree is reviewed in §1 of this paper, following the point of view of [12] and [1]. The simplicial case is equivalent to the case $\Lambda = \mathbb{Z}$; see [12, p. 430].) The study of such generalized trees began with Lyndon [8] and Chiswell [3], who took the abstract group-theoretic point of view. In [8], Lyndon raised the question of which groups can act freely on $\mathbb{R}$-trees, and how the actions may be classified.

By "classifying" actions of a group Γ on $\mathbb{R}$-trees we shall mean classifying the (translation) length functions (1.16) that they define on Γ. In [1] and [4], it is shown that for a suitably nondegenerate action, the length function contains all the essential information. In this introduction we shall say that two actions of Γ on $\mathbb{R}$-trees are *equivalent* if they define the same length function. An action of a finitely generated group Γ on T defines the zero length function if and only if some point of T is fixed by all of Γ; we shall call such an action *trivial*.

The study of generalized trees from the point of view of linear group theory was initiated independently by Tits, who recognized the role of such trees as buildings for rank-1 algebraic groups over fields with indiscrete

valuations. The link between the two points of view was recognized by Alperin and Moss [2].

In [12] and [9], it was shown that degenerations of homotopy-hyperbolic structures on a manifold M are closely related to certain representations of $\Gamma = \pi_1(M)$ in rank-1 algebraic groups over indiscretely valued fields, and hence to certain actions of Γ on generalized trees. This connection between dendrology and hyperbolic geometry was further exploited in [13].

The actions of a group Γ on a Λ-tree T that are relevant to the theory of [12] and [9] are not in general free actions; rather they have the property that the stabilizer of every nondegenerate segment (1.1) in T is a *small* subgroup of Γ, i.e., one that contains no nonabelian free subgroup. In this introduction we shall call such an action *small*. According to the theory of [9], the space $\mathrm{DF}_n(\Gamma)$ of characters of discrete faithful representations of Γ in $\mathrm{SO}(n,1)$ has a natural compactification in which the ideal points are defined by the length functions associated with nontrivial small actions of Γ on $\mathbb{R}$-trees.

Each of these actions is in fact a completion (1.10) of an action of Γ on a Λ-tree, where $\Lambda \subset \mathbb{R}$ is a group whose $\mathbb{Q}$-rank—i.e., the dimension of the $\mathbb{Q}$-vector space $\Lambda \otimes_{\mathbb{Z}} \mathbb{Q}$—is finite. Indeed, Λ arises as a subgroup of the value group of a valuation v on a finite extension F of the function field $\mathbb{R}(X)$, where X is the variety of characters of discrete faithful representations of Γ in $\mathrm{SO}(n,1)$, and $v|\mathbb{R} = 0$. The $\mathbb{Q}$-rank of Λ is thus at most $\dim_{\mathbb{R}}(X)$, the transcendence degree of F over $\mathbb{R}$.

In this introduction we shall say that an action of Γ on an $\mathbb{R}$-tree T is *defined over* a subgroup $\Lambda \subset \mathbb{R}$ if it is the completion of an action on some Λ-tree. In this case the length function takes values in Λ. We shall say that an action *has rank* $\leq r$ if it is defined over some $\Lambda \subset \mathbb{R}$ with $\mathbb{Q}$-rank $\leq r$.

The theory of [9] shows that it is important to be able to classify the finite-rank small actions of finitely generated groups on $\mathbb{R}$-trees. For example, if a given finitely generated Γ admits no nontrivial finite-rank small action on any $\mathbb{R}$-tree, then $\mathrm{DF}_n(\Gamma)$ is compact for any n.

A question related to the above questions about free and small actions, and of interest from the abstract group-theoretical point of view, is to characterize the groups that admit nontrivial actions on $\mathbb{R}$-trees.

In the simplicial case, i.e., the case of actions defined over $\mathbb{Z}$, answers to the above questions are accessible. It has long been known (and is included in the Bass-Serre theory) that a group can act freely on the real completion of a simplicial tree if and only if it is a free group. The Bass-Serre theory also answers the questions of which groups admit nontrivial actions, and which ones admit nontrivial small actions. Let us say that

a group Γ *splits* over a subgroup Δ if Γ is either a free product of two proper subgroups with the amalgamated subgroup Δ, or an HNN group with associated subgroup Δ. Then a group admits a nontrivial (small) action on some $\mathbb{Z}$-tree if and only if it splits over some (small) subgroup. Similarly, the Bass-Serre theory may be used to analyze the different free (or nontrivial, or small) actions of a given group.

Thus it is both natural and useful to try to develop a structure theory similar to the Bass-Serre theory for actions on Λ-trees, in particular for the case where Λ is a subgroup of $\mathbb{R}$ with finite $\mathbb{Q}$-rank. In this paper we develop such a theory for the case where the $\mathbb{Q}$-rank is at most 2. We show that our theory is strong enough to be applied to all the problems discussed above. Before discussing our main theorem, we wish to discuss its consequences with regard to these problems.

We begin with the problem about free actions. We restrict attention to finitely generated groups; the examples in [2] show that very exotic infinitely generated groups can admit free actions.

It is easy to show that the class of groups which act freely and without inversions on $\mathbb{R}$-trees is closed under the formations of free products. Since a free abelian subgroup of $\mathbb{R}$ acts on $\mathbb{R}$ (which is itself an $\mathbb{R}$-tree) by translations, it follows that any free product of finitely generated free abelian groups acts freely on some $\mathbb{R}$-tree. These are the examples pointed out in [8]. On the other hand, it is shown in [14] that most surface groups (i.e., fundamental groups of closed 2-manifolds) admit free actions on $\mathbb{R}$-trees.

It is an intriguing question whether every finitely generated group Γ which acts freely and without inversions on an $\mathbb{R}$-tree is a free product of free abelian groups and surface groups. In [13] it was shown that this question has an affirmative answer if one adds the hypothesis that Γ is the fundamental group of a 3-manifold (possibly with boundary). Likewise, in [10] it was shown that the answer is affirmative if Γ is assumed to be a free product of two free groups with a cyclic amalgamated subgroup. The following result gives an affirmative answer in the case where the action has rank ≤ 2.

Theorem A. *Let Λ be a subgroup of $\mathbb{R}$, and let Γ be a finitely generated group which acts freely and without inversions (1.16) on some Λ-tree. If Λ has $\mathbb{Q}$-rank 1 then Γ is free. If Λ has $\mathbb{Q}$-rank 2 then Γ is a free product of infinite cyclic groups and surface groups.*

(The condition that Γ act freely and without inversions on T is equivalent (1.19) to the assumption that the action may be completed to a free action on an $\mathbb{R}$-tree.)

This theorem will be proved in two stages in the present paper. The rank-1 case is included in Proposition 2.2, and the rank-2 case is Proposition 6.3.

Theorem A leaves open the problem of classifying the rank-1 and rank-2 free actions of the groups appearing in the conclusion of the theorem. Two obvious special cases of this question are the cases where Γ is a free group or a surface group.

If Γ is a free group, the only known free actions of rank ≤ 2 of Γ on an $\mathbb{R}$-tree T are those which are essentially simplicial, i.e., equivalent to actions which are topologically conjugate to simplicial actions on simplicial trees. It is not clear from our structure theorem whether all free actions of rank ≤ 2 are essentially simplicial; our theorem shows only that they are direct limits, in the strong sense described later in this introduction, of essentially simplicial actions. Bestivina and Handel have given examples of rank-3 actions of free groups that are not essentially simplicial.

On the other hand, in the case of a surface group, we have a complete description, in terms of Thurston's theory of measured foliations, of all the length functions defined by free actions—or even by nontrivial *small* actions—of rank ≤ 2. Thurston's theory is reviewed in §5, where it is shown that every measured foliation μ on a surface Σ defines an $\mathbb{R}$-tree T, the *dual tree* of μ, on which $\Gamma = \pi_1(\Sigma)$ acts in a natural way. The action of Γ on T is always small and nontrivial. (Since any small infinite-index subgroup of a surface group Γ is cyclic, a nontrivial action of Γ is small if and only if the stabilizer of every segment is cyclic.)

Theorem B. *Let Σ be a closed surface, let Λ be a subgroup of $\mathbb{R}$ with $\mathbb{Q}$-rank ≤ 2, and consider an arbitrary small, nontrivial action of $\Gamma = \pi_1(\Sigma)$ on a Λ-tree T. Then the induced action of Γ on the $\mathbb{R}$-completion of T is equivalent to its action on the dual tree of some measured foliation on Σ.*

This theorem is also proved in two stages in the paper. In the case where Λ has $\mathbb{Q}$-rank 1, Proposition 2.11 asserts that up to equivalence the action of Γ on T is defined by a system of disjoint simple closed curves on the surface, which is equivalent to a "rational measured foliation" in Thurston's sense. Proposition 6.10 is the rank-2 case of Theorem B.

It is a plausible conjecture[1] that Theorem B holds without the restriction to rank ≤ 2, i.e., that every nontrivial small action of $\pi_1(\Sigma)$ on an $\mathbb{R}$-tree is defined by a measured foliation. In addition to its connection with Lyndon's problem, this conjecture is also significant from the point of view of hyperbolic geometry. A celebrated result of Thurston's is equivalent to the

---

[1]*Note added in proof.* R. Skora has recently proved this conjecture.

assertion that the actions on $\mathbb{R}$-trees that describe the ideal points of the compactification of the Tiechmüller space $DF_2(\Gamma)$, where Γ is a surface group, are defined by measured foliations. The conjecture would provide a purely group-theoretical explanation of this, given the general theory of [12]. (The arguments used in [12] to derive Thurston's result from the general theory are somewhat indirect; the same is true of the alternative arguments in [11].) More importantly, the conjecture would provide information analogous to Thurston's theorem for the compactification of $DF_n(\Gamma)$, where Γ is a surface group and $n > 2$.

We next turn to the problem of determining which groups admit nontrivial small actions on Λ-trees. It is conceivable that the answer for $\Lambda = \mathbb{R}$ is the same as for $\Lambda = \mathbb{Z}$, i.e., that a group Γ which admits a nontrivial small action on an $\mathbb{R}$-tree always splits over a small subgroup. In [13], this conjecture was shown to be true in the case where Γ is the fundamental group of a 3-manifold; this implies Thurston's result that for any group Γ which does not split over a small subgroup, $DF_3(\Gamma)$ is compact. As the following result shows, the conjecture is also true in the the case where the action has rank ≤ 2, provided that Γ satisfies some additional mild conditions.

Theorem C. *Let Γ be a finitely presented group. Let Λ be a subgroup of $\mathbb{R}$, and suppose that Γ admits a nontrivial small action on some Λ-tree. Furthermore, suppose that either*

   (i) *Λ has $\mathbb{Q}$-rank 1, or*
   (ii) *Λ has $\mathbb{Q}$-rank 2, and every small subgroup of Γ is finitely-generated.*

Then Γ splits over a small subgroup.

The rank-1 case is included in Proposition 2.7, and the rank-2 case is Proposition 6.9.

The assumption that every small subgroup of Γ is finitely generated is a harmless one for applications to hyperbolic geometry, since it holds whenever $DF_n(\Gamma) \neq \varnothing$. On the other hand, the assumption that Γ is finitely presented is restrictive. It seems quite possible that the theorem would become false if Γ were assumed only to be finitely generated rather than finitely presented.

Now consider the problem of characterizing those finitely presented groups that admit nontrivial actions on $\mathbb{R}$-trees. It is natural to ask whether the answer for $\mathbb{R}$-trees is the same as for $\mathbb{Z}$-trees: that is, whether every group that admits a nontrivial action on an $\mathbb{R}$-tree splits over some subgroup. Theorem D below answers this affirmatively in the rank-1 case, but gives only a partial answer in the rank-2 case. Let us say that an action

of a group Γ on a Λ-tree T satisfies the *ascending chain condition* if for every nested sequence $I_1 \supset I_2 \supset \cdots$ of segments of T with a common midpoint, the corresponding sequence of stabilizers $G_1 \subset G_2 \subset \cdots \subset \Gamma$ is eventually constant, i.e., $G_i = G_{i+1}$ for all sufficiently large i.

Theorem D. *Let Λ be a subgroup of $\mathbb{R}$, and let Γ be a finitely generated group which admits a nontrivial action on some Λ-tree T. Suppose that either*

   (i) Λ *has $\mathbb{Q}$-rank 1, or*
  (ii) Λ *has $\mathbb{Q}$-rank 2 and the action of Γ on T satisfies the ascending chain condition.*

Then Γ splits over some subgroup.

The rank-1 case of this result is again included in Proposition 2.7; the rank-2 case is Proposition 6.8.

The above results are derived from a general structure theorem for actions of rank ≤ 2. The structure theorem is stated and proved in two parts: the rank-1 case is Theorem 2.1 and the rank-2 case is Theorem 6.2. Theorem 2.1 asserts that an arbitrary action of rank 1 is a direct limit, in a very strong sense (see §1) of simplicial actions. This implies, for example, that if l is the length functions defined by a rank-1 action of a finitely presented group Γ, then there exists a sequence $(l_i)_{i \geq 0}$ of $\mathbb{Z}$-valued length functions defined by simplicial actions, and a sequence $(n_i)_{i \geq 0}$ of positive integers, such that for any given $g \in \Gamma$ we have $l(g) = l_i(g)/n_i$ for all sufficiently large i.

The rank-2 version of the structure theorem applies to rank-2 actions satisfying the ascending chain condition explained above. It asserts that such an action is a direct limit, in the same strong sense, of actions defined in a concrete geometric way. These geometrically defined actions are a common generalization of simplicial actions on the actions defined by measured foliations on surfaces. They arise from "measured foliations", in a suitably generalized sense, on spaces that we call "singular surfaces": these are simplicial complexes in which each 1-simplex is incident either to exactly two 2-simplices or to none.

The main technique that we use in the proofs of the structure theorems is to associate complexes with trees. Given a Λ-tree T, where Λ is a subgroup of $\mathbb{R}$, we construct for each rank-n free abelian subgroup L of Λ a simplicial n-complex $X(L)$. The vertices of $X(L)$ are the points of T, and a 1-simplex is defined by a pair of points whose distance is a unimodular element of L (i.e., an element of some basis of L). A 2-simplex is determined by a triple $\{x, y, z\}$, where y lies on the segment

joining x to z, and where the distances $d(x, y)$ and $d(y, z)$ form a basis of L.

If Λ has $\mathbb{Q}$-rank n we may write it as a monotone union of rank-n free abelian groups L_i. The corresponding complexes $X_i = X(L_i)$ form a direct system of n-complexes. In the case $n = 1$ we obtain a direct system of simplicial trees $\widetilde{X}_i$ from the X_i by a universal covering construction. We can interpret $\widetilde{X}_i$ as an L_i-tree, and the direct system consisting of the $\widetilde{X}_i$ converges strongly to T.

When Λ has $\mathbb{Q}$-rank 2, we show that any subcomplex Y of X_i satisfying suitable finiteness condition can be mapped in a natural way into a Γ_i-invariant singular surface $\Sigma \subset X_i$. The metric on T induces a natural measured foliation on Σ. By varying i and $Y \subset X_i$ we obtain a direct system of singular surfaces with measured foliations; these give rise to the "geometrically defined trees" that were mentioned in the above discussion of the structure theorem.

The assumption that Λ has $\mathbb{Q}$-rank 2 is used in two important ways in the above argument. First, it is an easy consequence of the definition of the X_i that for any 1-simplex τ in X_i there are exactly two 2-simplices of which τ is the "longest" edge; here the "length" of a 1-simplex is the distance between the two points of T that define it. This property of the X_i is the starting point for the construction that "replaces" the subcomplex Y of X_i by the singular surface Σ. For $n > 2$, there is no natural way to construct complexes X_i having this property.

The second way in which the restriction $n = 2$ is used is in the proof that the trees associated with the $\widetilde{\Sigma}_j$ converge strongly to T; this depends on showing (9.15) that the fundamental groups of the complexes X_i (with any consistent choice of base points) converge to the trivial group. The proof of the latter fact uses the classical theory of continued fractions, which deals with the structure of rank-2 free abelian subgroups of $\mathbb{R}$, and does not have a straightforward generalization to higher rank.

In any case it is not clear what the correct generalization of the structure theorem to rank > 2 would be. Indeed, if Theorem 2.1 went through without change when Λ had $\mathbb{Q}$-rank > 2, the proof of 6.3 would show that every finitely generated group that acts freely on a Λ-tree was a free product of cyclic groups and surface groups. But this is false when Λ has $\mathbb{Q}$-rank > 2. Indeed, Λ may itself be regarded as a Λ-tree, and any finitely generated subgroup of Λ acts freely on Λ by translations; but a free abelian group of rank > 2 is not a free product of cyclic groups and surface groups. (A free abelian group of rank 2 is the fundamental group of a 2-torus.)

In §1 we introduce a category of trees and study the notion of a strong direct limit. In §2 we state the structure theorem for the rank-1 case and show that it implies Theorems A, B, C and D in this case. §§3 and 4 are devoted to laying the foundations for a theory of complexes associated with trees; §4 concludes with the proof of the structure theorem for the rank-1 case.

In §5 we define singular surfaces and measured foliations, and construct the tree associated with a measured foliation. §6 contains the statement of the structure theorem for the rank-2 case and the proof that it implies the rank-2 versions of Theorems A, B, C and D. §§7–10 are preparation for the proof of the rank-2 structure theorem, which is given in §11. §7 shows how certain maps from singular surfaces to trees define measured foliations; §8 contains a crucial result on rank-2 free abelian subgroups of $\mathbb{R}$, which is a slight refinement of a standard theorem on continued fractions; §9 contains the basic properties of the complexes X_i in the rank-2 case; and §10 gives the argument for replacing a subcomplex Y of X_i be a singular surface with a natural measured foliation.

We are indebted to J. Wood for providing us with the proof of a crucial special case of Lemma 2.3; to A. O. L. Atkin for pointing out Theorem 175 of [6]; and to H. Masur and S. Kerckhoff for their patience in explaining the theory of measured foliations to us. We have also profited from conversations with J. Baldwin, H. Bass, M. Culler, G. Mess, J. Morgan, L. Mosher, M. Ronan, J. Smillie and K. Vogtmann.

1. Strong convergence

1.1.   Let Λ be any ordered abelian group. Recall from [12] that a Λ-*metric space* is a set X with a distance function $d: X \times X \to \Lambda$ satisfying the usual formal axioms for a metric space.

Let X and X' be Λ-metric spaces. A map of sets $f: X \to X'$ will be called an *isometric embedding* if it preserves distances (and is therefore injective). If f is bijective it will be called an *isometry*.

Let X be a Λ-metric space. Recall from [12] that a *(closed) segment* in X is a subset which is isometric to a (possibly degenerate) closed interval in Λ. A segment has two well-defined endpoints, or one endpoint if it is degenerate. Recall that a Λ-*tree* is a nonempty Λ-metric space T satisfying the following axioms.

(T1) For any two points x, $y \in X$ there is a unique closed segment with endpoints x and y.

(T2) If two segments σ, $\sigma' \subset T$ have a common endpoint x, then $\sigma \cap \sigma'$ is a segment having x as one endpoint.

(T3) If the segments σ σ' intersect precisely in the common endpoint x, then $\sigma \cup \sigma'$ is a segment.

If X is an arbitrary Λ-metric space with distance function d, and if x and y are points of X, we shall denote by $[x, y] = [x, y]_X$ the set of all points $z \in X$ satisfying $d(x, z) + d(z, y) = d(x, y)$. If x and y are points of a Λ-tree T, one checks immediately that $[x, y]$ is the unique segment with endpoints x and y.

1.2. The following characterization of Λ-trees will be useful in §5.

Proposition. *Let T be a Λ-metric space. Let $\mathfrak{S}$ be a set of segments in T such that the following conditions hold.*

(1) *For any two points x, $y \in T$, there is a unique segment $S(x, y)$ which belongs to $\mathfrak{S}$ and has endpoints x and y.*

(2) *For any three points x, y, $z \in T$ we have $S(x, y) \cap S(x, z) = S(x, w)$ for some $w \in T$.*

(3) *If x, y, $z \in T$ satisfy $S(x, y) \cap S(x, z) = \{x\}$, then $S(x, y) \cup S(x, z) = S(y, z)$.*

Then T is a Λ-tree and $S(x, y) = [x, y]$ for all x, $y \in T$. (Thus every segment in T belongs to $\mathfrak{S}$.)

If $\Lambda = \mathbb{R}$, then condition (2) may be replaced by the following condition.

$(2_{\mathbb{R}})$ *Every subsegment of a segment in $\mathfrak{S}$ belongs to $\mathfrak{S}$.*

Proof. To prove the first assertion we need only show that if $\mathfrak{S}$ satisfies (1)–(3) then $S(x, y) = [x, y]$ for all x, $y \in T$; it will then follow from the definition that T is a Λ-tree. By condition (1) we have $S(x, y) \subset [x, y]$. Now consider any point $z \in [x, y]$. By (2) we have $S(x, z) \cap S(z, y) = S(z, w)$ for some $w \in T$. It follows from (3) that $S(x, y) = S(x, w) \cup S(w, y)$. In particular $w \in S(x, y)$. By (1) we have $d(x, y) = d(x, w) + d(w, y)$. If $w \neq z$ it then follows that $d(x, y) < d(x, z) + d(z, y)$, contradicting the assumption that $z \in [x, y]$. Hence $w = z$ and therefore $z \in S(x, y)$. This shows that $[x, y] = S(x, y)$, and the first assertion is proved.

Now suppose that $\Lambda = \mathbb{R}$ and that $\mathfrak{S}$ satisfies (1) and $(2_{\mathbb{R}})$. We shall complete the proof of the proposition by showing that $\mathfrak{S}$ satisfies (2) as well. Let x, y and z be points of T. Then $J = S(x, y) \cap S(x, z)$ is a compact subset of $S(x, y)$ containing x. If u and v are any two points of J, then by (1) and $(2_{\mathbb{R}})$ we have $S(u, v) \subset J$, so that J is connected. Hence J is a segment having x as an endpoint; by $(2_{\mathbb{R}})$ we have $J \in S$, and (2) is established.

1.3. It will be convenient to work with a somewhat more general class of Λ-metric spaces than Λ-trees. These spaces will turn out to be canonically isometric to subsets of Λ-trees (see 1.10 and 1.14).

A Λ-metric space T is called a Λ-*pretree* if it satisfies the following axioms.

(PT1) For any two points x, $y \in T$, the set $[x, y]$ is isometric to a subset of an interval in Λ.

(PT2) For any three points x, y, $z \in T$ we have $[x, y] \cap [x, z] = [x, w]$ for some $w \in T$.

(PT3) If x, y, $z \in T$ satisfy $[x, y] \cap [x, z] = \{x\}$, then $[x, y] \cup [x, z] = [y, z]$.

It follows from (PT1) that for any x, $y \in T$ there is an isometry of $[x, y]$ onto a subset of an interval $I \subset \Lambda$ mapping x and y onto the endpoints of I. We call $[x, y]$ the *presegment with endpoints x and y*.

It is not hard to show that (PT1) is redundant, i.e., that it follows from (PT2) and (PT3). We shall not need this fact.

It is easy to show that the intersection of two presegments in a pretree is either the empty set or a presegment.

Any Λ-tree is a Λ-pretree. If Λ_0 is a subgroup of Λ, then any Λ_0-pretree, and in particular any Λ_0-tree, is clearly a Λ-pretree.

Let T be a Λ-pretree and let T' be a subset of T. We may regard T' as a Λ-metric space. If T' is itself a Λ-pretree we call it a *subpretree* of T.

1.4. Proposition. *Let T' be a subset of a Λ-pretree T. Then the following conditions are equivalent.*

(i) *T' is a subpretree of T.*

(ii) *For any three points x, y, $z \in T'$, we have $[x, y]_T \cap [x, z]_T = [x, w]_T$ for some $w \in T'$.*

Proof. Note that for any x, $y \in T'$ we have $[x, y]_T \cap T' = [x, y]_{T'}$. It follows that any subset T' of T satisfies axiom (PT1). Now suppose that (ii) holds. Then for any x, y, $z \in T'$ there is a point $w \in T'$ such that $[x, y]_T \cap [x, z]_T = [x, w]_T$; hence

$$[x, y]_{T'} \cap [x, z]_{T'} = [x, y]_T \cap [x, z]_T \cap T' = [x, w]_T \cap T' = [x, w]_{T'}.$$

This establishes (PT2). If $[x, y]_{T'} \cap [x, z]_{T'} = \{x\}$, then, defining w as in the argument just given, we have $[x, w]_{T'} = \{x\}$ and hence $w = x$. Therefore $[x, y]_T \cap [x, z]_T = \{x\}$. Since T is a Λ-pretree it follows that

$[x, y]_T \cup [x, z]_T = [y, z]$. Hence

$$[x, y]_{T'} \cup [x, z]_{T'} = ([x, y]_T \cup [x, z]_T) \cap T' = [y, z]_{T'}.$$

This establishes (PT3) and completes the proof that (ii) $\Rightarrow$ (i).

Conversely, suppose that (i) holds, and let x, y, z be points of T'. Set $\alpha = d(x, y) + d(x, z) - d(y, z)$. The pretree axioms, applied to T, show that $[x, y]_T \cap [x, z]_T = [x, w]_T$, where w is a point of T such that $2d(x, w) = \alpha$. Similarly, by applying the pretree axioms to T', we obtain a point $w \in T'$ such that $[x, y]_{T'} \cap [x, z]_{T'} = [x, w]_{T'}$ and $2d(x, w') = \alpha$. Since in particular w and w' lie in $[x, z]_T$, which can be isometrically identified with a subset of an interval in Λ in such a way that x is an endpoint, it follows that $w = w'$. Hence (i) $\Rightarrow$ (ii). q.e.d.

The following two simple facts about pretrees will be needed later in this section.

1.5. Proposition. *Let S be a subset of a Λ-pretree T, and let x_1 and x_2 be points of T. Suppose that each x_i lies in a presegment σ_i whose endpoints are in S. Then $[x_1, x_2]$ is contained in a presegment whose endpoints are in S.*

Proof. Let y_i and z_i denote the endpoints of σ_i. It follows easily from the axioms for a Λ-pretree that at least one of the presegments $[x_1, y_1]$ and $[x_1, z_1]$ meets $[x_1, x_2]$ only in the endpoint x_1. After changing notation if necessary we may assume that $[x_1, y_1] \cap [x_1, x_2] = \{x_1\}$. Similarly we may assume that $[x_2, y_2] \cap [x_1, x_2] = \{x_2\}$. The Λ-pretree axoms now imply that $[y_1, y_2] = [y_1, x_1] \cup [x_1, x_2] \cup [x_2, y_2]$. In particular $[x_1, x_2] \subset [y_1, y_2]$. But y_1 and y_2 belong to S. q.e.d.

1.6. Proposition. *Let σ and σ' be presegments in a Λ-pretree T such that $\sigma \cap \sigma' \neq \varnothing$. Then $\sigma \cap \sigma'$ is a presegment. If the endpoints of σ and σ' lie in a subpretree T' of T, then the endpoints of $\sigma \cap \sigma'$ also lie in T'.*

Proof. Set $\sigma = [x_1, x_2]$ and $\sigma' = [y_1, y_2]$. Choose a point $z \in \sigma \cap \sigma'$. For $i = 1, 2$, set $\sigma_i = [x_i, z]$ and $\sigma'_i = [y_i, z]$. Then $\sigma \cap \sigma'$ is the union of the four sets $\sigma_i \cap \sigma'_j$, each of which—according to axiom (PT2)—is a presegment contained in σ and having z as an endpoint. Using the isometric identification of σ with an interval in Λ, we conclude that $\sigma \cap \sigma'$ is a presegment.

Now let w be an endpoint of $\sigma \cap \sigma'$. Then $\sigma \cap \sigma'$ is contained either in $[x_1, w]_T$ or in $[w, x_2]_T$. After changing notation if necessary we may assume that $\sigma \cap \sigma' \subset [x_1, w]_T$. Similarly we may assume that $\sigma \cap \sigma' \subset [y_1, w]_T$. The pretree axioms then imply that $[x_1, y_2] \cap [x_1, x_2] = [x_1, w]$. By 1.3 it now follows that if the x_i and y_i lie in a subpretree T' of T then w also lies in T'.

1.7. Let T and T' be Λ-pretrees. A map of sets $f: T \to T'$ will be called a *morphism* from T to T' if each presegment in T can be written as a finite union of presegments $[x_0, x_1], \cdots, [x_{n-1}, x_n]$ in such a way that $f|[x_{i-1}, x_i]: [x_{i-1}, x_i] \to T'$ is an isometric embedding of Λ-metric spaces for $i = 1, \cdots, n$. (We may suppose the x_i to be chosen so that $[x_{i-1}, x_i] \cap [x_i, x_{i+1}] = \{x_i\}$.) With this definition of morphism, Λ-pretrees form a category. A map between two Λ-trees will be called a morphism of Λ-trees if it is a morphism of Λ-pretrees; thus Λ-trees form a full subcategory of the category of Λ-pretrees.

Note that our notion of morphism differs from the one used in [1], where morphisms are defined to be isometric embeddings. An isometric embedding of Λ-pretrees is obviously a morphism in our sense, but not conversely. On the other hand, a morphism is always a distance-decreasing map; that is, a morphism $f: T \to T'$ satisfies $d(f(x), f(y)) \leq d(x, y)$ for all $x, y \in T$. (Here and elsewhere we write d for the distance function on a Λ-pretree if it is clear from the context which Λ-pretree is involved.)

It follows that the isomorphisms in the category of Λ-pretrees are isometries. Hence when we say that a group acts on a Λ-pretree, it will always be understood that it acts by isometries. The full group of self-isometries of a Λ-pretree T will be denoted $\operatorname{Aut}(T)$.

Note that if $f: T \to T'$ is a morphism of Λ-trees and if $\sigma \subset T$ is a segment such that $f|\sigma$ is 1-1, then $f|\sigma$ is an isometric embedding.

1.8. The following characterization of morphisms will be useful in §7.

Proposition. *Let T_0 and T_1 be Λ-pretrees, and let $f: T_0 \to T_1$ be a map of sets. Suppose that for any two points $x, y \in T$ there exist points $x = x_0, \cdots, x_n = y$ such that $f|[x_{i-1}, x_i]$ is an isometric embedding for $i = 1, \cdots, n$. Then f is a morphism of Λ-pretrees.*

Proof. We must show that for any $x, y \in T$ we may express $[x, y]$ as a finite union of presegments on which f restricts to an isometric embedding. We choose $x_0, \cdots, x_n$ as in the hypothesis of the proposition and argue by induction on n. If we set $z = x_{n-1}$, then the pretree axioms imply that $[x, z] \cap [z, y] = [z, w]$ for some $w \in T$, and that $[x, y] = [x, w] \cup [w, y]$. By the induction hypothesis, $[x, z]$ is a finite union of presegments on which f restricts to an isometric embedding; hence the same is true of $[x, w]$. On the other hand, since $f|[z, y]$ is an isometric embedding, so is $f|[w, y]$. The assertion follows.

1.9. In [1], Alperin and Bass study Λ-trees in terms of base points. Given a Λ-metric space X and a base point $x_0 \in X$, they define a function

$\wedge\colon X \times X \to \frac{1}{2}\Lambda$ by $x \wedge y = \frac{1}{2}(d(x, x_0) + d(y, x_0) - d(x, y))$. According to Theorem 3.17 of [1], X is isometric to a subspace of a Λ-tree if and only if the following conditions hold:

   (a) $x \wedge y \in \Lambda$ for all x, $y \in X$; and

   (b) $x \wedge z \geq \min(x \wedge y, y \wedge z)$ for all x, y, $z \in X$.

The proof given in [1] of the necessity of (a) and (b) works equally well if X is a subspace of a Λ-pretree. In particular, if T is a Λ-pretree, and if we choose a base point $x_0 \in T$ and define $\wedge\colon T \times T \to \frac{1}{2}\Lambda$ as above, then (a) and (b) hold. It follows that T is isometric to a subspace of a Λ-tree $\overline{T}$. Furthermore, the proof of Theorem 3.17 of [1] gives an explicit construction of such a Λ-tree $\overline{T}$; and it follows from the construction that every point of $\overline{T}$ lies in a segment with endpoints in T. This proves the following result.

   1.10. **Proposition.** *Let T be a Λ-pretree. Then there exist a Λ-tree $\overline{T}$ and a map $j\colon T \to \overline{T}$ such that*

   (i) *j is an isometric embedding of Λ-metric spaces, and*

   (ii) *every point of $\overline{T}$ lies in a segment whose endpoints are in $j(T)$.*

A pair $(\overline{T}, j)$ consisting of a Λ-tree and a map having the above properties will be called a Λ-*completion* of T. (We shall say that $\overline{T}$ is a Λ-completion when it is clear which map j is involved.)

   1.11. **Proposition.** *Let $(\overline{T}, j)$ be a completion of a Λ-pretree T. Then for any Λ-tree T' and any morphism $\psi\colon T \to T'$ of Λ-pretrees, there is a unique morphism $\overline{\psi}\colon \overline{T} \to T'$ of Λ-trees such that $\overline{\psi} \circ j = \psi$.*

   1.12. The proof of 1.11 will be based on the following result.

   Lemma. *Let $(\overline{T}, j)$ be a completion of a Λ-pretree T. Then any segment in $\overline{T}$ is contained in a segment whose endpoints are in $j(T)$.*

   Proof. Using the isometry j, we identify T with a Λ-metric subspace of $\overline{T}$. Consider a segment $[x_1, x_2]_{\overline{T}}$. By the definition of a completion, each x_i lies in a segment σ_i whose endpoints are in T. We now apply 1.5, letting T and $\overline{T}$ play the roles of S and T respectively, to conclude that $[x_1, x_2]_{\overline{T}}$ is contained in a segment whose endpoints are in T. q.e.d.

   1.13. *Proof of* 1.11. Using j, we identify T with a Λ-metric subspace of $\overline{T}$.

Consider an arbitrary point $x \in \overline{T}$. Note that since $\overline{T}$ is a completion of T and ψ is a morphism, there is a segment $\sigma \subset \overline{T}$, whose endpoints y and z lie in T, such that $x \in \sigma$, and such that $\psi|\sigma \cap T$ is an isometric embedding. There is a unique isometry $\overline{\psi}_\sigma$ of σ onto the segment $[\psi(y), \psi(z)] \subset T'$ which extends $\psi|\sigma \cap T$. If $\overline{\psi}\colon T \to T'$ is a morphism extending ψ, then since $\overline{\psi}$ is distance-decreasing we must have $\overline{\psi}|\sigma = \overline{\psi}_\sigma$, and hence $\overline{\psi}(x) = \overline{\psi}_\sigma(x)$. This proves uniqueness.

To prove existence, we again consider an arbitrary point $x \in \overline{T}$, and choose a segment σ as above. We wish to set $\overline{\psi}(x) = \overline{\psi}_\sigma(x)$. We must check that $\overline{\psi}$ is well defined.

For $i = 1, 2$, let σ_i be a segment whose endpoints y_i and z_i lie in T, such that $x \in \sigma_i$, and such that $\psi|\sigma_i \cap T$ is an isometric embedding. By 1.3, $\sigma_3 = \sigma_1 \cap \sigma_2$ is a segment whose endpoints y_3 and z_3 lie in T. For $1 \le i \le 3$, there is a unique isometry $\overline{\psi}_i$ of σ_i onto the segment $[\psi(y_i), \psi(z_i)] \subset T'$ which extends $\psi|\sigma_i \cap T$. Clearly $\overline{\psi}_1|\sigma_3 = \overline{\psi}_3 = \overline{\psi}_2|\sigma_2$. In particular $\overline{\psi}_1(x) = \overline{\psi}_2(x)$, so that $\overline{\psi}$ is well defined.

If $x \in T$ we can take the segment σ appearing in the definition of $\overline{\psi}(x)$ to be $\{x\}$; hence $\overline{\psi}|T = \psi$. It remains to show that $\overline{\psi}$ is a morphism. To do this, we consider an arbitrary segment $\sigma \subset \overline{T}$. By 1.12, σ is contained in a segment σ' whose endpoints x, y lie in T. Since ψ is a morphism we may write $[x, y]_T = \sigma' \cap T$ in the form $[x_0, x_1]_T \cup \cdots \cup [x_{n-1}, x_n]_T$, where $x_0 = x$, $x_n = y$, $[x_{i-1}, x_i]_T \cap [x_i, x_{i+1}]_T = \{x_i\}$ for $0 < i < n$, and $\psi|[x_{i-1}, x_i]_T$ is an isometric embedding for $1 \le i \le n$. Now $\sigma' = [x_0, x_1]_{\overline{T}} \cup \cdots \cup [x_{n-1}, x_n]_{\overline{T}}$; and the definition of $\overline{\psi}$ says that $\overline{\psi}|[x_{i-1}, x_i]_{\overline{T}}$ is an isometric embedding for $1 \le i \le n$. This proves that $\overline{\psi}$ is a morphism.   q.e.d.

1.14.   It follows from 1.11 that the completion of a Λ-pretree T is unique in the sense that if $(\overline{T}_1, j_1)$ and $(\overline{T}_2, j_2)$ are completions then there is a unique isometry $i: \overline{T}_1 \to \overline{T}_2$ such that $i \circ j_1 = j_2$. The Λ-completion of a Λ-pretree T will be denoted ΛT. If T is a Λ_0-tree, where Λ_0 is a subgroup of Λ (cf. 1.3), we may say that ΛT is obtained from T by *base change*. If one compares 1.11 with Proposition 4.4 of [1], it is clear that this is consistent with the terminology of [1].

If T is a Λ-pretree, we shall always understand T to be identified with a subpretree of ΛT via the map j. It follows from 1.11 that any morphism $f: T \to T'$ of Λ-pretrees extends uniquely to a morphism $\Lambda f: \Lambda T \to \Lambda T'$. (Thus $T \mapsto \Lambda T$ is a functor from the category of Λ-pretrees to the category of Λ-trees.)

1.15.   In studying group actions on trees, it will be necessary to consider group actions on objects in various other categories. Furthermore, it will often be necessary to compare actions of different groups on different objects within a single category. The following terminology will be convenient.

Let $\mathscr{C}$ be a category. By an *object with symmetry in* $\mathscr{C}$ we mean a triple $\mathscr{X} = (X, \Gamma, \rho)$, where X is an object in $\mathscr{C}$, Γ is a group and

ρ is a left action of Γ on X, i.e., a homomorphism from Γ to the group of automorphisms of X.

Let $\mathscr{X} = (X, \Gamma, \rho)$ and $\mathscr{X}' = (X', \Gamma', \rho')$ be objects with symmetry in $\mathscr{C}$. A *morphism* from $\mathscr{X}$ to $\mathscr{X}'$ is a pair (f, ω), where $\omega: \Gamma \to \Gamma'$ is a homomorphism and $f: X \to X'$ is a morphism such that $f \circ \rho(g) = \rho'(\omega(g)) \circ f$ for all $g \in \Gamma$; here $\circ$ denotes composition in the category $\mathscr{C}$. If $\Gamma = \Gamma'$ and ω is the identity, we shall often say that the map f is Γ-*equivariant*.

With the above definition of morphism, the objects with symmetry in $\mathscr{C}$ form a category. For example, we have the category of sets with symmetry, the category of Λ-trees with symmetry and the category of Λ-pretrees with symmetry.

1.16. Recall from [12] and [1] that a group Γ is said to act *without inversions* on a Λ-tree T if whenever a segment $[x, y] \subset T$ is invariant under an element g of Γ, either g fixes $[x, y]$ pointwise, or $d(x, y)$ is divisible by 2 in Λ (in which case g fixes the midpoint of $[x, y]$). It is shown in [12] that if Γ acts on T without inversions then $l(g) = \min_{x \in T} d(x, g \cdot x) \in \Lambda$ exists for every $g \in \Gamma$. As in [12], we call the function $l: \Gamma \to \Lambda$ the *length function* defined by the action of Γ on T.

1.17. If $\mathscr{T} = (T, \Gamma, \rho)$ is a Λ-pretree with symmetry, it follows from 1.14 that the action of Γ on T extends uniquely to an action on ΛT, so that we have a Λ-tree with symmetry $(\Lambda T, \Gamma, \bar{\rho})$. This Λ-tree with symmetry will be denoted $\Lambda\mathscr{T}$. Clearly $\mathscr{T} \mapsto \Lambda\mathscr{T}$ is a functor from the category of Λ-pretrees with symmetry to the category of Λ-trees with symmetry.

1.18. If $\Lambda_0 \subset \Lambda$ and if a group Γ acts without inversion on a Λ_0-tree T, then it follows from Corollary 6.16 and Lemma 6.15 of [1] that Γ acts without inversions on ΛT, and that the actions of Γ on T and ΛT define the same length function; in particular, the length function defined by the action of Γ on ΛT takes its values in Λ_0.

1.19. Suppose that a group Γ acts without inversions on a Λ-tree. It is clear that Γ acts freely on T if and only if the associated length function is nonzero on every nontrivial element of Γ. Hence it follows from 1.18 that if $\Lambda_0 \subset \Lambda$, and if a group Γ acts freely and without inversions on a Λ_0-tree T, then Γ also acts freely and without inversions on ΛT.

1.20. The remainder of this section is devoted to studying certain direct limits in the category of Λ-pretrees or of Λ-pretrees with symmetry.

Consider a direct system $(T_i; f_{ij})$ in the category of Λ-pretrees. (Such a system is understood to be indexed by some nonempty filtered ordered

set. Of course the f_{ij} are understood to be morphisms.) We shall write d_i to denote the distance function on T_i.

The direct system $(T_i; f_{ij})$ will be said to *converge strongly* if for every index i and for all x, $y \in T_i$, there exists some $k \geq i$ such that for every $j \geq k$, the map $f_{kj}|f_{ik}([x, y])$ is an isometric embedding of $f_{ik}([x, y])$ into T_j. Note that this implies in particular that for any x, $y \in T_i$ the sequence $(d_j(f_{ij}(x), f_{ij}(y)))_{j \geq i}$ is eventually constant, i.e., that there exists some $k \geq i$ such that for all $j \geq k$ we have $d_j(f_{ij}(x), f_{ij}(y)) = d_k(f_{ik}(x), f_{ik}(y))$.

1.21. Assuming that the direct system $(T_i; f_{ij})$ converges strongly, let T denote the direct limit of $(T_i; f_{ij})$ in the category of sets, and for each index i let $f_i: T_i \to T$ denote the canonical map. We define a function $d: T \times T \to \Lambda$ as follows. Given x, $y \in T$ we may write $x = f_i(x_i)$ and $y = f_i(y_i)$ for some positive integer i and some x_i, $y_i \in T_i$. Strong convergence implies that $(d_j(f_{ij}(x_i), f_{ij}(y_i)))_{j \geq i}$ is eventually constant. We define $d(x, y)$ to be the value of $d_j(f_{ij}(x_i), f_{ij}(y_i))$ for large j. This definition is at once seen to be independent of the choices of i, x_i and y_i.

Since the T_i are Λ-metric spaces, it is immediately clear that the set T becomes a Λ-metric space if we take d as the distance function. Furthermore, since the f_{ij} are distance-decreasing (1.7), the f_i are clearly also distance-decreasing.

1.22. Proposition. *The Λ-metric space T is a Λ-pretree, and the maps f_i are morphisms in the category of Λ-pretrees.*

Proof. We first show that T is isometric to a subspace of a Λ-tree. This will be proved using the result of Alperin and Bass discussed in 1.9. Thus we choose a base point $x_0 \in T$, and define $\wedge: T \times T \to \frac{1}{2}\Lambda$ by $x \wedge y = \frac{1}{2}(d(x, x_0) + d(y, x_0) - d(x, y))$. We must show that $\wedge$ satisfies conditions (a) and (b) of 1.9. To prove (b), consider three points x, y, $z \in T$. It follows from the definition of the metric on T that there exist an index i and points x_{0i}, x_i, y_i, $z_i \in T_i$ which are mapped by f_i to x_0, x, y and z respectively, and such that $f_i|\{x_{0i}, x_i, y_i, z_i\}$ is an isometric embedding. The base point $x_{i0} \in T_i$ defines a function $\wedge_i: T_i \times T_i \to \frac{1}{2}\Lambda$ satisfying (b). Hence we have

$$x \wedge z = x_i \wedge z_i \geq \min(x_i \wedge y_i, y_i \wedge z_i) = \min(x \wedge y, y \wedge z).$$

This shows that T satisfies (b). The same method establishes (a).

Thus we may identify T isometrically with a subspace of a Λ-tree $\overline{T}$. According to 1.4, in order to show that T is a Λ-pretree, we need only

prove that if x, y and z are points of T and if $[x, y]_{\overline{T}} \cap [x, z]_{\overline{T}} = [x, w]_{\overline{T}}$ for some $w \in \overline{T}$, then $w \in T$. To prove this, we first note that by the pretree axioms we have $d(x, w) = \frac{1}{2}(d(x, y) + d(x, z) - d(y, z))$. On the other hand, by the definition of the metric on T, there exist an index i and points x_i, y_i, $z_i \in T_i$ which are mapped by f_i to x, y and z respectively, and such that $f_i|\{x_i, y_i, z_i\}$ is an isometric embedding. We have $[x_i, y_i] \cap [x_i, z_i] = [x_i, w_i]$ for some $w_i \in T_i$. The pretree axioms give

$$d(x_i, w_i) = \tfrac{1}{2}(d(x_i, y_i) + d(x_i, z_i) - d(y_i, z_i)) = d(x, w).$$

Since the morphisms f_{ij} are necessarily distance-decreasing, the definition of the metric on T implies that the map f_i is distance-decreasing. Since $d(x, y) = d(x_i, y_i)$, it follows that f_i embeds $[x_i, y_i]$ isometrically in $[x, y]_T \subset [x, y]_{\overline{T}}$. But we have $w \in [x_i, y_i]$, $w \in [x, y]$ and $d(x_i, w_i) = d(x, w)$. Therefore $f_i(w_i) = w$, so that in particular $w \in T$. This completes the proof that T is a Λ-pretree.

To prove that f_i is a morphism for each index i, we consider an arbitrary presegment $\sigma \subset T_i$. By the definition of strong convergence there is an index $k \geq i$ such that for every $j \geq k$, the map $f_{kj}|f_{ik}(\sigma)$ is an isometric embedding. By the definition of the metric on T, the map $f_k|f_{ik}(\sigma)$ is an isometric embedding. Since f_{ik} is a morphism, we may write σ as a finite union of presegments each of which is isometrically embedded in T_k by f_{ik}. Hence f_i embeds each of these presegments isometrically in T. This shows that f_i is a morphism.   q.e.d.

The Λ-pretree T given by 1.22 will be called the *limit* of the strongly convergent system $(T_i; f_{ij})$.

1.23.   Remark.   It is easy to show that T is the direct limit of $(T_i; f_{ij})$ in the category of Λ-pretrees. However, a direct system of Λ-pretrees may well have a direct limit without being strongly convergent. For example, for each positive real number r with $0 < r < 1$, let T_r denote the $\mathbb{R}$-tree

with $d(x_r, z_r) = d(y_r, z_r) = r$ and $d(z_r, w_r) = 1 - r$. Whenever $0 < s < r < 1$, there is a unique morphism $f_{rs}: T_r \to T_s$ mapping x_r, y_r and z_r to x_s, y_s and z_s respectively. Hence $(T_r; f_{rs})$ is a direct system of $\mathbb{R}$-trees; it is easily seen that the unit interval in $\mathbb{R}$, regarded as an $\mathbb{R}$-tree, is a direct limit of this system. However, the system clearly does not converge strongly.

1.24.   The following characterization of a strongly convergent system of Λ-pretrees and of its limit will be useful.

Proposition.   *Let (T_i, f_{ij}) be a direct system in the category of Λ-pretrees. Let T be a Λ-pretree, and for each index i let $f_i: T_i \to T$ be a morphism of Λ-pretrees. Suppose that the following conditions hold:*

(i) $f_i = f_j \circ f_{ij}$ *whenever* $i \le j$;
(ii) $T = \bigcup_i f_i(T_i)$;
(iii)   *given any index i and any two points x, $y \in T_i$, there exists some $j \ge i$ such that $d(f_{ij}(y)) = d(f_i(x), f_i(y))$.*

Then the direct system (T_i, f_{ij}) converges strongly, and its limit is isometric to T. Furthermore, the isometric identification of T with the limit of the T_i may be made in such a way that the f_i are the canonical morphisms of the T_i to their limit.

Proof.   Observe that since a morphism of Λ-pretrees is distance-decreasing, hypothesis (iii) implies:

(iii′)   *given any index i and any two points x, $y \in T_i$, we have*

$$d(f_{ij}(x), f_{ij}(y)) = d(f_i(x), f_i(y))$$

for all sufficiently large $\ge i$

To prove that (T_i, f_{ij}) converges strongly, we consider an arbitrary index i and an arbitrary presegment $[x, y] \subset T_i$. Since f_i is a morphism we may write $[x, y] = [x, y]_{T_i}$ in the form $[x_0, x_1] \cup \cdots \cup [x_{n-1}, x_n]$, where $x_0 = x$, $x_n = y$, $[x_{r-1}, x_r] \cap [x_r, x_{r+1}] = \{x_r\}$ for $0 < r < n$, and $f_i|[x_{r-1}, x_r]$ is an isometric embedding for $1 \le r \le n$. By (iii)′ there is an index $k \ge i$ such that for all nonnegative integers r, $s \le n$ we have $d(f_{ik}(x_r), f_{ik}(x_s)) = d(f_i(x_r), f_i(x_s))$. We claim that for every $j \ge k$, the map $f_{kj}|f_{ik}([x, y])$ is an isometric embedding; this will prove strong convergence.

Set $S = \{f_{ik}(x_1), \cdots, f_{ik}(x_n)\}$. Then according to our choice of k, the map $f_k|S$ is an isometric embedding of S in T. On the other hand, note that since $f_i|[x_{r-1}, x_r]$ is an isometric embedding for $1 \le r \le n$, and since f_k is distance-decreasing, $f_{ik}|[x_{r-1}, x_r]$ is also an isometric embedding for $1 \le r \le n$.

Let u_1 and u_2 be any two points of $f_{ik}([x, y])$. For $m = 1, 2$, set $u_m = f_{ik}(w_m)$, where $w_m \in [x, y]$. Then each w_m lies in one of the presegments $[x_{r-1}, x_r]$, $1 \le r \le n$. Since $f_{ik}|[x_{r-1}, x_r]$ is an isometric embedding it follows that each u_i lies in one of the presegments $[f_{ik}(x_{r-1}), f_{ik}(x_r)]$. Thus each u_i lies in a presegment with endpoints

in S. Hence by 1.5, it follows that $[u_1, u_2]$ is contained in a presegment σ with endpoints in S. Since $f_k|S$ is an isometric embedding, and since f_k is distance-decreasing, it follows that $f_k|\sigma$ is an isometric embedding. In particular $d(f_k(u_1), f_k(u_2)) = d(u_1, u_2)$. Finally, for any $j \geq k$, since f_{kj} and f_j are distance-decreasing, we must have $d(f_{kj}(u_1), f_{kj}(u_2) = d(u_1, u_2)$. This shows that $f_{kj}|f_{ik}([x, y])$ is an isometric embedding, and completes the proof of strong convergence.

Next note that for any given index i, it follows from hypothesis (iii) that two points of T_i have the same image under f_i if and only if they have the same image under f_{ij} for some $j \geq i$. This, together with hypothesis (i) and (ii), allows us to identify T with the direct limit of the T_i in the category of sets, in such a way that the f_i are the canonical maps of the T_i to their direct limit. Hypothesis (iii) now shows that the given metric on T_i coincides with the metric on the set-theoretic direct limit constructed in 1.21. Thus T is the limit of the strongly convergent system (T_i, f_{ij}). q.e.d.

1.25. Now consider a direct system $(\mathscr{T}_i; \mathscr{f}_{ij})$, in the category of Λ-pretrees with symmetry, where $\mathscr{T}_i = (T_i, \Gamma_i, \rho_i)$ and $\mathscr{f}_{ij} = (f_{ij}, \omega_{ij})$. We shall say that such a sequence *converges strongly* if the underlying direct system (T_i, f_{ij}) of Λ-pretrees converges strongly. Let T denote the limit of this strongly convergent system, and let Γ denote the direct limit of the direct system (Γ_i, ω_{ij}) of groups. For each index i, let $f_i: T_i \rightarrow T$ denote the canonical morphism, and let $\omega_i: \Gamma_i \rightarrow \Gamma$ denote the natural homomorphism. Then there is a unique action ρ of Γ on T such that (f_i, ω_i) is a morphism from $\mathscr{T}_i$ to $\mathscr{T} = (T, \Gamma, \rho)$ for each i. We call $\mathscr{T}$ the *limit* of the given strongly convergent system $(\mathscr{T}_i; \mathscr{f}_{ij})$.

1.26. The following characterization of a strongly convergent system of Λ-pretrees with symmetry, and of its limit, is an immediate consequence of Proposition 1.24.

Proposition. *Let $(\mathscr{T}_i; \mathscr{f}_{ij})$ be a direct system in the category of Λ-pretrees with symmetry, where $\mathscr{T}_i = (T_i, \Gamma_i, \rho_i)$ and $\mathscr{f}_{ij} = (f_{ij}, \omega_{ij})$. Let $\mathscr{T} = (T, \Gamma, \rho)$ be a Λ-pretree with symmetry, and for each index i let $f_i = (f_i, \omega_i): \mathscr{T}_i \rightarrow \mathscr{T}$ be a morphism of Λ-pretrees with symmetry. Suppose that the following conditions hold.*

(i) *$\mathscr{f}_i = \mathscr{f}_j \circ \mathscr{f}_{ij}$ whenever $i \leq j$.*

(ii) *$T = \bigcup_i f_i(T_i)$.*

(iii) *Given any $i \geq 1$ and any two points x, $y \in T_i$, there exists some $j \geq i$ such that $d(f_{ij}(x), f_{ij}(y)) = d(f_i(x), f_i(y))$.*

(iv) *The ω_i induce an isomorphism of $\varinjlim \Gamma_i$ onto Γ.*

Then the system $(\mathscr{T}_i; f_{ij})$ *converges strongly, and its limit is isomorphic to* $\mathscr{T}$ *. Furthermore, the isomorphic identification of* $\mathscr{T}$ *with the limit of the* $\mathscr{T}_i$ *may be made in such a way that the* f_i *are the canonical morphisms of the* $\mathscr{T}_i$ *to their limit.*

1.27. If $(\mathscr{T}_i; f_{ij})$ is a direct system in the category of Λ-pretrees with symmetry, then by 1.17 we have system $(\Lambda\mathscr{T}_i; \Lambda f_{ij})$ in the category of Λ-trees with symmetry.

Proposition. *If* $(\mathscr{T}_i; f_{ij})$ *converges strongly, then* $(\Lambda\mathscr{T}_i; \Lambda f_{ij})$ *also converges strongly, and its limit is isomorphic to the* Λ-*completion of the limit of* $(\mathscr{T}_i; f_{ij})$.

Proof. Let us write $\mathscr{T}_i = (T_i, \Gamma_i, \rho_i)$ and $f_{ij} = (f_{ij}, \omega_{ij})$.

Let $\mathscr{T} = (T, \Gamma, \rho)$ denote the limit of the system $(\mathscr{T}_i; f_{ij})$, and for each i let $f_i = (f_i, \omega_i): \mathscr{T}_i \to \mathscr{T}$ denote the natural morphism. We shall prove the proposition by showing that conditions (i)–(iv) of 1.26 are satisfied when $(\mathscr{T}_i; f_{ij})$, $\mathscr{T}$ and f_i are replaced by $(\Lambda\mathscr{T}_i; \Lambda f_{ij})$, $\Lambda\mathscr{T}$ and Λf_i. Conditions (i) and (iv) are clear. To check (ii) we consider a point $x \in \Lambda T$. Let $[y, z]_{\Lambda T}$ be a segment having endpoints in T and containing x. By the definition of the metric on T there exist an index i and points y_i, $z_i \in T_i$ such that $f(y_i) = y$, $f(z_i) = z$ and $d(y, z) = d(y_i, z_i)$. Hence the morphism Λf_i restricts to an isometry of $[y_i, z_i]_{\Lambda T_i}$ onto $[y, z]_{\Lambda T}$. In particular, $x \in f_i(\Lambda T_i)$. This proves (ii).

To prove (iii) we consider an index i and two points x, $y \in \Lambda T_i$. By 1.12 there is a segment σ in ΛT_i which contains $[x, y]$ and has its endpoints in T_i. By the definition of strong convergence and the definition of the metric on T, there is an index $j \geq i$ such that $f_j | f_{ij}(\sigma)$ is an isometric embedding. In particular we have $d(f_{ij}(x), f_{ij}(y)) = d(f_i(x), f_i(y))$. q.e.d.

1.28. We conclude this section with a simple result on length functions. Let X be a set and let $l_1, l_2, \ldots$ be a sequence of real-valued functions on X. We shall say that $l_1, l_2, \ldots$ *converges strongly* if for each $x \in X$ the sequence $l_1(x), l_2(x), \ldots$ is eventually constant. A strongly convergent sequence (l_i) has a limit function $l: X \to \mathbb{R}$, which can be defined for each $x \in X$ by $l(x) = l_i(x)$ for sufficiently large i.

Proposition. *Let* Γ *be a group, and let* $(\mathscr{T}_i; f_{ij})$ *be a strongly convergent direct system of* Λ-*trees with symmetry, where* $\mathscr{T}_i = (T_i, \Gamma, \rho_i)$ *and* $f_{ij} = (f_{ij}, 1)$. *Let* $\mathscr{T} = (T, \Gamma, \rho)$ *denote the limit of the system. Suppose that* Γ *acts without inversions on each* T_i. *Then* Γ *also acts without inversions on* T, *and the length functions defined by the* $\mathscr{T}_i$ *converge strongly to the length function defined by* $\mathscr{T}$.

Proof. Let $\mathcal{f}_i = (f_i, 1): T_i \to T$ denote the canonical morphism.

To show that Γ acts without inversions on T, we must show that if x and y are points of T such that $\gamma \cdot x = y$ and $\gamma \cdot y = x$ for some $\gamma \in \Gamma$, then $d(x, y)$ is divisible by 2 in Λ. For some index i there is a point $x_i \in T_i$ with $f_i(x_i) = x$. Set $x_i' = \gamma^2 \cdot x_i$. Then $f_i(x_i') = x$. Hence for large enough $j \geq i$ we have $f_{ij}(x_i) = f_{ij}(x_i') = x_j$, say. Set $y_j = \gamma \cdot x_j$; then $\gamma \cdot y_j = x_j$. Since Γ acts without inversions on T_j, it follows that $d(x_j, y_j)$ is divisible by 2 in Λ. But for large enough $j \geq i$ we have $d(x_j, y_j) = d(x, y)$. This proves that Γ acts without inversions on T.

Now let d_i and l_i (resp. d and l) denote the distance function and length function for $\mathcal{T}_i$ (resp. $\mathcal{T}$). Let $g \in \Gamma$ be given. Since the f_i and f_{ij} are distance-decreasing it is clear from the definition of the length functions that $l_i(g) \geq l_j(g) \geq l(g)$ whenever $i \leq j$. Now let x be a point of T such that $d(x, \rho(g)x) = l(g)$. For some index m there is a point $y \in T_m$ such that $f_m(y) = x$. By strong convergence, for some $n \geq m$ we have $d(f_{mn}(y), f_{mn}(\rho_m(g)y)) = d(x, \rho(g)x)$. Thus $l(g) \geq l_n(g)$. Hence $l_i(g) = l(g)$ for all $i \geq n$.   q.e.d.

1.29.   *For the rest of this paper, all direct systems will be understood to be indexed by the natural numbers.*

2. The rank-1 case: statement of results

In this section we state the structure theorem (Theorem 2.1) for group actions on Λ-trees, where Λ is a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. We also give more formal statements of the rank-1 versions of Theorems A, B, C and D of the introduction. We prove in this section that the latter results follow from the structure theorem. Along the way we formulate a result, Proposition 2.6, which, in the case of a finitely presented group, gives a partial reinterpretation of the structure theorem in terms of the Bass-Serre theory of graphs of groups. The proof of the structure theorem itself will be given in §4.

2.1.   Theorem.   *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1, and let $\mathcal{T} = (T, \Gamma, \rho)$ be a Λ-tree with symmetry. Then $\mathcal{T}$ is the limit of a strongly convergent direct system $(\mathcal{T}_i; \mathcal{f}_{ij})$ in the category of Λ-pretrees with symmetry, where each $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$ is a Λ_i-tree with symmetry for some cyclic subgroup Λ_i of Λ, and $\Lambda_i \subset \Lambda_j$ whenever $i \leq j$. Furthermore, if $\mathcal{f}_i = (f_i, \omega_i)$ denotes the canonical morphism from $\mathcal{T}_i$ to $\mathcal{T}$, the kernel of ω_i acts freely and without inversions on T_i for each i.*

Theorem 2.1 will be proved in §4. For the rest of this section it will be assumed.

2.2.   The following result is the rank-1 case of Theorem A of the introduction.

Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. Let Γ be a finitely generated group which acts freely, without inversions, on some Λ-tree. Then Γ is a free group.*

Furthermore, T is the limit of a strongly convergent system (T_i, f_{ij}) in the category of Λ-pretrees, where each T_i is a Λ_i-tree for some infinite cyclic subgroup Λ_i of Λ; and there exist free actions without inversions of Γ on the T_i such that the f_{ij} and the canonical morphisms $f_i: T_i \to T$ are Γ-equivariant.

2.3.   The following group-theoretical lemma is needed to deduce Proposition 2.2 from Theorem 2.1. It is stated in a general enough form to be used for other applications later in this section and in §6. We are grateful to John Wood for pointing out a proof of the lemma in the special case where Γ is finitely generated and all the Γ_i are free, which is the case needed for the proof of 2.2. The proof given below in case (ii) was suggested by Wood's argument.

Lemma. *Let (Γ_i, ω_{ij}) be a direct system in the category of groups. Set $\Gamma = \varinjlim \Gamma_i$. Suppose that either*

(i)   *Γ is finitely presented, or*
(ii)   *Γ is finitely generated, and there exist an integer $n \geq 1$ and a field K such that for each i, every finitely generated subgroup of Γ_i is isomorphic to a subgroup of $\mathrm{GL}_n(K)$.*

Then for every sufficiently large index i, there is a subgroup Γ_i^* of Γ_i such that the canonical homomorphism $\omega_i: \Gamma_i \to \Gamma$ maps Γ_i^* isomorphically onto Γ.

Proof. Since Γ is finitely generated, there is some index l such that ω_l is surjective. We may assume that $l = 1$. There is a finitely generated subgroup Γ_l^* of Γ_l such that $\omega_l(\Gamma_l^*) = \Gamma$. After replacing each Γ_i by $\omega_{li}(\Gamma_l^*)$, we may assume that the Γ_i are finitely generated and that the ω_{ij} are all surjective. Under this assumption we shall prove the lemma by showing that ω_i is an isomorphism for all sufficiently large i.

We fix a homomorphism h of a finitely generated free group F onto Γ_l. For each i, let R_i denote the kernel of $\omega_{li} \circ h: F \to \Gamma_i$. Likewise, let R denote the kernel of $\omega_l \circ h: F \to \Gamma$. Then $R_i \subset R_j$ whenever $i \leq j$, and $R = \bigcup_i R_i$.

Suppose that (i) holds. Since $\Gamma \approx F/R$ is finitely presented, the normal subgroup R of F is the normal closure of a finite set of elements of F.

Hence for sufficiently large i we must have $R_i = R$. Then ω_i is an isomorphism.

Now suppose that (ii) holds. Since we are now in the case where the Γ_i are finitely generated, this means that each Γ_i is isomorphic to a subgroup of $GL_n(K)$. Let us write r for the rank of the free group F. Fixing a basis for F establishes a bijective correspondence between the set of all representations of F in $GL_n(K)$ and the Cartesian power $GL_n(K)^r$. We regard the latter set as a subset of the affine space $M_n(K)^r$ and endow it with the Zariski topology. For each index i, let V_i denote the set of all points of $GL_n(K)^r$ that correspond to representations whose kernels contain R_i. Then V_i is a closed subset of $GL_n(K)$, and $V_i \supset V_j$ for $i \leq j$.

By the Hilbert Basis Theorem, there is an index m such that $V_i = V_m$ for all $i \geq m$. This means that whenever the kernel of a representation of F contains R_m it must actually contain R_i for all $i \geq m$. But according to condition (ii), for each i there is a representation ρ_i of F whose kernel is precisely R_i. Hence we must have $R_i = R_m$ for all $i \geq m$. This means that for $i \geq m$ we have $R = R_i$, so that ω_i is an isomorphism. q.e.d.

2.4. *Proof of Proposition* 2.2. Let $(\mathcal{T}_i; f_{ij})$ be the direct system, and $\Lambda_i \subset \Lambda$ the subgroups, given by Theorem 2.1. We write $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$ and $f_{ij} = (f_{ij}, \omega_{ij})$. For each index i, let $f_i = (f_i, \omega_i): \mathcal{T}_i \to \mathcal{T}$ be the canonical morphism. Since Λ_i is a cyclic subgroups of Λ, we can identify T_i with the 0-skeleton of a simplicial tree K_i, and the action of Γ_i on T_i is the restriction of a simplicial action.

The morphism f_i is in particular a morphism of sets with symmetry. Since Γ acts freely on T, it follows that any element of Γ_i that fixes a point of T_i lies in the kernel of ω_i. But by 2.1, the kernel of ω_i acts freely on T_i. Therefore the action of Γ_i on T_i is free.

If τ is any 1-simplex of K_i, the endpoints of τ constitute a segment in T_i whose length is the positive generator of Λ_i. It follows from the definition (1.7) of a morphism of Λ-pretrees that f_i embeds such a segment isometrically in T. Hence if an element g of Γ_i leaves τ invariant then $\omega_i(g)$ must leave a segment in T invariant. Since Γ acts freely and without inversions on T this implies that $g \in \ker \omega_i$. Since $\ker \omega_i$ acts freely and without inversions on T_i, we have $g = 1$. This shows that Γ_i acts without inversions on K_i in the sense of [15].

Since Γ_i acts freely and without inversions on a simplicial tree, it is a free group (see [15, Chapter I, Theorem 4]). On the other hand, by the definition of the limit of a strongly convergent system of trees with symmetry, Γ is the direct limit of the directed system of groups (Γ_i, ω_{ij}).

It follows from [15, Chapter I, Example 1.5.3] that a finitely generated free group is isomorphic to a subgroup of $GL_2(\mathbb{C})$. Hence we can apply 2.3 to conclude that Γ is isomorphic to a subgroup of one of the Γ_i and is therefore itself a free group.

To prove the last assertion of the proposition, note that by 2.3 there exists, for every sufficiently large i, say for $i \geq m$, a subgroup Γ_i^* of Γ_i such that ω_i maps Γ_i^* isomorphically onto Γ. Restricting the action of Γ_i gives an action of Γ_i^* on T_i, which we can then push forward via the isomorphism $\omega_i|\Gamma_i^*$ to define an action of Γ on T_i. Since the f_i and f_{ij} are morphisms of sets with symmetry, f_i and f_{ij} are Γ-equivariant, in terms of the pushed-forward actions, for $m \leq i \leq j$.

2.5. The next result, Proposition 2.6, will be used in deducing the rank-1 case of Theorems B, C, and D of the introduction from Theorem 2.1. The statement of the result involves the notion of a graph of groups. Recall from [15] that a *graph of groups* $\mathscr{G}$ is a connected 1-dimensional CW-complex G in which each cell c is labelled with a group $\mathscr{G}(c)$, and for each oriented edge e with terminal vertex v, a homomorphism is specified from $\mathscr{G}(e)$ to $\mathscr{G}(v)$. In [15] a group $\pi_1(\mathscr{G})$ is associated with $\mathscr{G}$, the *fundamental group* of the graph of groups $\mathscr{G}$. For each cell c of G, the group $\mathscr{G}(c)$ is identified isomorphically with a subgroup of $\pi_1(\mathscr{G})$; this identification is canonical modulo inner automorphisms of $\pi_1(\mathscr{G})$. Thus a cell of G determines a conjugacy class of subgroups of $\pi_1(\mathscr{G})$. We shall call a subgroup of $\pi_1(\mathscr{G})$ a *vertex group* (or *edge group*) if it is in the conjugacy class of subgroups determined by some vertex (or edge) of G.

2.6. Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. Let Γ be a finitely presented group which acts without inversions on a Λ-tree T. Then Γ can be identified isomorphically with the fundamental group of a graph of groups in such a way that*

(i) *each vertex group in Γ is contained in the stabilizer of some point in T, and*
(ii) *each edge group in Γ is contained in the stabilizer of a nondegenerate arc in T.*

Furthermore, T is the limit of a strongly convergent system (T_i, f_{Ij}) in the category of Λ-pretrees, where each T_i is a Λ_i-tree for some infinite cyclic subgroup Λ_i of Λ; and there exist actions without inversions of Γ on the T_i such that the f_{ij} and the canonical morphisms $f_i: T_i \to T$ are Γ-equivariant.

Proof. Let $(\mathcal{T}_i; \mathcal{f}_{ij})$ be the direct system, and $\Lambda_i \subset \Lambda$ the subgroups, given by Theorem 2.1. We write $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$ and $\mathcal{f}_{ij} = (f_{ij}, \omega_{ij})$. For each index i, let $\mathcal{f}_i = (f_i, \omega_i): \mathcal{T}_i \to \mathcal{T}$ be the canonical morphism. By the definition of the limit of a strongly convergent system of Λ-pretrees with symmetry, Γ is the direct limit of the system of groups (Γ_i, ω_{ij}). Since Γ is finitely presented, we can apply 2.3 to obtain a subgroup Γ_i^* of Γ_i, for all sufficiently large i, say $i \geq m$, such that ω_i maps Γ_i^* isomorphically onto Γ. As in the last step of the proof of 2.2 we can then conclude that there are actions of Γ on the T_i for $i \geq m$ such that f_i and f_{ij} are Γ-equivariant for $m \leq i \leq j$. This proves the last assertion of the proposition.

To prove the other assertions we use the Bass-Serre theory [15]. As in the proof of 2.2, we regard T_m as the 0-skeleton of a simplicial tree K_m. Then Γ acts without inversions on the barycentric subdivision K_m^* of K_m. According to [15, Chapter I, Theorem 13] we can identify Γ isomorphically with the fundamental group of a graph of groups $\mathcal{G}$ in such a way that each vertex group in Γ is the stabilizer of some vertex of K_m^*, and each edge group in Γ is the stabilizer of an edge in K_m^*. We claim that conditions (i) and (ii) of the statement of Proposition 2.6 then hold.

Consider any vertex group $\mathcal{G}(v) \subset \Gamma$. Since $\mathcal{G}(v)$ fixes a vertex of K_m^*, it either fixes a vertex of K_m or leaves an edge of K_m invariant. If $\mathcal{G}(v)$ fixes a vertex of K_m, i.e., a point $x \in T_m$, then since f_i is Γ-equivariant, $\mathcal{G}(v)$ is contained in the stabilizer of $f_m(x) \in T$. Now suppose that $\mathcal{G}(v)$ leaves invariant an edge of K_m with endpoints x, $y \in T_m$. Then the set $\{f_m(x), f_m(y)\}$, and hence the segment $[f_m(x), f_m(y)]$, are invariant under $\mathcal{G}(v)$. Since Γ acts without inversions on T, some point of $[f_m(x), f_m(y)]$ must be fixed by $\mathcal{G}(v)$. This establishes condition (i).

Now consider an edge group $\mathcal{G}(e)$. Since $\mathcal{G}(e)$ fixes an edge of K_m^*, it must fix (pointwise) some edge of K_m with endpoints x, $y \in T_m$. The distance in T_m between x and y is the positive generator of the infinite cyclic group Λ_m. Hence by the definition of a morphism of trees, f_m maps $[x, y] \subset T_m$ isometrically onto $[f_m(x), f_m(y)] \subset T$. It follows that the latter segment is nondegenerate and, by the Γ-equivariance of f_m, is fixed by $\mathcal{G}(e)$. This establishes condition (ii).   q.e.d.

2.7.   The next result combines the rank-1 cases of Theorems C and D of the introduction. Recall from the introduction that a group Γ is said to *split over* a subgroup Δ if Γ is either a free product of two proper

subgroups with amalgamated subgroup Δ, or an HNN extension with associated subgroup Δ.

Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. Suppose that a finitely presented group Γ admits a nontrivial action, without inversions, on a Λ-tree T. Then Γ splits over some subgroup. If in addition the stabilizer of every nondegenerate segment of T is a small subgroup of Γ, then Γ splits over a small subgroup.*

2.8. The proof of 2.7 is based on the following lemma.

Lemma. *Let $\mathscr{G}$ be a graph of groups. Suppose that $\pi_1(\mathscr{G})$ is finitely generated and that every vertex group in $\pi_1(\mathscr{G})$ is a proper subgroup. Then there is some edge e of $\mathscr{G}$ such that $\pi_1(\mathscr{G})$ splits over $\mathscr{G}(e)$.*

Proof. Since $\Gamma = \pi_1(\mathscr{G})$ is finitely generated, there is a finite subgraph $\mathscr{G}_0$ of the graph of groups $\mathscr{G}$ such that $\pi_1(\mathscr{G}_0)$ maps isomorphically onto $\pi_1(\mathscr{G})$. Among all finite subgraphs with this property, we may suppose $\mathscr{G}_0$ to have been chosen to have the smallest possible number of vertices. If the underlying graph G_0 of $\mathscr{G}_0$ is nonsimply connected then it has a nonseparating edge e. It follows from the definition of $\pi_1(\mathscr{G}_0)$ that Γ is a HNN group with $\mathscr{G}(e)$ as amalgamated subgroup. If G_0 is simply connected it has an endpoint v. Since the vertex groups in $\pi_1(\mathscr{G})$ are proper subgroups we have $\mathscr{G}(v) \neq \Gamma$ and hence $\{v\} \neq G_0$. Thus v is incident to a unique edge e of G_0, and the graph G_1 obtained by removing v and e from G_0 is nonempty and connected. If $\mathscr{G}_1$ is the graph of groups corresponding to G_1, the definition of $\pi_1(\mathscr{G}_0)$ implies that $\Gamma = \pi_1(\mathscr{G}_1) *_{\mathscr{G}(e)} \mathscr{G}(v)$. But $\pi_1(\mathscr{G}_1) \neq \Gamma$ by the minimality of $\mathscr{G}_0$. Thus Γ is a nontrivial free product with amalgamation, the amalgamated subgroup being $\mathscr{G}(e)$. q.e.d.

2.9. *Proof of* 2.7. If Γ admits a nontrivial action on Λ-tree T, where the $\mathbb{Q}$-rank of Λ is 1, then according to 2.6, Γ may be identified with $\pi_1(\mathscr{G})$, where $\mathscr{G}$ is a graph of groups, in such a way that (i) every vertex group is a proper subgroup of Γ and (ii) every edge group is contained in the stabilizer of a nondegenerate segment in T. The conclusion now follows from Lemma 2.8. q.e.d.

2.10. The final result of this section, Proposition 2.11, includes the rank-1 case of Theorem B of the introduction. However, it will be stated in elementary form, without reference to the theory of measured foliations.

Let F be a closed (piecewise-linear or smooth) surface. If $C \subset F$ is a closed 1-manifold whose components are two-sided, homotopically nontrivial simple closed curves, then C determines a function $l_C \colon \pi_1(F) \to \mathbb{N}$, which assigns to any element γ of $\pi_1(F)$ the geometric intersection number with C of the homotopy class of closed curves corresponding to

the conjugacy class of γ. According to [12, Proposition III.2.3] l_C is the length function for a nontrivial action of $\pi_1(F)$ on a $\mathbb{Z}$-tree.

2.11. Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. Let F be a closed surface and set $\Gamma = \pi_1(F)$. Suppose that Γ acts on a Λ-tree T in such a way that the stabilizer of every nondegenerate segment in T is a cyclic subgroup of Γ. Then the length function defined by the action has the form $l = \lambda \cdot l_C$, where λ is a positive element of Λ, $C \subset F$ is some compact two-sided 1-manifold whose components are homotopically nontrivial, " $\cdot$ " denotes multiplication of real numbers.*

2.12. The proof of 2.11 depends on the following lemma. Let us define a *geometric length function* on $\pi_1(F)$ to be a function of the form l_C, where $C \subset F$ is some compact two-sided 1-manifold whose components are homotopically nontrivial.

Lemma. *Let Γ be the fundamental group of a closed surface. Then there exists a finite set Φ of a conjugacy classes in Γ such that any two geometric length functions which agree on Φ are equal.*

Proof. The lemma asserts that there exist closed curves $\sigma_1, \cdots, \sigma_N$ in F such that if C and C' are any two compact two-sided 1-manifolds in F without homotopically trivial components, and the geometric intersection numbers $\mathrm{I}(C, \sigma_i)$ and $\mathrm{I}(C', \sigma_i)$ are equal for all i, then $\mathrm{I}(C, \sigma) = \mathrm{I}(C', \sigma)$ for every closed curve σ in F. This follows from results proved in [5]. (In [5, Exposé 6, §IV.5 and Appendix], it is shown that there are simple closed curves $\sigma_1, \cdots, \sigma_N$ such that any two measured foliations μ and μ' which satisfy $\mathrm{I}(\mu, \sigma_i) = \mathrm{I}(\mu', \sigma_i)$ for and i are equivalent via Whitehead moves; this implies that they satisfy $\mathrm{I}(\mu, \sigma) = \mathrm{I}(\mu', \sigma)$ for every σ. To deduce our assertion we need only apply the latter result to the measured foliations associated to C and C' by the construction of [5, Exposé 5, §III.1]. These arguments can easily be translated into an elementary proof of the lemma not referring to measured foliations.)   q.e.d.

2.13. *Proof of 2.11.* Let $(T_i; f_{ij})$ be the direct system, and $\Lambda_i \subset \Lambda$ the subgroups, given by Theorem 2.1. We may write $\Lambda_i = \lambda_i \mathbb{Z}$ for some $\lambda_i > 0$. We can identify T_i with the 0-skeleton of a simplicial tree K_i so that the action of Γ on T_i is the restriction of a simplicial action. Since the natural morphisms $f_i : T_i \to T$ are Γ-equivariant, the stabilizer in Γ of each edge in K_i is cyclic. It therefore follows from [12, Theorem III.2.6] that the length function l_i defined by the action of Γ on T_i has the form $\lambda_i \cdot l_i'$, where l_i' is a geometric length function.

Let Φ be the finite set given by 2.12. According to 1.28, the l_i converge strongly to the length function l defined by the action of Γ on T. In particular there is a natural number m such that $l_m|\Phi = l_i|\Phi = l|\Phi$ for

all $i \geq m$. If we write $\lambda_i = a_i/b_i$ with a_i, $b_i \in \mathbb{N}$, then the geometric length functions $b_i b_m l_i = a_i b_m l_i'$ and $b_i b_m l_m = a_m b_i l_m'$ agree on Φ and are therefore equal. Hence $l_i = l_m$ for all $i \geq m$. By strong convergence it follows that $l = l_m = \lambda_m \cdot l_m'$.

3. Complexes associated to trees

The proof of Theorem 2.1, which all be given in §4, is an illustration of the principal method of this paper: one studies group actions on a Λ-tree T, where Λ is a subgroup of $\mathbb{R}$, by associating a simplicial complex to T in a natural way. Groups that act on T then act in a natural way on the associated complex. In this section we shall give a general discussion of complexes associated with trees.

3.1. We begin by fixing some conventions involving topological spaces, simplicial complexes, paths and homotopies that will be used for the rest of the paper. Recall that a *path* γ in a topological space X is a map of the unit interval into X. The image of the map will be called the *support* of γ and will be denoted $|\gamma|$. *Homotopy* of paths will always mean fixed-endpoint homotopy; in particular, the assertion that two paths are homotopic includes the assertion that they have the same endpoints. A path which is a 1-1 map is called a *parametrized arc*. The *inverse* $\bar{\gamma}$ of a path γ is defined by $\bar{\gamma}(t) = \gamma(1 - t)$. If $\gamma_1, \cdots, \gamma_n$ are paths with $\gamma_1(1) = \gamma_{i+1}(0)$ for $1 \leq i < n$, their *composition*, denoted $\gamma_1 * \cdots * \gamma_n$, is the path γ defined by $\gamma|[(i - 1)/n, i/n] = \gamma_i \circ h_i$, where h_i is the unique orientation-preserving affine map of $[(i - 1)/n, i/n]$ onto $[0, 1]$.

If T is an $\mathbb{R}$-tree, it follows from Proposition II.1.13 of [12] that the image of a parametrized arc $\gamma \colon [0, 1] \to T$ is a segment $[x, y] \subset T$. For this reason a parametrized arc in an $\mathbb{R}$-tree will often be called a *parametrized segment*, or—when the support σ is named explicitly—a *parametrization* of σ.

It will be necessary to distinguish between *abstract* and *geometric* simplicial complexes. The main reason for being careful about the distinction is that we will often be considering nonsimplicial maps between geometric simplicial complexes.

Recall that a geometric simplicial complex is a cell complex in which each (closed) cell has the structure of an affine simplex in a Euclidean space, and the inclusion map of a face into a cell is always an affine map. A geometric simplicial complex will be regarded as a topological space, with the usual weak topology.

On the other hand, an abstract simplicial complex is defined by a set, the set of vertices, having a distinguished class of finite subsets called simplices, such that every subset of a simplex is a simplex. The geometric realization of an abstract simplicial complex is a geometric simplicial complex in the above sense, and conversely.

In an abstract simplicial complex, an n-simplex whose vertices are $x_1, \cdots, x_{n+1}$ is by definition $\{x_1, \cdots, x_{n+1}\}$. In a geometric simplicial complex, an n-simplex whose vertices are $x_1, \cdots, x_{n+1}$ will be denoted $\Delta(x_1, \cdots, x_{n+1})$.

If X is an abstract or geometric simplicial complex and i is a nonnegative integer, we shall write $sk_i(X)$ for the i-skeleton of X.

A simplicial complex—whether abstract or geometric—is *not* assumed to be locally finite.

The term "complex", when unmodified, will mean "geometric simplicial complex".

3.2.   There is a familiar category in which the objects are complexes and the morphisms are simplicial maps. An object with symmetry in this category will be called a *complex with simplicial symmetry*.

Let (X, Γ, ρ) be a complex with simplicial symmetry. We shall say that X is *finite mod* Γ if there are only finitely many orbits for the action of Γ on the set of simplices of Y.

For our purposes it will often be necessary to consider categories in which more general maps appear as morphisms.

Let X and X' be geometric simplicial complexes. Let k be an integer ≥ 0. By a k-*tame map* from X to X' we shall mean a continuous map $f: X \to X'$ such that for every closed simplex Δ and X whose dimension is at most k, $f(\Delta)$ is a subcomplex of X with $\dim f(\Delta) \leq \dim \Delta$. In particular a k-tame map carries $sk_i(X)$ into $sk_i(X')$ for every $i \leq k$.

A k-tame map is k'-tame for any $k \leq k$. A map is 0-tame if and only if it is continuous and takes vertices to vertices. A simplicial map is in particular k-tame for every k.

Let X be a geometric simplicial complex. For $k = 0$ or 1, we shall say that a path $\gamma: [0, 1] \to X$ is k-*tame* if it is a k-tame map of simplicial complexes, where $[0, 1]$ is triangulated with two 0-simplices and a single 1-simplex.

A path is 0-tame if and only if its endpoints are vertices of X. A parametrized arc γ is 1-tame if and only if it is 0-tame and has support contained in $sk_1(X)$.

For every k there is a k-*tame category of* (*geometric*) *simplicial complexes* in which the morphisms are all k-tame maps.

3.3. By 1.15 we also have a *k-tame category of complexes with symmetry*. For any k, we may regard complexes with simplicial symmetry as forming a full subcategory of the k-tame category of complexes with symmetry, the *k-tame category of complexes with simplicial symmetry*. It is this category that will be most important for our purposes. Explicitly, an object in this category is a triple (X, Γ, ρ) consisting of a complex and an action of a group by simplicial automorphisms; on the other hand, a morphism $(f, \omega): (X, \Gamma, \rho) \to (X', \Gamma', \rho')$ is a morphism of sets with symmetry such that f is a k-tame map.

Let X be a complex and let v be a base vertex of X. The *universal cover* of X with respect to v will be defined to be the set $\tilde{X}$ of homotopy classes of paths with initial point v. If X is connected, this is the standard definition of the universal cover; in general, it gives the universal cover (in the standard sense) of the connected component of X containing v. In any event we shall always regard $\tilde{X}$ as a geometric complex, with the unique triangulation for which the covering projection $p: \tilde{X} \to X$ is a simplicial map. Furthermore, $\tilde{X}$ will always be understood to be equipped with the standard base vertex $\tilde{v}$ which is the homotopy class of the constant path at v.

Now let $\mathscr{X} = (X, \Gamma, \rho)$ be a complex with simplicial symmetry and let v be a base vertex for X. Let Γ_0 denote the stabilizer in Γ of the connected component of X containing v. Let $\tilde{X}$ denote the universal cover of X with respect to v, and let $p: \tilde{X} \to X$ be the covering projection. We shall say that a simplicial automorphism h of $\tilde{X}$ *covers* an element $g \in \Gamma_0$ if we have $p(h(\tilde{x})) = g \cdot p(\tilde{x})$ for every point $\tilde{x} \in \tilde{X}$. Let $\tilde{\Gamma}$ denote the set of all ordered pairs of the form (g, h), where g is an element of Γ_0 and h is a simplicial automorphism of $\tilde{X}$ covering g. Then $\tilde{\Gamma}$ is a group under the composition law $(g, h) \circ (g', h') = (gg', hh')$, and a natural action of $\tilde{\Gamma}$ on $\tilde{X}$ is defined by setting $(g, h) \cdot \tilde{x} = h(\tilde{x})$. Thus we have a complex with simplicial symmetry $\widetilde{\mathscr{X}} = (\tilde{X}, \tilde{\Gamma}, \tilde{\rho})$, which we call the *universal cover* of $\mathscr{X}$ with respect to v.

3.4. We define a homomorphism $\theta: \tilde{\Gamma} \to \Gamma$ by setting $\theta(g, h) = g \in \Gamma_0 \subset \Gamma$. By elementary covering space theory, the image of θ is all of Γ_0. Its kernel is the group of covering transformations of $\tilde{X}$. Note also that $\not{p} = (p, \theta): \widetilde{\mathscr{X}} \to \mathscr{X}$ is a simplicial morphism of complexes with symmetry; we shall call it the *covering projection*.

3.5. For $i = 1, 2$, let $\mathscr{X}_i = (X_i, \Gamma_i, \rho_9)$ be a complex with symmetry and let $v_i \in X_i$ be a base vertex. Let $\widetilde{\mathscr{X}_i} = (\tilde{X}_i, \tilde{\Gamma}_i, \tilde{\rho}_i)$ denote the universal cover of $\mathscr{X}_i$ with respect to v_i, let $\tilde{v}_i$ denote the base vertex of

$\widetilde{\mathscr{X}_i}$, and let $\rho_i \colon \widetilde{\mathscr{X}_i} \to \mathscr{X}_i$ denote the corresponding covering projection. A k-tame morphism $\mathscr{f} \colon \widetilde{\mathscr{X}_1} \to \widetilde{\mathscr{X}_2}$ will be said to *cover* a k-tame morphism $\mathscr{f} \colon \mathscr{X}_1 \to \mathscr{X}_2$ if we have $\rho_2 \circ \widetilde{\mathscr{f}} = \mathscr{f} \circ \rho_1$.

It follows easily from the lifting criterion of elementary covering space theory that if $\mathscr{f} = (f, \omega) \colon \mathscr{X}_1 \to \mathscr{X}_2$ is a k-tame morphism such that $f(v_1) = v_2$, then there exists a unique k-tame morphism $\widetilde{\mathscr{f}} = (\tilde{f}, \tilde{\omega}) \colon \widetilde{\mathscr{X}_1} \to \widetilde{\mathscr{X}_2}$ which covers $\mathscr{f}$ and satisfies $\tilde{f}(\tilde{v}_1) = \tilde{v}_2$.

3.6. We now turn to the discussion of complexes associated with trees. Let T be a Λ-tree, where Λ is a subgroup of $\mathbb{R}$. Then $\mathbb{R}T$ is an $\mathbb{R}$-tree; in particular it is an $\mathbb{R}$-metric space, and will be understood to be endowed with the usual metric topology. Recall (1.14) that we regard T as a subset of $\mathbb{R}T$. Any segment in $\mathbb{R}T$ has the structure of an affine simplex of dimension ≤ 1.

By a *geometric T-complex* we mean a pair (X, ϕ), where X is a simplicial complex and $\phi \colon X \to \mathbb{R}T$ is a continuous map such that

(i) $\phi(\mathrm{sk}_0(X)) \subset T$, and

(ii) for each positive-dimensional closed simplex Δ of X, the set $\phi(\Delta)$ is a nondegenerate segment, and ϕ restricts to an affine map of Δ onto $\phi(\Delta)$.

The segment $\phi(\Delta)$ will be called the *support* of Δ and will be denoted $|\Delta|$. Note that its endpoints lie in T. We define the *length* of the simplex Δ to be the length of $|\Delta|$. The length of any simplex is a positive element of Λ.

We shall often say that X is a T-complex when it is understood which map ϕ is involved.

An *abstract T-complex* is a pair $(\mathfrak{X}, \Phi)$, where $\mathfrak{X}$ is an abstract simplicial complex and Φ is a map of the set of vertices of $\mathfrak{X}$ to the tree T such that for each simplex $\mathfrak{D}$ of $\mathfrak{X}$ the map restricts to a 1-1 map of $\mathfrak{D}$ into a segment of T. If $\mathfrak{D}$ is a simplex of the abstract T-complex $\mathfrak{X}$, the smallest segment of $\mathbb{R}T$ containing $\Phi(\mathfrak{D})$ is called the *support* of $\mathfrak{D}$ and is denoted $|\mathfrak{D}|$. Its endpoints belong to $\Phi(\mathscr{D})$.

Let $(\mathfrak{X}, \Phi)$ be an abstract T-complex, and let X denote the geometric realization of $\mathfrak{X}$. For any closed simplex Δ of X, the set of vertices of Δ is identified with a simplex $\mathfrak{D}$ of $\mathfrak{X}$. The map $\Phi|\mathfrak{D}$ extends uniquely to an affine map $\phi^\Delta \colon \Delta \to |\mathfrak{D}|$, which we may regard as a map of Δ into $\mathbb{R}T$. If Δ' is a face of Δ, it is clear that $\phi^\Delta|\Delta' = \phi^{\Delta'}$. Hence we can define a continuous map $\phi \colon X \to \mathbb{R}T$ by setting $\phi|\Delta = \phi^\Delta$ for every simplex Δ

of X. Clearly (X, ϕ) is a geometric T-complex; we call it the *geometric realization* of $(\mathfrak{X}, \Phi)$.

The term "T-complex", when unmodified, will mean "geometric T-complex".

Let (X, ϕ) and (X', ϕ') be T-complexes. By a T-*map* from X to X' we mean a 0-tame map $J: X \to X'$ such that $\phi' \circ J = \phi$. Note that for any $k \geq 0$ there is a k-*tame category of* T-*complexes*, in which the objects are T-complexes and the morphisms are the k-tame T-maps.

If (X, ϕ) is an abstract or geometric T-complex such that ϕ maps the $\mathrm{sk}_0(X)$ bijectively onto T, then we shall often identify $\mathrm{sk}_0(X)$ with T. We shall say that the T-complex (X, ϕ) is *natural* if ϕ maps $\mathrm{sk}_0(X)$ bijectively onto T and the action of $\mathrm{Aut}(T)$ on T extends to an action on X. This extended action, which is then clearly unique, will be called the *natural action* of $\mathrm{Aut}(T)$ on X. The geometric realization of a natural abstract T-complex is a natural geometric T-complex.

Note that if (X, ϕ) is a natural geometric T-complex, then the map $\phi: X \to \mathbb{R}T$ is $\mathrm{Aut}(T)$-equivariant, where $\mathrm{Aut}(T)$ is understood to act in the natural way on $\mathbb{R}T$ (see 1.17).

Note also that if (X, ϕ) is a natural geometric T-complex then any 1-simplex of X is mapped homeomorphically onto its support by ϕ.

Let (X, ϕ) and (X', ϕ') be two natural geometric T-complexes. Since ϕ and ϕ' map $\mathrm{sk}_0(X)$ and $\mathrm{sk}_0(X')$ bijectively onto T, there is a unique T-map from the $\mathrm{sk}_0(X)$ to $\mathrm{sk}_0(X')$. Hence any two 0-tame T-maps from X to X' agree on $\mathrm{sk}_0(X)$.

3.7. Let (X, ϕ) be a geometric T-complex, and let $\gamma: [0, 1] \to \mathbb{R}T$ be a parametrized segment. By a *lift* of γ to X we mean a path $\hat{\gamma}: [0, 1] \to X$ such that $\phi \circ \hat{\gamma} = \gamma$. Of course any lift of a parametrized segment in $\mathbb{R}T$ is a parametrized arc in X. For $k = 0$ or 1, a lift of γ will be called k-tame if it is a k-tame path in X. Since ϕ maps vertices of X to points of T, a necessary condition for a parametrized segment to have a 0-tame lift is that its endpoints lie in T.

If the geometric T-complex (X, ϕ) is natural then any two 0-tame lifts of a given parametrized segment must have the same endpoints.

If $\hat{\gamma}$ is a k-tame lift of γ then the inverse of $\hat{\gamma}$ is a k-tame lift of the inverse of γ. Similarly, if γ and γ' are parametrized segments such that the composition $\gamma * \gamma'$ is defined, then for any k-tame lifts $\hat{\gamma}$ and $\hat{\gamma}'$ of γ and γ', the composition $\hat{\gamma} * \hat{\gamma}'$ is defined and is a k-tame lift of $\gamma * \gamma'$.

3.8. Let (X, ϕ) and (X', ϕ') be geometric T-complexes, let $\hat{\gamma}$ be a k-tame lift to X of a parametrized segment γ in $\mathbb{R}T$, and let $J: X \to X'$ be a k-tame T-map. Then $J \circ \hat{\gamma}$ is clearly a k-tame lift of γ to X'.

3.9. The following notions will be important in §9. Let γ be a parametrized segment in $\mathbb{R}T$, and let $\hat{\gamma}$, $\hat{\gamma}'$ be 1-tame lifts of γ to X. A *T-homotopy* from $\hat{\gamma}$ to $\hat{\gamma}'$ is a homotopy $\Xi\colon [0, 1] \times [0, 1] \to X$, constant on 0 and 1, such that $\Xi_0 = \hat{\gamma}$, $\Xi_1 = \hat{\gamma}'$, and Ξ_t is a lift of γ for every $t \in [0, 1]$. The T-homotopy Ξ will be called *tame* if its image is a subcomplex of X. If there exists a tame T-homotopy from $\hat{\gamma}$ to $\hat{\gamma}'$ we shall say that $\hat{\gamma}$ and $\hat{\gamma}'$ are *tamely T-homotopic*. Tame T-homotopy is plainly an equivalence relation among 1-tame lifts of γ.

3.10. We shall now show how to associate a natural 1-dimensional T-complex with any free abelian subgroup of Λ.

Let L be a free abelian group of finite rank. Recall that an element α of L is called *unimodular* if it is an element of some basis of L. Equivalently, α is unimodular if and only if it cannot be written in the form $n\alpha_0$ with $\alpha_0 \in L$ and $n > 1$.

If L is a free abelian subgroup of Λ, we define a 1-dimensional abstract T-complex $(\mathfrak{U}, \Psi) = (\mathfrak{U}(L), \Psi^L)$ as follows. The vertices of $\mathfrak{U}$ are the points of T. A 1-simplex is a two-element subset $\{x, y\}$ of T such that $d(x, y)$ is a unimodular element of L. The map Ψ is the identity map of T. One checks immediately that $(\mathfrak{U}, \Psi)$ is a natural T-complex.

The geometric realization of $(\mathfrak{U}(L), \Psi^L)$ will be denoted $(U, \psi) = (U(L), \psi^L)$. Note that the length of any 1-simplex of U is a positive unimodular element of L.

3.11. Let $\gamma\colon [0, 1] \to \mathbb{R}T$ be a parametrized segment with endpoints in T. There is a simple way of cataloguing all 1-tame lifts of γ to U. Identifying T with $\mathrm{sk}_0(U)$, we may regard $x = \gamma(0)$ and $y = \gamma(1)$ as vertices of U. Let $\hat{\gamma}$ be a 1-tame lift of γ, and let $0 = a_0 < a_1 < \cdots < a_n = 1$ be the points of $[0, 1]$ that are mapped to vertices of U by $\hat{\gamma}$. Set $x_i = \hat{\gamma}(a_i)$ for $i = 0, \cdots, n$, so that $x_0 = x$ and $x_n = y$. Then $\Delta(x_{i-1}, x_i)$ is a 1-simplex of $\mathfrak{U}$ for $i = 1, \cdots, n$; hence $\alpha_i = d(x_{i-1}, x_i)$ is a positive unimodular element of L. Since γ is a homeomorphism of $[0, 1]$ onto the segment $[x, y]$, we have $d(x, y) = \alpha_1 + \cdots + \alpha_n$. Thus any 1-tame lift of γ determines an ordered partition of $d(x, y)$ into positive unimodular elements of L. This is at once seen to define a bijective correspondence between 1-tame lifts of γ and ordered partitions of $d(x, y)$ into positive unimodular elements of L.

In particular, a parametrized segment γ in $\mathbb{R}T$ admits a 1-tame lift to U if and only if the endpoints x and y of γ lie in T and $d(x, y)$ is an element of L.

If $d(x, y)$ does belong to L, there is a unique way to write $d(x, y)$ in the form $n\alpha$, where α is unimodular and n is a nonnegative integer.

The 1-tame lift γ corresponding to the partition $d(x, y) = \alpha + \cdots + \alpha$ will be called the *standard lift* of γ to $U(L)$.

If L is free abelian of rank 1, there is only one way to write an element of L as a sum of positive unimodular elements. Hence a parametrized segment in $\mathbb{R}T$ can have only the standard lift to U.

3.12. Now let L and L' be free abelian subgroups of Λ with $L \subset L'$. We shall define a canonical 1-tame T-map $H = H_{L, L'} : U(L) \to U(L')$.

Identifying the 0-skeletons of $U(L)$ and $U(L')$ with T, we define H to be the identity on $\mathrm{sk}_0(U(L))$. If τ is any 1-simplex of $U(L)$, fix a homeomorphism $j \colon [0, 1] \to \tau$. Then $\gamma = \psi \circ j$ is a parametrized segment whose endpoints $x = \gamma(0)$ and $y = \gamma(1)$ lie in T. We have $d(x, y) \in L \subset L'$. Hence there is a standard lift $\hat{\gamma}$ of γ to $U(L')$. Define $H^{\tau} = \hat{\gamma} \circ j^{-1}$. Since $\hat{\gamma}$ is a 1-tame lift of γ, we have $\psi \circ H^{\tau} = \psi | \tau$, and H^{τ} fixes the endpoints of τ. It is clear that H^{τ} is independent of the choice of the homeomorphism j. Hence we can extend the identity map of $\mathrm{sk}_0(U(L))$ to a 1-tame T-map $H \colon (U(L) \to U(L')$ by setting $H | \tau = H^{\tau}$ for every 1-simplex τ of $U(L)$.

Since the definition of H is canonical, it is clear that H is $\mathrm{Aut}(T)$-equivariant with respect to the natural actions of $\mathrm{Aut}(T)$ on $U(L)$ and $U(L')$.

3.13. Let $\mathscr{T} = (T, \Gamma, \rho)$ be a Λ-tree with symmetry, where Λ is a subgroup of $\mathbb{R}$. We define a (geometric) $\mathscr{T}$-*complex* to be a quintuple $\mathscr{X} = (X, \Gamma_0, \rho_0, \phi, \omega)$, where (X, Γ_0, ρ_0) is a complex with simplicial symmetry, $\phi \colon X \to T$ is a map that makes (X, ϕ) a geometric T-complex, and $\omega \colon \Gamma_0 \to \Gamma$ is a group homomorphism that makes (ϕ, ω) a morphism of sets with symmetry.

For $i = 0, 1$, let $\mathscr{X}_i = (X_i, \Gamma_i, \rho_i, \phi_i, \omega_i)$ be a $\mathscr{T}$-complex. A $\mathscr{T}$-*map* from $\mathscr{X}_0$ to $\mathscr{X}_1$ is a 0-tame morphism of complexes with symmetry $\mathscr{J} = (J, \kappa) \colon (X_0, \Gamma_0, \rho_0) \to (X_1, \Gamma_1, \rho_1)$ such that J is a T-map and $\omega_1 \circ \kappa = \omega_0$. (In particular it follows that $(\phi_1, \omega_1) \circ (J, \kappa) = (\phi_0, \omega_0)$.)

For any fixed $\mathscr{T}$, and for any $k \geq 0$, we have a k-*tame category of* $\mathscr{T}$-*complexes* in which the morphisms are all k-tame $\mathscr{T}$-maps.

3.14. Suppose that we are given a group $\Lambda \subset \mathbb{R}$, a Λ-tree with symmetry $\mathscr{T} = (T, \Gamma, \rho)$ and a natural T-complex (X, ϕ). There is a unique action $\overline{\rho}$ of Γ on X that makes ϕ a Γ-equivariant map. Thus $(X, \Gamma, \overline{\rho}, \phi, 1)$ is a $\mathscr{T}$-complex.

If (X_0, ϕ_0) and (X_1, ϕ_1) are two natural T-complexes and $J \colon X_0 \to X_1$ is an $\mathrm{Aut}(T)$-equivariant T-map, then $(J, 1)$ is a $\mathscr{T}$-map between the corresponding $\mathscr{T}$-complexes.

3.15.   Let $\mathcal{T}$ be any Λ-tree with symmetry, let $\mathcal{X} = (X, \Gamma_0, \rho_0, \phi, \omega)$ be an arbitrary $\mathcal{T}$-complex, and let v be a base vertex of X. Let $(\tilde{X}, \tilde{\Gamma}_0, \tilde{\rho}_0)$ denote the universal cover of the complex with symmetry (X, Γ_0, ρ_0), let $\tilde{v}$ denote the base vertex of $\tilde{X}$, and let (p, θ) denote the covering projection. If we set $\tilde{\phi} = \phi \circ p$ and $\tilde{\omega} = \omega \circ \theta$, then $\widetilde{\mathcal{X}} = (\tilde{X}, \tilde{\Gamma}_0, \tilde{\rho}_0, \tilde{\phi}_0, \tilde{\omega})$ is a $\mathcal{T}$-complex. It is called the *universal cover* of $\mathcal{X}$ with respect to v.

Now for $i = 0, 1$ let $\mathcal{X}_i = (X_i, \Gamma_i, \rho_i, \phi_i, \omega_i)$ be a $\mathcal{T}$-complex and let v_i be a base vertex in X_i. Let $\mathcal{J} : \mathcal{X}_0 \to \mathcal{X}_1$ be a k-tame $\mathcal{T}$-map such that $\mathcal{J}(v_0) = v_1$. Let $\widetilde{\mathcal{X}}_i = (\tilde{X}_i, \tilde{\Gamma}_i, \tilde{\rho}_i, \tilde{\phi}_i, \tilde{\omega}_i)$ denote the universal cover of $\mathcal{X}_i$ with respect to v_i. By 3.5, there is a unique base-point-preserving k-tame morphism of complexes with symmetry $\widetilde{\mathcal{J}} : (\tilde{X}_0, \tilde{\Gamma}_0, \tilde{\rho}_0) \to (\tilde{X}_1, \tilde{\Gamma}_1, \tilde{\rho}_1)$ which covers $\mathcal{J}$. It follows from the definitions $\widetilde{\mathcal{J}}$ is a $\mathcal{T}$-map; it is said to be *induced by* $\mathcal{J}$.

4. Proof of the structure theorem in the rank-1 case

In this section we shall give the proof of Theorem 2.1. The method of proof is to write the value group as a monotone union of rank-1 free abelian subgroups L_i, and to use the construction of 3.10 to define a direct system $(\mathcal{U}_i, \mathcal{H}_{ij})$ of 1-dimensional $\mathcal{T}$-complexes corresponding to the L_i. One then replaces the system $(\mathcal{U}_i, \mathcal{H}_{ij})$ by a direct system of 1-connected 1-dimensional $\mathcal{T}$-complexes and 1-tame maps via the universal covering construction of 3.15, and interprets the 1-skeletons of the latter complexes as L_i-trees; this gives the directed system of pretrees with symmetry $\mathcal{T}_i$ appearing in the statement of 2.1.

Much of the relevant material on direct systems of $\mathcal{T}$-complexes will be needed again in §11 to prove the analogue of 2.1 in the case where Λ has Q-rank 2. In order to avoid extensive repetition, we shall discuss this material in a form that is general enough to cover both the rank-1 and rank-2 cases. The study of direct systems of $\mathcal{T}$-complexes will occupy most of the present section.

4.1.   Let (X_i, J_{ij}) be any direct system (see 1.29) in the 0-tame category of geometric simplicial complexes, and let v_i be a base vertex in X_i for each i. We shall say that the v_i are *consistent* if $J_{ij}(v_i) = v_j$ whenever $i \leq j$. If the X_i are nonempty, it is obvious that consistent vase vertices exist for the system.

4.2.   Let Λ be a subgroup of $\mathbb{R}$ which is exhibited as a monotone union $\bigcup_{i=1}^{\omega} L_i$, where each L_i is a free abelian subgroup of Λ and $L_i \subset L_{i+1}$.

(The inclusions need not be proper.) Let T be a Λ-tree. By 3.10, for each $i \geq 1$, we have a T-complex $U_i = U(L_i)$, and by 3.12, for $1 \leq i \leq j$ we have a 1-tame T-map $H_{ij} = H_{L_i, L_j} : U_i \to U_j$. Thus we have a direct system (U_i, H_{ij}) in the 0-tame category of geometric simplicial complexes. Let us choose consistent base vertices for the U_i. In terms of these base vertices we can form the direct limit $\varinjlim \pi_n(U_i)$ for any $n \geq 0$. For $n > 1$, we have $\pi_n(U_i) = 0$ for all i since the U_i are 1-dimensional. Thus the interesting cases are $n = 0$ and 1.

4.3. Proposition. *We have* $\varinjlim \pi_0(U_i) = 0$.

Proof. We must show that for any $i \geq 1$ and any vertices x, y of U_i, there is some $j \geq i$ such that $H_{ij}(x)$ and $H_{ij}(y)$ are in the same connected component of U_j. We can regard x and y as points of T. Let γ be a parametrization of the segment $[x, y] \subset \mathbb{R}T$. Choose $j \geq i$ so that $d(x, y) \in L_j$. Then the standard lift of γ to U_j is an arc joining $H_{ij}(x)$ to $H_{ij}(y)$ in U_j. q.e.d.

In general, $\varinjlim \pi_1(U_i)$ need not be 0. However, we shall see (4.5) that it is 0 when the $\mathbb{Q}$-rank of Λ is 1 (in which case L_i has rank 1 for large i). This is a key step in the proof of Theorem 2.1. We shall prove it as an application of a more general result which will be used again in §9.

4.4. Consider the following general situation. Let Λ be a subgroup of $\mathbb{R}$, let T be a Λ-tree, and let $((X_i, \phi_i); J_{ij})$ be a direct system in the 0-tame category of T-complexes. Let us fix consistent base vertices for the X_i. The following result gives a general criterion for triviality of $\varinjlim \pi_1(X_i)$.

Proposition. *Suppose that the direct system* $((X_i, \phi_i); J_{ij})$ *satisfies the following conditions:*

(1) *Every arc in $\mathbb{R}T$ with endpoints in T admits a tame lift to some X_i.*

(2) *Given any $i \geq 1$, any arc γ in T, and any two tame lifts $\hat{\gamma}$ and $\hat{\gamma}'$ of γ to X_i, the paths $J_{ij} \circ \hat{\gamma}$ and $J_{ij} \circ \hat{\gamma}'$ are homotopic in X_j for some $j \geq i$.*

Then we have $\varinjlim \pi_1(X_i) = 0$.

Proof. Suppose that (1) and (2) hold. Then we shall prove the following assertion:

(4.4.1). *Let $i \geq 1$ be given, and let $\gamma: [0, 1] \to X_i$ be a path whose endpoints are vertices of X_i. Then for some $j \geq i$ (and hence for all sufficiently large j, see 3.8), the path $J_{ij} \circ \gamma$ is homotopic in X_j either to a constant or to a 1-tame lift of some parametrized segment in $\mathbb{R}T$.*

It follows from (4.4.1) that if γ is a closed loop γ based at a vertex of X_i then $J_{ij}(\gamma)$ is homotopic to a constant. Hence proving (4.4.1) will complete the proof of the proposition

Any path between two vertices in a complex is homotopic either to a constant or to a path of the form $\gamma_1 * \cdots * \gamma_r$, where $r > 0$, and where each γ_p, $1 \leq p \leq r$, is an affine homeomorphism of a closed interval onto a 1-simplex in X_i. Hence it is enough to prove (4.4.1) for paths of the latter form. For such paths we shall proceed by induction on r.

If $r = 1$ then γ is an affine homeomorphism of an interval onto a 1-simplex τ in X. By the definition of a T-complex, ϕ maps τ homeomorphically onto a segment in $\mathbb{R}T$; thus γ is itself a 1-tame lift of a parametrized segment in $\mathbb{R}T$.

Now suppose that $r > 1$. Then the induction hypothesis implies that γ has the form $\zeta * \varepsilon$, where for sufficiently large k each of the paths $J_{ik} \circ \zeta$ and $J_{ik} \circ \varepsilon$ is homotopic in X_k to a constant or to a 1-tame lift of a parametrized segment in $\mathbb{R}T$. We may assume that $J_{ik} \circ \zeta$ and $J_{ik} \circ \varepsilon$ are homotopic to 1-tame lifts $\hat{\omega}$ and $\hat{\eta}$ of parametrized segments ω and η in $\mathbb{R}T$. Thus $J_{ik}(\gamma)$ is homotopic to $\hat{\omega} * \hat{\eta}$. Let x and y denote, respectively, the initial point of ω and the terminal point of η. The terminal point of ω, which is also the initial point of η, will be denoted by u.

Since T is a tree, we have $[x, u] \cap [y, u] = [z, u]$ for some $z \in T$. We have $[x, z] \cap [z, y] = \{z\}$, and hence $[x, z] \cup [z, y] = [x, y]$. Let us consider the case where x, y, u and z are all distinct. We can regard ω as a reparametrization of $\omega_1 * \omega_2$, where ω_1 is a parametrization of the segment $[x, z] \subset \mathbb{R}T$, and ω_2 is a parametrization of the segment $[z, u]$. Likewise, η is a reparametrization of $\eta_1 * \eta_2$, where η_1 is a parametrization of the segment $[u, z]$ and ω_2 is a parametrization of the segment $[x, y]$. Since ω_2 and η_1 are homeomorphisms of $[0, 1]$ onto the segment $[u, z]$, the path η_1 is a reparametrization of the inverse of ω_2.

Since $[x, z] \cap [z, y] = \{z\}$ and $[x, z] \cup [z, y] = [x, y]$, the path $\omega_1 * \eta_2$ is a parametrization of the segment $[x, y]$. By hypothesis (1) and 3.8, for sufficiently large k the parametrized segments ω_i and η_i admit 1-tame lifts $\hat{\omega}_1$ and $\hat{\eta}_i$ to X_k. By 3.7 we may take $\hat{\eta}_1$ to be a reparametrization of the inverse of $\hat{\omega}_2$. In particular, $\hat{\omega}_2 * \hat{\eta}_1$ is homotopic to a constant.

Now by 3.7, the parametrized segment ω has a 1-tame lift which is a reparametrization of $\hat{\omega}_1 * \hat{\omega}_2$ and is therefore homotopic to $\hat{\omega}_1 * \hat{\omega}_2$. Hypothesis (2) then implies that for large enough $j \geq k$, the paths $J_{kj}(\hat{\omega})$ and $J_{kj}(\hat{\omega}_1 * \hat{\omega}_2)$ are homotopic. Similarly, for large enough j the paths

$J_{kj}(\hat{\eta})$ and $J_{kj}(\hat{\eta}_1 * \hat{\eta}_2)$ are homotopic. Hence $J_{ij}(\gamma)$ is homotopic to $J_{kj}(\hat{\omega}_1 * \hat{\omega}_2 * \hat{\eta}_1 * \hat{\eta}_2)$, and therefore to $J_{kj}(\hat{\omega}_1 * \hat{\eta}_2)$. But by 3.8 the latter path is a 1-tame of the parametrized segment $\omega_1 * \eta_2$.

This completes the induction in the case where x, y, z and w are all distinct. The other cases are handled by exactly the same method and are slightly simpler.

4.5. Corollary. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 1. Let us write Λ as a monotone union $\bigcup_{i=1}^{\infty} L_i$, where each L_i is infinite cyclic and $L_i \subset L_{i+1}$. Then with the conventions of 4.2, we have $\varinjlim \pi_1(U_i) = 0$.*

Proof. If a parametrized segment γ in $\mathbb{R}T$ has endpoints x, $y \in T$, then $d(x, y) \in \Lambda$ and hence $d(x, y) \in L_i$ for some i. Hence by 3.11, γ admits a 1-tame lift to U_i. Since L_i is free abelian of rank 1, it follows from 3.11 that the lift of γ is unique. The result therefore follows from Proposition 4.4.

4.6. Now let us consider a Λ-tree with symmetry $\mathscr{T} = (T, \Gamma, \rho)$, where Λ is a subgroup of $\mathbb{R}$, and a direct system $(\mathscr{X}_i; \mathscr{J}_{ij})$ in the 0-tame category of $\mathscr{T}$-complexes. Let $\mathscr{X}_i = (X_i, \Gamma_i, \rho_i, \phi_i, \omega_i)$ and $\mathscr{J}_{ij} = (J_{ij}, \kappa_{ij})$. We shall say that $(\mathscr{X}_i; \mathscr{J}_{ij})$ is *abundant* if it satisfies the following conditions.

(i) For every point $x \in T$ there exist an index i and a vertex v of X_i such that $\phi_i(v) = x$.

(ii) Given an index i, a parametrized segment γ in $\mathbb{R}T$, and two vertices v_0 and v_1 of X_i such that $\phi_i(v_t) = \gamma(t)$ for $t = 0, 1$, there exist an index $j \geq i$ and a 0-tame lift $\hat{\gamma}$ of γ to X_j such that $\hat{\gamma}(t) = J_{ij}(v_t)$ for $t = 0, 1$.

(iii) The natural map from $\varinjlim \Gamma_i$ to Γ induced by the ω_i is an isomorphism.

4.7. In order to motivate the above definition, let us consider a group $\Lambda \subset \mathbb{R}$ which is given in the form $\bigcup_{i=1}^{\infty} L_i$, where each L_i is a free abelian subgroup of Λ and $L_i \subset L_j$ whenever $i \leq j$. Let $\mathscr{T} = (T, \Gamma, \rho)$ be any Λ-tree with symmetry. Let us write (U_i, ψ^i) for the T-complex $(U(L_i), \psi^{L_i})$. As in 4.2, setting $H_{ij} = H_{L_i, L_j}$, we have a direct system (U_i, H_{ij}) of T-complexes and 1-tame maps. Since the U_i are natural T-complexes, and since the H_{ij} are $\mathrm{Aut}(T)$-equivariant, according to 3.14 we have a $\mathscr{T}$-complex $\mathscr{U}_i = (U_i, \Gamma, \rho_i, \psi_i, 1)$ for each i, and a 1-tame morphism $\mathscr{H}_{ij} = (H_{ij}, 1): \mathscr{U}_i \to \mathscr{U}_j$ of $\mathscr{T}$-complexes whenever $i \leq j$. Thus we have a direct system $(\mathscr{U}_i; \mathscr{H}_{ij})$ in the 1-tame category of $\mathscr{T}$-complexes.

Proposition. *The direct system* $(\mathscr{U}_i\,;\mathscr{H}_{ij})$ *is abundant.*

Proof. Since the U_i are natural, each ψ_i maps $\mathrm{sk}_0(U_i)$ bijectively to T; in particular, condition (i) of 4.6 holds. That condition (ii) holds is immediate from the proof of 4.3. Condition (iii) is trivial in this case since we have $\Gamma_i = \Gamma$ and $\omega_i = \kappa_{ij} = 1$ for all i and j. q.e.d.

A key step in the proof of Theorem 2.1 will be to show that the universal cover of the $\mathscr{T}$-complexes in a certain abundant direct system again form an abundant direct system. This will follow from Proposition 4.8 below, which will also be used in §11.

4.8. Let Λ be a subgroup of $\mathbb{R}$, let $\mathscr{T}$ be a Λ-tree with symmetry, and let $(\mathscr{X}_i, \mathscr{F}_{ij})$ be a direct system in the k-tame category of $\mathscr{T}$-complexes for some $k \geq 0$. Suppose that we are given consistent base vertices v_i for the $\mathscr{X}_i$. Let $\widetilde{\mathscr{X}}_i$ denote the universal cover of $\mathscr{X}_i$ with respect to v_i. By 3.15, the $\mathscr{F}_{ij}$ induce k-tame T-maps $\widetilde{\mathscr{F}}_{ij} \colon \widetilde{\mathscr{X}}_i \to \widetilde{\mathscr{X}}_j$, so that we have a direct system $(\widetilde{\mathscr{X}}_i, \widetilde{\mathscr{F}}_{ij})$ in the k-tame category of $\mathscr{T}$-complexes. We shall call it the *universal cover* of the direct system $(\mathscr{X}_i, \mathscr{F}_{ij})$ with respect to the v_i.

Proposition. *Let* Λ *be a subgroup of* $\mathbb{R}$, *let* $\mathscr{T} = (T, \Gamma, \rho)$ *be a* Λ-*tree with symmetry, and let* $(\mathscr{X}_i, \mathscr{F}_{ij})$ *be an abundant direct system in the* 0-*tame category of* $\mathscr{T}$-*complexes. Let* X_i *denote the underlying (geometric) simplicial complex of* $\mathscr{X}_i$, *and let* $v_i \in X_i$ *be consistent base vertices. Suppose that for* $n = 0$, 1, *we have* $\varinjlim \pi_n(X_i, v_i) = 0$. *Then the universal cover of* $(\mathscr{X}_i, \mathscr{F}_{ij})$ *with respect to the* v_i *is an abundant direct system.*

Proof. Let $\mathscr{X}_i = (X_i, \Gamma_i, \rho_i, \phi_i, \omega_i)$ and $\mathscr{F}_{ij} = (J_{ij}, \kappa_{ij})$. Let V_i denote the component of X_i containing v_i.

To check condition (i) of 4.6, consider an arbitrary point $x \in T$. Since $(\mathscr{X}_i, \mathscr{F}_{ij})$ is abundant, there exist an index l and a vertex v of X_l such that $\phi_l(v) = x$. Since $\varinjlim \pi_0(X_i, v_i) = 0$, there is an index $i \geq l$ such that $v \in V_i$. If $\tilde{v}$ is a vertex of $\widetilde{X}_i$ such that $p_i(\tilde{v}) = v$, then $\tilde{\phi}_i(\tilde{v}) = x$. This proves (i).

Next we verify (ii). We are given an index i, a parametrized segment γ in $\mathbb{R}T$, and two vertices $\tilde{v}_0$ and $\tilde{v}_1$ of $\widetilde{X}_i$ such that $\tilde{\phi}_i(\tilde{v}_t) = \gamma(t)$ for $t = 0$, 1. Set $v_t = \tilde{\phi}_i(\tilde{v}_t)$. Since $(\mathscr{X}_i, \mathscr{F}_{ij})$ is abundant, there exist an index $k \geq i$ and a 0-tame lift $\hat{\gamma}$ of γ to X_k such that $\hat{\gamma}(t) = J_{ik}(v_t)$ for $t = 0$, 1. Since $\varinjlim \pi_0(X_i, v_i) = 0$, there is an index $k' \geq k$ such that $\phi_{k'}(\hat{\gamma}([0, 1])) \subset V_{k'}$. Hence we may assume k to be chosen so that $\phi_k(\hat{\gamma}([0, 1])) \subset V_k$.

Now let $\tilde{\gamma}_0$ be a lift of $\hat{\gamma}$ to $\widetilde{X}_k$ such that $\tilde{\gamma}_0(0) = \tilde{v}_0$. Set $\tilde{v}_1' = \gamma_0(1)$. Then $\tilde{v}_1'$ and $\tilde{v}_1$ are in the same fiber in the covering space $\widetilde{X}_k$ of V_k. Since $\varinjlim \pi_1(X_i, v_i) = 0$, there is an index $j \geq k$ such that $J_{kj}(\tilde{v}_1') = J_{kj}(\tilde{v}_1)$. The path $\tilde{\gamma} = J_{kj}(\tilde{\gamma}_0)$ is a 0-tame lift of γ to $\widetilde{X}_j$ such that $\hat{\gamma}(t) = J_{ij}(v_t)$ for $t = 0$, 1. This establishes (ii).

To prove (iii) we first show that the natural map $\omega \colon \varinjlim \Gamma_i \to \Gamma$ is surjective. Consider an arbitrary element $g \in \Gamma$. Since $(\mathscr{X}_i, \mathscr{J}_{ij})$ is abundant, there is an index l such that $g = \omega_l(g_l)$ for some g_l in Γ_l. Since $\varinjlim \pi_0(X_i, v_i) = 0$, there is an index $k \geq l$ such that $g_k = k_{lk}(g_l)$ belongs to the stabilizer G_l of V_l in $\Gamma_{l'}$. Since by 3.4 the natural homomorphism $\theta_l \colon \widetilde{\Gamma}_l \to \Gamma_l$ has image G_l, it follows that $g \in \omega_l(\widetilde{\Gamma}_l)$; this shows that ω is surjective.

It remains to prove that ω is injective. Consider an index i and an element $\tilde{g}$ of $\widetilde{\Gamma}_i$ such that $\omega_i(\tilde{g}) = 1$. Set $g = \theta_i(\tilde{g}) \in G_i$. Since $(\mathscr{X}_i, \mathscr{J}_{ij})$ is abundant, there is an index $k \geq i$ such that $\kappa_{ik}(g) = 1$. Hence $\tilde{\kappa}_{ik}(\tilde{g})$ is a covering transformation in the covering space $\widetilde{X}_k$ of V_k. Since $\varinjlim \pi_1(X_i, v_i) = 0$ it follows that $\tilde{\kappa}_{ij}(\tilde{g}) = 1$ for some $j \geq k$. Thus ω is injective. q.e.d.

4.9. *Proof of Theorem* 2.1. We are given a Λ-tree with symmetry $\mathscr{T} = (T, \Gamma, \rho)$, where $\Lambda \subset \mathbb{R}$ has $\mathbb{Q}$-rank 1. Let us write Λ as a union $\bigcup_{i=1}^{\infty} L_i$, where each L_i is an infinite cyclic subgroup of Λ and $L_i \subset L_{i+1}$. We shall denote the positive unimodular element of L_i by λ_i. Using the L_i we can define, as in 4.7, a direct system $(\mathscr{U}_i; \mathscr{H}_{ij})$ in the 1-tame category of $\mathscr{T}$-complexes, where $\mathscr{U}_i = (U_i, \Gamma, \rho_i, \psi_i, 1)$ and $\mathscr{H}_{ij} = (H_{ij}, 1)$. Recall that this system is indexed by the positive integers, that the U_i are 1-dimensional, and that the system is abundant by Proposition 4.7.

As we observed in (4.1), we can choose consistent base vertices v_i in the U_i. In terms of the v_i we can form the universal cover $(\widetilde{\mathscr{U}}_i; \widetilde{\mathscr{H}}_{ij})$ of $(\mathscr{U}_i; \mathscr{H}_{ij})$, which is again a direct system in the 1-tame category of $\mathscr{T}$-complexes. Let us write $\widetilde{\mathscr{U}}_i = (\widetilde{U}_i, \widetilde{\Gamma}_i, \tilde{\rho}_i, \tilde{\psi}_i, \tilde{\omega}_i)$ and $\widetilde{\mathscr{H}}_{ij} = (\widetilde{H}_{ij}, \tilde{\omega}_{ij})$. According to the definitions, $(\widetilde{U}_i, \widetilde{\Gamma}_i, \tilde{\rho}_i)$ is the universal covering space of (U_i, Γ, ρ_i), where V_i is the component of U_i containing the vertex v_i. For each i we let $\wp_i = (p_i, \theta_i) \colon (\widetilde{U}_i, \widetilde{\Gamma}_i, \tilde{\rho}_i) \to (U_i, \Gamma, \rho_i)$ denote the covering projection. Note that since the group homomorphism underlying each $\mathscr{H}_{ij}$ is the identity, we have $\tilde{\omega}_i = \theta_i$.

By 4.3 and 4.5 we have $\lim_{i \to \infty} \pi_n(U_i) = 0$ for $n = 0$, 1. Hence by 4.8, the direct system $(\widetilde{\mathscr{U}}_i; \widetilde{\mathscr{H}}_{ij})$ is abundant.

The $\widetilde{U}_i$ are 1-connected 1-complexes. Hence each $T_i = \mathrm{sk}_0(\widetilde{U}_i)$ is a $\mathbb{Z}$-tree under the standard metric in which the distance between two vertices $\tilde{v}$ an $\tilde{w}$ is 1 if and only if $\tilde{v}$ and $\tilde{w}$ are joined by an edge. Multiplying this standard metric by λ_i, we obtain a metric d_i that makes T_i an L_i-tree. The action $\tilde{\rho}_i$ of $\widetilde{\Gamma}_i$ on $\widetilde{U}_i$ restricts to an action on T_i. Thus $\mathscr{T}_i = (T_i, \widetilde{\Gamma}_i, \tilde{\rho}_i)$ is an L_i-tree with symmetry.

Since $\psi_i \colon U_i \to \mathbb{R}T$ maps $\mathrm{sk}_0(U_i)$ into T, the map $\tilde{\psi}_i$ restricts to a map $f_i \colon T_i \to T$ for each i. Likewise, since the H_{ij} are 1-tame and hence 0-tame, $\widetilde{H}_{ij}$ restricts to a map $f_{ij} \colon T_i \to T_j$.

It follows from the definition of a $\mathscr{T}$-map that $(\tilde{\psi}_i \tilde{\omega}_i)$ and (H_{ij}, ω_{ij}) are morphisms of complexes with symmetry. Hence $f_i = (f_i, \tilde{\omega}_i)$ and $f_{ij} = (f_{ij}, \tilde{\omega}_{ij})$ are morphisms of sets with symmetry.

We claim that f_i is a morphism of Λ-pretrees, so that f_i is a morphism of Λ-pretrees with symmetry. Since T_i is defined from the 1-connected 1-complex $\widetilde{U}_i$, this amounts to saying that if two vertices $\tilde{v}$ and $\tilde{w}$ are joined by an edge in $\widetilde{U}_i$, then $d(\tilde{\psi}_i(\tilde{v}), \tilde{\psi}_i(\tilde{w})) = \lambda_i$. But if $\tilde{v}$ and $\tilde{w}$ are joined by an edge, the vertices $v = p_i(\tilde{v})$ and $w = p_i(\tilde{w})$ in U_i are joined by an edge in U_i. By the definition of U_i this means that $d(\psi_i(v), \psi_i(w)) = d(\tilde{\psi}_i(\tilde{v}), \tilde{\psi}_i(\tilde{w}))$ is a unimodular element of L_i, so that $d(\tilde{\psi}_i(\tilde{v}), \tilde{\psi}_i(\tilde{w})) = \lambda_i$. This proves the claim.

We claim that the maps $f_{ij} \colon T_i \to T_j$ are also morphisms of Λ-pretrees, so that the $f_{ij} = (f_{ij}, \tilde{\omega}_{ij})$ are morphisms of Λ-pretrees with symmetry. To prove this it is enough to show that if two vertices $\tilde{v}$ and $\tilde{w}$ are joined by an edge in $\widetilde{U}_i$, then the distance in T_j between $f_{ij}(v)$ and $f_{ij}(w)$ is λ_i. Let us write $\lambda_i = n\lambda_j$, where n is a positive integer. We must show that the unique simplicial arc in $\widetilde{U}_j$ joining $\widetilde{H}_{ij}(\tilde{v})$ to $\widetilde{H}_{ij}(\tilde{w})$ consists of exactly n closed 1-simplices.

Let $\tilde{\tau}$ denote the closed edge in $\widetilde{U}_i$ joining $\tilde{v}$ to $\tilde{w}$. Set $v = p_i(\tilde{v})$, $w = p_i(\tilde{w})$, and $\tau = p_i(\tilde{\tau})$. It follows from the definition (3.12) of H_{ij} that $H_{ij}(\tau)$ is the image of the standard lift to U_j of a parametrization of the segment $[v, w] \subset \mathbb{R}T$. Hence $H_{ij}(\tau)$ is a simplicial arc in U_i consisting of exactly n closed 1-simplices. It follows that $\widetilde{H}_{ij}(\tilde{\tau})$ is a simplicial arc in $\widetilde{U}_i$ consisting of exactly n closed 1-simplices. This proves the claim.

We now have a direct system $(\mathscr{T}_i; f_{ij})$ of Λ-pretrees with symmetry, and we have morphisms $f_i \colon \mathscr{T} \to \mathscr{T}$. We shall prove the first assertion of Theorem 2.1 by showing that $(\mathscr{T}_i; f_{ij})$ converges strongly and that its limit can be identified isomorphically with $\mathscr{T}$ in such a way that the f_i

are the canonical morphisms. To do this we must check that conditions (i)–(iv) of 1.26 are satisfied.

Condition (i) holds since $\widetilde{\mathscr{H}}_{ij}$ is a $\mathscr{T}$-map for $1 \le i \le j$. The verification of conditions (ii)–(iv) uses the fact that the direct system $(\widetilde{\mathscr{U}}_i; \widetilde{\mathscr{H}}_{ij})$ is abundant. Indeed, conditions (ii) and (iv) of 1.26 follow respectively from conditions (i) and (iii) of Definition 4.6. In order to check condition (iii) of 1.26 we must consider an integer $i \ge 1$ and two points $\tilde{v}_0, \tilde{v}_1 \in T_i$, which we may regard as vertices of $\widetilde{U}_i$. For $t = 0, 1$, set $v_t = p_i(\tilde{v}_t)$ and $x_t = \tilde{\phi}_i(\tilde{v}_t) = \phi_i(v_t)$. Let γ be a parametrized segment in $\mathbb{R}T$ with $\gamma(t) = x_t$.

By condition (ii) of 4.6, there exist an index $j \ge i$ and a 0-tame lift $\hat{\gamma}$ of γ to $\widetilde{U}_j$ such that $\hat{\gamma}(t) = \widetilde{H}_{ij}(\tilde{v}_t) = f_{ij}(\tilde{v}_t)$ for $t = 0, 1$. The path $p_j \circ \hat{\gamma}$ is a 0-tame lift of γ to U_j. By 3.11, it must be the standard lift. Hence the arc $p_j(\hat{\gamma}([0, 1])) \subset U_j$ is made up of n closed simplices, where n is the nonnegative integer defined by $d(x, y) = n\lambda_j$. Therefore the arc $\hat{\gamma}([0, 1]) \subset \widetilde{U}_j$ is also made up of n closed simplices. By the definition of the metric in T_j this means that $d(f_{ij}(\tilde{v}_0), f_{ij}(\tilde{v}_1)) = n\lambda_j = d(x, y)$. This establishes condition (iii) of 1.26 and proves the first assertion of Theorem 2.1.

To prove the second assertion of the theorem, we recall that $\hat{\omega}_i = \theta_i$, and that by 3.4, the kernel of θ_i is the group of covering transformations of $\widetilde{U}_i$. Thus $\ker \hat{\omega}_i$ acts freely on $\widetilde{U}_i$ and hence on $\mathbb{R}T$.

5. Measured foliations

It was recognized by the authors of [14] that Thurston's theory of measured foliations on surfaces gives a natural class of group actions on $\mathbb{R}$-trees. As we explained in the introduction, the actions arising from measured foliations will play a central role in the structure theorem for group actions on Λ-trees where $\Lambda \subset \mathbb{R}$ and $\mathbb{Q}$-rank 2.

It will be necessary to rework some of the material developed in [14]. One reason is that the presentation in [14] is not in terms of measured foliations, which for our purposes are the natural objects, but in terms of measured geodesic laminations; these objects are in a sense equivalent to measured foliations, but the equivalence is nontrivial. A second reason for reworking the material in [14] is that we need it in a more general context: our foliations are defined not on 2-manifolds but on spaces that we call "singular surfaces", which are locally Euclidean of dimension at most 2 except at isolated points. We begin with a discussion of such spaces; for convenience we shall take them to be triangulated.

5.1.   If δ is an i-simplex in a geometric simplicial complex Y, we define the *valence* if δ (in Y) to be the (cardinal) number of $(i + 1)$-simplices in Y of which δ is a face. The valence will be denoted $\mathrm{val}_Y(\delta)$. Since our simplicial complexes are not assumed to be locally finite, $\mathrm{val}_Y(\delta)$ may be a finite or infinite cardinal. A simplex δ is said to be *isolated* (in Y) if $\mathrm{val}_Y(\delta) = 0$, and *free* (in Y) if $\mathrm{val}_Y(\delta) = 1$.

A geometric simplicial complex Σ will be called a (*triangulated*) *singular surface* if Σ has dimension at most 2 and if every 1-simplex of Σ has valence 0 or 2.

If Σ is a singular surface, then the link of an arbitrary vertex of Σ is a complex of dimension ≤ 1 in which every 0-simplex has valence 0 or 2. Hence the components of the link of a vertex are points and combinatorial 1-manifolds. The link may have infinitely many components, and the 1-manifold components may include complexes homeomorphic to either $\mathbb{R}$ or S^1.

Thus a singular surface Σ is locally Euclidean of dimension of most 2 at all points except vertices.

If the link of every vertex of Σ is S^1 (so that Σ is locally Euclidean of dimension 2) we call Σ a (*triangulated*) *surface*. If the link of every vertex is homeomorphic to either S^1 or $\mathbb{R}$ we call Σ a (*triangulated*) *surface with points at infinity*.

If Σ is a surface with points at infinity, and if we let E denote the set of *points at infinity* of Σ—i.e., vertices of Σ whose links are homeomorphic to $\mathbb{R}$—then $\overset{\circ}{\Sigma} = \Sigma - E$ is a 1-connected 2-manifold; we call it the *interior* of Σ.

5.2.   Let x be a point of singular surface Σ. Let S denote the closed star of x, i.e., the union of all the closed simplices of Σ containing x. By a *local branch* of Σ at x we shall mean the closure of a connected component of $S - x$.

If x is a vertex of Σ, we may identify S with the cone over the link of x. Hence in this case each local branch of Σ is identified with the cone over a connected component of the link of x. Thus a local branch at any vertex x is either a closed arc with x as an endpoint, a closed disk with x as an interior point, or a homeomorphic copy of $\{0\} \cup (D^2 \cap H^2) \subset \mathbb{R}^2$, where D^2 is the closed unit disk and H^2 is the open upper half-plane. (The latter set is the cone over $\mathbb{R}$; here x corresponds to 0.)

If x is not a vertex, there is only one local branch at x, namely the closed star of x itself. This unique local branch is PL homeomorphic to a closed disk, with x an interior point.

In any case, the closed star of a point x is the union of all local branches at x, and any two local branches intersect precisely at x.

5.3. Now suppose that the singular surface Σ is 1-connected. We define a *global branch* of Σ to be a subcomplex which is the closure of a connected component of $\Sigma - \mathrm{sk}_0(\Sigma)$. (Thus Σ is the union of its global branches, unless it consists of a single vertex.) It follows from the 1-connectedness of Σ that the global branches are 1-connected, and that any two global branches meet in a single point. The 1-connectedness of Σ also implies that for any vertex x of Σ and any global branch $\mathscr{B}$ containing x, the link of x in $\mathscr{B}$ is connected. It follows that the link of x in $\mathscr{B}$ is a connected component of the link of x in Σ, and therefore that the star of x in $\mathscr{B}$ is a local branch of x in Σ.

Hence every global branch of a 1-connected singular surface is either a closed 1-simplex or a 1-connected surface with points at infinity.

5.4. For future reference we insert a few remarks here about groups acting on surfaces with points at infinity.

Let Σ be a 1-connected triangulated surface with points at infinity and suppose that a group Γ admits an effective simplicial action on Σ such that Σ is finite mod Γ. Then Γ is clearly finitely generated. The interior $\overset{\circ}{\Sigma}$ of Σ is clearly Γ-invariant. It is easy to show that the action of Γ on $\overset{\circ}{\Sigma}$ must be properly discontinuous, in the sense that each point of $\overset{\circ}{\Sigma}$ has a neighborhood V such that $g\dot{V} \cap V = \varnothing$ for all but finitely many $g \in \Gamma$.

If x is a vertex belonging to $\overset{\circ}{\Sigma}$, so that the link of x is a triangulated 1-sphere L, then the stabilizer Γ_x acts faithfully and simplicially on L, and is therefore a finite cyclic or finite dihedral group. Similarly, if x is a vertex in E, then Γ_x acts faithfully and simplicially on a triangulated real line L; and L is finite $\mathrm{mod}\,\Gamma_x$ since Σ is finite $\mathrm{mod}\,\Gamma$. Thus Γ_x is an infinite cyclic or infinite dihedral group.

By a *planar discontinuous group* we shall mean a pair $(\Gamma, \mathscr{P})$, where Γ is a group, $\mathscr{P}$ is a (possibly empty) family of cyclic and dihedral subgroups of Γ which is closed under conjugation, and Γ admits an effective simplicial action on some 1-connected surface with points at infinity such that $\mathscr{P}$ is the set of all subgroups of Γ that are stabilizers of simplices. We shall often omit $\mathscr{P}$ from the notation and say that Γ is a planar discontinuous group. The finite subgroups in $\mathscr{P}$ are called *elliptic subgroups* of Γ; the infinite subgroups in $\mathscr{P}$ are *parabolic subgroups*.

For a classification of planar discontinuous groups, see Chapter 4 of [17].

We now turn to the discussion of measured foliations on singular surfaces. Our definitions are straightforward generalizations of those given in [5] for the case of (nonsingular) surfaces except that we work in the piecewise-linear (PL) category.

5.5. Let X be a topological space. By a *decomposition* of X we mean a collection of nonempty, pairwise disjoint subsets whose union is X. If $\mathscr{L}$ is a decomposition of X, every subset Z of X has an *induced decomposition* consisting of the path components of intersection of Z with the sets in $\mathscr{L}$.

Let Z_1 and Z_2 be polyhedra in simplicial complexes X_1 and X_2; thus Z_i is a subcomplex in some subdivision of X_i. Let $\mathscr{L}_i$ be a decomposition of Z_i. We shall say that $\mathscr{L}_1$ and $\mathscr{L}_2$ are *piecewise-linearly homeomorphic* if there is a PL homeomorphism of Z_1 onto Z_2 carrying the elements of $\mathscr{L}_1$ onto those of $\mathscr{L}_2$.

Let Z be a polyhedron in a simplicial complex. A decomposition of Z will be called a *product decomposition* if it is PL homeomorphic to the decomposition of some product complex $A \times B$ by the sets $A \times \{b\}$ for $b \in B$. The spaces A and B, which are determined up to homeomorphism by the decomposition, will be called the *divisor* and *quotient* of the product decomposition. The projection $q = q_Z : Z \to B$ is determined modulo PL homeomorphisms of B.

5.6. Let R be a connected combinatorial 1-manifold and let C denote the cone over R. Let K be a discrete subset of R such that every component of $R - K$ is relatively compact. Suppose that K has cardinality $k \geq 2$; note that k is finite if $R \approx S^1$ and $k = \aleph_0$ if $R \approx \mathbb{R}$. The cone over K is a subcomplex P of C. If D is the closure of any component of $C - P$ then D is PL homeomorphic to a closed disk, and $D \cap P$ is a closed arc; thus D has a product decomposition $\mathscr{L}_D$ in which $D \cap P$ is a single element, and the quotient is an arc. There is a unique decomposition $\mathscr{L}$ of C which induces the decomposition $\mathscr{L}_D$ on D for each component D of $C - P$. This decomposition $\mathscr{L} = \mathscr{L}_k$ is uniquely determined by k up to PL homeomorphism.

Let x be a point of a polyhedron Z. A decomposition of Z will be called a *standard k-prong decomposition centered at* x if it is PL homeomorphic to $\mathscr{L}_k$, under a PL homeomorphism that maps x onto the cone point in $\mathscr{L}_k$ (c.f. figure, next page).

Note that a standard 2-prong decomposition is simply a product decomposition in which the divisor and quotient are arcs.

A decomposition of Z will be called a *standard decomposition centered at* x if either (i) it is standard k-prong decomposition centered at x for

$$k = 3 \qquad\qquad\qquad k = \aleph_0$$

some $k \geq 2$, or (ii) Z is a closed arc and the elements of the decomposition are points.

5.7. Let Σ be a singular surface. A decomposition $\mathscr{F}$ of Σ will be called a *foliation with nondegenerate singularities*, or more briefly a *foliation*, if for every point $x \in \beta$ and every local branch β at x, there is a neighborhood Z of x in β whose induced decomposition is a standard decomposition centered at x. We call Z a *standard neighborhood*. (Our standard neighborhoods generalize the *domaines de cartes* of [5, Exposé 5].) If P denotes the element of the induced decomposition of Z containing x, the closures of the components of $Z - P$ are called *sectors* in the standard neighborhood Z of x. The induced decomposition of each sector is a product decomposition. The components of $P - \{x\}$ are called *prongs*. (Thus if $\dim \beta = 1$, there are no prongs.)

A point of Σ will be called *regular* (with respect to $\mathscr{F}$) if it has a neighborhood in Σ whose induced decomposition is a product decomposition; otherwise it will be called *singular*. The singular points form a discrete closed subset of Σ.

5.8. If $\mathscr{F}$ is a foliation of a singular surface Σ, an element of the decomposition $\mathscr{F}$ will be called a *leaf*. On the other hand, an element of the induced decomposition of the set of regular points of Σ will be called a *classical leaf*. In the case where Σ is a surface, our notion of classical leaf corresponds to the notion of "leaf" used in [5].

5.9. Let $\mathscr{F}$ be a foliation of a singular surface Σ. A path γ in Σ is said to be *monotonic (with respect to $\mathscr{F}$)* if $|\gamma|$ is contained in a subset Z of Σ whose induced decomposition is a product decomposition $Z = A \times B$ and $q_Z \circ \gamma$ is a monotonic map of $[0, 1]$ into B (in the sense that the inverse image of every point in B is a connected subset of $[0, 1]$). If $q_Z \circ \gamma$ is one-to-one we shall say that γ is *strictly monotonic*.

A path γ in Σ is *piecewise monotone* if some reparametrization of γ is a composition of finitely many monotonic paths. An arc $A \subset \Sigma$ is *piecewise monotonic* if a parametrization of A is piecewise-monotonic.

5.10. Now let D be a compact subset of Σ whose induced decomposition is a product decomposition with quotient B, and let η be a PL homeomorphism of B onto a closed interval in $\mathbb{R}$. Then there is a non-negative real-valued function μ_η on the set of all monotonic paths in D which assigns to each monotonic path γ the Euclidean length of the interval $\eta(q_D(|\gamma|))$. A function on the set of monotonic paths in D will be called a *standard transverse measure* in D if it has the form μ_η for some PL homeomorphism η.

Suppose now that $\mathscr{F}$ is a foliation (with nondegenerate singularities) of the singular surface Σ and that μ is a function that assigns a nonnegative number $\mu(\gamma) \in \Lambda$ to each path which is piecewise-monotonic with respect to $\mathscr{F}$. The pair $(\mathscr{F}, \mu)$ is called a *measured foliation* of Σ if the following conditions hold:

(i) μ is additive under composition of piecewise-monotonic paths; and

(ii) for every point $x \in \Sigma$ and every local branch β at x, there exist a standard neighborhood Z of x in β such that for each sector D of Z, the restriction of μ to the set of monotonic paths in D is a standard transverse measure.

Let $(\mathscr{F}, \mu)$ be a measured foliation on Σ. Then a path γ in Σ is contained in a leaf of $\mathscr{F}$ if and only if $\mu(\gamma)$ is defined and equal to 0. For this reason we shall generally suppress $\mathscr{F}$ from the notation and denote the measured foliation by μ.

It is clear that the measure of a piecewise-monotonic path with respect to a measured foliation is invariant under reparametrization. In particular, a piecewise-monotonic arc A has a well-defined measure $\mu(A)$.

5.11. It is often convenient to define a measured foliation on a singular surface Σ by local data. Suppose that for every point x of Σ and every local branch β at x we are given a neighborhood U_β of x in β, a standard decomposition $\mathscr{F}_\beta$ of U_β, and a standard measure μ_D on each sector D of U_β. Suppose that for any two points x, x', any two local branches β and β' at x and x' and any two sectors D and D' in U_β and $U_{\beta'}$, the following compatibility conditions hold:

(i) $\mathscr{F}_\beta$ and $\mathscr{F}_{\beta'}$ induce the same decomposition of $D \cap D'$, and

(ii) for any monotonic path γ in $D \cap D'$ we have $\mu_D(\gamma) = \mu_{D'}(\gamma)$.

Then there is a unique measured foliation of Σ which induces the decompositions $\mathscr{F}_\beta$ and the measures μ_D for all β and D.

5.12.   By a *measured foliated singular surface* we shall mean a pair (Σ, μ), where Σ is a singular surface and μ is a measured foliation. In particular we have the notion of a *measured foliated surface*, and of a *measured foliated surface with points at infinity*.

In arguments involving measured foliated surfaces we shall often quote results from [5]. This is legitimate because the arguments given in [5] for the smooth category are equally easy to interpret in the PL category. (Alternatively, one could show that every PL measured foliated surface can be smoothed.)

5.13.   Let (Σ, μ) be a measured foliated singular surface. Two piecewise-monotonic paths γ and γ' in Σ will be said to be *equivalent* if (i) $\mu(\gamma) = \mu(\gamma')$ and (ii) for any leaf L of (Σ, μ) we have $L \cap |\gamma| \neq \varnothing$ if and only if $L \cap |\gamma'| \neq \varnothing$.

5.14.   Let (Σ, μ) be a measured foliated singular surface, and let $\gamma: [0, 1] \to \Sigma$ be a path. We shall say that γ is *taut* if $\gamma^{-1}(L)$ is a connected subset of $[0, 1]$ for every leaf L of (Σ, μ). It is easy to see that any taut path is piecewise monotonic.

5.15.   Our next result, Lemma 5.16, gives a local characterization of taut paths up to equivalence; it depends on the following definition, which generalizes the one given in [5, Exposé 5].

Let γ be a path in the measured foliated singular surface (Σ, μ). We shall say that γ is *quasi-transverse* if it has a reparametrization $\gamma_1 * \cdots * \gamma_n$, where each γ_i is either a strictly monotonic path or a parametrized arc in a leaf, and for each $i < n$ there exists an $\varepsilon > 0$ such that one of the following alternatives holds:

   (i) $\gamma_i([1 - \varepsilon, 1])$ and $\gamma_{i+1}([0, \varepsilon])$ lie in distinct local branches at $\gamma_i(1)$; or
   (ii) there is a standard neighborhood Z of $\gamma_i(1)$ in a two-dimensional local branch of Σ such that
       (a) $\gamma_i([1 - \varepsilon, 1])$ and $\gamma_{i+1}([0, \varepsilon])$ are interior to different sectors of Z or
       (b) one of $\gamma_i([1 - \varepsilon, 1])$ and $\gamma_{i+1}([0, \varepsilon])$ lies in a prong of Z and the other is interior to a sector not containing that prong.

It is obvious that any quasi-transverse arc is piecewise monotonic.

5.16.   Lemma.   *A piecewise-monotonic path in a 1-connected measured foliated singular surface (Σ, μ) is taut if and only if it is equivalent to a quasi-transverse arc.*

Proof.   Let us say that a piecewise-monotonic path is in *standard position* if it is PL and has the form $\gamma = \gamma_1 * \cdots * \gamma_n$, where each γ_i is

either a strictly monotonic arc or an arc whose interior is contained in a classical leaf. The integer n will be called the *complexity* of γ. Given any piecewise-monotonic path, we can obtain an equivalent path in standard position by approximation and reparametrization. On the other hand, it follows from the definitions that a taut path which is in standard position, and has minimal complexity among all paths in its equivalence class, is a quasi-transverse arc. Hence any taut path is equivalent to a quasi-transverse arc.

To prove the converse, it suffices to show that every quasi-transverse arc γ is taut. If γ is not taut then there exist a subarc β of γ, a leaf L, and an arc $\alpha \subset L$ having the same endpoints as β, such that $|\gamma|\cap|\alpha| = \beta(\{0, 1\})$. Since Σ is a 1-connected singular surface, the PL simple closed curve $|\alpha| \cup |\beta|$ bounds a disk $D \subset \Sigma$.

Let E be the disk obtained by doubling D along $|\beta|$. (Thus E is the union of two copies of D, with the two copies of $|\beta|$ identified.) Then E has a natural decomposition: it is the unique decomposition that induces on each copy of D the decomposition that is induced by the foliation of Σ. Likewise, the 2-sphere S obtained by doubling E along its boundary has a natural decomposition. Since $|\alpha|$ is contained in a leaf and β is quasi-transverse, the natural decomposition of F is a foliation with nondegenerate singularities. But it follows from [5, Exposé 5, §I.6] that no such foliation exists on a 2-sphere. (This argument is a slight generalization of the one given on p. 77 of [5], where $|\alpha|$ is assumed to be contained in a *classical* leaf.)

5.17. Lemma. *In a* 1*-connected measured foliated singular surface, every quasi-transverse path is an arc.*

Proof. Let γ be a quasi-transverse path in (Σ, μ). It is clear from the definition that the map $\gamma: [0, 1] \to \Sigma$ is locally one-to-one. Hence if γ is not an arc there is a subpath γ' of γ which is a simple loop; that is, we have $\gamma'(0) = \gamma'(1)$, but $\gamma|(0, 1)$ is one-to-one. Now we argue as in the proof of 5.16: the 1-sphere $|\gamma|$ must bound a disk $D \subset \Sigma$; doubling D along its boundary we obtain a 2-sphere whose natural decomposition is a foliation, and again we have a contradiction to [5, Exposé 5, §I.6]

5.18. Proposition. *Let L_0 and L_1 be leaves in a* 1*-connected measured foliated singular surface* (Σ, μ). *Let γ and γ' be taut paths with $\gamma(i)$, $\gamma'(i) \in L_i$ for $i = 0, 1$. Then γ and γ' are equivalent.*

Proof. According to 5.16 we may assume that γ and γ' are quasi-transverse arcs. Let us first consider the case in which $|\gamma| \cap |\gamma'| = \varnothing$. For $i = 0, 1$, let α_i be an arc in L_i joining $\gamma(i)$ to $\gamma'(i)$. The PL 1-sphere $|\alpha_1| \cup |\gamma| \cup |\alpha_2| \cup |\gamma'|$ must bound a disk $R \subset \Sigma$; doubling R along $|\gamma| \cup |\gamma'|$

gives an annulus A, and doubling A along its boundary gives a torus T. The natural decomposition of T is a foliation with nondegenerate singularities. By [5, Exposé 5, §I.6], T has no singularities. Hence the natural decomposition of A is a foliation in which the boundary curves are leaves; according to [5, Exposé 6, §I.1], this is a product decomposition. Therefore the induced decomposition of R is also a product decomposition; it follows at once that γ and γ' are equivalent in this case.

Now let us replace the assumption $|\gamma| \cap |\gamma'| = \varnothing$ by the weaker assumption that $|\gamma|$ and $|\gamma'|$ have disjoint interiors. For $i = 0$, 1, the endpoints $\gamma(i)$ and $\gamma'(i)$ have neighborhoods in $|\gamma|$ and $|\gamma'|$ which are respectively contained in sectors D_i and D_i' in local branches of Σ. If $D_0 \neq D_0'$ then $\overline{\gamma} * \gamma'$ is a quasi-transverse path; by 5.16 and 5.17 it is taut. Since $\overline{\gamma} * \gamma'$ has its endpoints in the leaf L_1, it must be entirely supported in L_1; hence γ and γ' are supported in L_1 and are therefore equivalent. Thus we may assume that $D_0 = D_0'$; similarly we may assume that $D_1 = D_1'$. This allows us to approximate γ and γ' by equivalent quasi-transverse arcs which are disjoint from each other, thus reducing the argument in this case to the case already proved.

Finally we turn to the general case where γ and γ' are arbitrary quasi-transverse arcs. We may assume that $|\gamma|$ and $|\gamma'|$ meet in finitely many subarcs; hence after reparametrization we may write $\gamma = \gamma_1 * \cdots * \gamma_n$ and $\gamma' = \gamma_1' * \cdots * \gamma_n'$, where for each i the arcs $|\gamma_i|$ and $|\gamma'|$ either have disjoint interiors or are equal. By the case already proved, γ_i and γ_i' are equivalent. Hence γ and γ' are equivalent.

5.19. A measured foliated singular surface (Σ, μ) will be called *complete* if for every piecewise-monotonic path γ in Σ there is a taut path γ_0 such that $\mu(\gamma_0) \leq \mu(\gamma)$ and $\gamma_0(i) = \gamma(i)$ for $i = 0$, 1. (By 5.16, we may take γ_0 to be a quasi-transverse PL arc.)

Now consider a measured foliated singular surface (Σ, μ) which is 1-connected and complete. Let L and L' be leaves of (Σ, μ). By completeness, there exists a taut path joining a point of L to a point of L'. By 5.18 any two such paths have the same length. We define the *distance* between the leaves L and L' to be the length of a taut path joining a point of L to a point of L'.

It follows from 5.18 and the definition of completeness that the distance between L and L' is the minimum length of any piecewise-monotonic path joining a point of L to a point of L'.

It is therefore clear that with this definition of distance the set $T = T_\mu$ of leaves of (Σ, μ) is a metric space. We call T the *leaf space* of (Σ, μ).

There is a natural continuous map $\chi = \chi_\mu \colon \Sigma \to T$ which maps each point to the leaf containing it.

5.20. Theorem. *Let (Σ, μ) be a measured foliated singular surface which is complete and 1-connected. Then $T = T_\mu$ is an $\mathbb{R}$-tree. If γ is a taut path in Σ with endpoints u and v, then $\chi_\mu \circ \gamma$ is a (weakly) monotone map of $[0, 1]$ onto the segment $[\chi_\mu(u), \chi_\mu(v)]_T$. Furthermore, if V is any nonempty subset of Σ containing all the singular points of μ, then $T_0 = \chi_\mu(V)$ is an $\mathbb{R}$-pretree.*

Proof. Set $\chi = \chi_\mu$. Let $\mathfrak{G}$ denote the set of all subsets of T of the form $\chi(|\gamma|)$ where γ is a taut path in Σ. We shall show that every set in $\mathfrak{G}$ is a segment and that $\mathfrak{G}$ satisfies conditions (1), $(2_\mathbb{R})$ and (3) of 1.2. This will prove that T is an $\mathbb{R}$-tree.

Let γ be any taut path in Σ. We define a map $f\colon [0, 1] \to \mathbb{R}$ by $f(t) = l(\gamma|[0, t])$. The definition of the distance function on T shows that $|f(s) - f(t)| = d(\chi(s), \chi(t))$ for any $s, t \in [0, 1]$. Hence there is a well-defined isometric embedding $\overline{f}$ of $\chi(|\gamma|)$ into $\mathbb{R}$ given by $\overline{f}(\chi(\gamma(t))) = f(t)$. Thus $\chi(|\gamma|)$ is a segment. This shows that the sets in $\mathfrak{G}$ are segments.

To check condition (1) of 1.2, we consider two points $x, y \in T$. By completeness there is a segment in $\mathfrak{G}$ having x and y as endpoints. It follows from 5.18 that such a segment is unique. This establishes condition (1).

Condition $(2_\mathbb{R})$ follows from the observation that a subpath of a taut path is taut.

To prove (3) we consider three points $x, y, z \in T$ such that $S(x, y) \cap S(y, z) = \{y\}$. Thus if we write $x = \chi(u)$, $y = \chi(v)$ and $z = \chi(w)$, and if γ and ζ are taut paths joining u to v and v to w respectively, then the only leaf meeting both $|\gamma|$ and $|\zeta|$ is the leaf containing v. It therefore follows directly from the definition of tautness (5.14) that $\gamma * \zeta$ is taut. Hence $S(x, z) = \chi(|\gamma * \zeta|) = S(x, y) \cup S(y, z)$, and (iii) is established.

This completes the proof that T is an $\mathbb{R}$-tree. It also follows from this proof that if γ is a taut path in Σ with endpoints u and v, then $\chi_\mu \circ \gamma$ is a monotonic map of $[0, 1]$ onto $[\chi_\mu(u), \chi_\mu(v)]_T$.

To prove the last assertion of the proposition, we must show that if x_0, x_1, x_2 and y are points of T such that $[x_0, x_1] \cap [x_0, x_2] = [x_0, y]$, and if the x_i belong to $\chi_\mu(V)$, then $y \in \chi_\mu(V)$. Set $x_i = \chi_\mu(v_i)$, where $v_i \in V$. By completeness, there are quasi-transverse arcs γ_i joining v_0 to v_i for $i = 1, 2$. We have $\chi_\mu(|\gamma_i|) = [x_0, x_i]$; hence $\chi_\mu(\gamma_i(t_i)) = y$ for some $t_i \in [0, 1]$.

Assume that $y \notin \chi_\mu(V)$. Then we have $0 < t_i < 1$ for $i = 1, 2$, and the leaf containing the $\gamma_i(t_i)$ contains no singular points of the foliation. Hence if α is an arc in this leaf having the $\gamma_i(t_i)$ as endpoints, α has a neighborhood Z in Σ whose induced decomposition is a product decomposition. Hence for some $\varepsilon > 0$, the $\gamma_i|(t_i - \varepsilon, t_i + \varepsilon)$ are strictly monotone in Z. It follows that y is an interior point of $[x_0, x_1] \cap [x_0, x_2]$; this is a contradiction.   q.e.d.

5.21.   A foliation $\mathscr{F}$ of a singular surface Σ will be called *tame* if no simplex of Σ is contained in a leaf, and if for every 2-simplex δ of Σ, the decomposition of int δ induced by $\mathscr{F}$ is a product decomposition in which the elements form a parallel family of open line segments in the affine structure of δ. In particular, if $\mathscr{F}$ is tame then every singular point of μ is a vertex of Σ. A leaf of a tame foliated singular surface will be called 0-*tame* if it contains at least one vertex.

Now let Λ be a subgroup of $\mathbb{R}$. By a Λ-*foliated singular surface* we shall mean a measured foliated singular surface (Σ, μ) such that the underlying foliation of μ is tame, and such that for any taut path γ in Σ whose endpoints are vertices, we have $\mu(\gamma) \in \Lambda$. Note that the notion of an $\mathbb{R}$-foliated singular surface is slightly stronger than the notion of a measured foliated singular surface, since tameness is built into the definition.

5.22.   If (Σ, μ) is a Λ-foliated singular surface which is 1-connected and complete, then it follows from 5.20 and 5.21 that $T_0 = \chi_\mu(\mathrm{sk}_0(\Sigma))$ is a Λ-pretree. We shall call T_0 the *preleaf space* of (Σ, μ).

The tree T_μ is an $\mathbb{R}$-completion of T_0. Indeed, it follows from the definition of a tame foliation that every point lies on a monotone arc which is contained in a simplex and has its endpoints at vertices. Hence by 5.16 and 5.20, every point of T_μ lies in a segment whose endpoints are in T_0.

5.23.   Suppose that (Σ_0, μ) and (Σ_1, μ_1) are Λ-foliated singular surfaces. By a *morphism* from (Σ_0, μ_0) to (Σ_1, μ_1) we will mean a 0-tame map $W: \Sigma_0 \to \Sigma_1$ with the property that every taut path γ in Σ_0 may be reparametrized as a composition $\gamma_1 * \cdots * \gamma_n$, where

(i) $\gamma_i(1)$ lies in a 0-tame leaf for $i = 1, \cdots, n - 1$, and
(ii) each path $W \circ \gamma_i$ $(1 \leq i \leq n)$ is taut in Σ_1 and has the same length as γ_i.

This defines the *category of Λ-foliated singular surfaces*.

If (Σ, μ) is any Λ-foliated singular surface, we denote by $\mathrm{Aut}(\Sigma, \mu)$ the group of automorphisms of (Σ, μ) in the category of Λ-foliated singular surfaces. An element of $\mathrm{Aut}(\Sigma, \mu)$ is a 0-tame self-homeomorphism

of Σ which maps each piecewise-monotonic path in Σ to a piecewise-monotonic path of the same length. In particular an automorphism maps leaves to leaves. We denote by $\text{Aut}^S(\Sigma, \mu)$ the subgroup of $\text{Aut}(\Sigma, \mu)$ consisting of those automorphisms which are simplicial self-homeomorphisms of Σ.

By 1.15, we have a *category of Λ-foliated singular surfaces with symmetry*. An object in the latter category is a quadruple $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ consisting of a Λ-foliated complex and an action of a group by automorphisms. If Γ acts by simplicial automorphisms, i.e., if $\rho(\Gamma) \subset \text{Aut}^S(X, l)$, we call $\mathscr{S}$ a Λ-*foliated singular surface with simplicial symmetry*. We regard such objects as forming a full subcategory of the category of measured foliated complexes with symmetry (cf. 3.3).

5.24. Proposition. *Let (Σ_0, μ_0) and (Σ_1, μ_1) be Λ-foliated singular surfaces which are complete and 1-connected. Let T_i denote the preleaf space of (Σ_i, L_i), and set $\chi_i = \chi_{\mu_i} : \Sigma_i \to T_i$. Let $W : \Sigma_0 \to \Sigma_1$ be a 0-tame map. Then W is a morphism from (Σ_0, μ_0) to (Σ_1, μ_1) if and only if there is a morphism of Λ-pretrees $f : T_0 \to T_1$ such that the diagram*

$$
\begin{array}{ccc}
\Sigma_0 & \xrightarrow{\;W\;} & \Sigma_1 \\
\chi_0 \downarrow & & \downarrow \chi_1 \\
\mathbb{R}T_0 & \xrightarrow{\;\mathbb{R}f\;} & \mathbb{R}T_1
\end{array}
$$

commutes. Furthermore, if W is a morphism then f is uniquely determined by W.

Proof. First suppose that W is a morphism. If γ is a path in a leaf of Σ_0 then the definition of morphism implies that $W \circ \gamma$ is a piecewise-monotonic path of length 0 in Σ_1; thus W maps each leaf of Σ_0 into a leaf of Σ_1. Hence there is a (unique) map of sets $\overline{f} : \mathbb{R}T_0 \to \mathbb{R}T_1$ such that the diagram

$$
\begin{array}{ccc}
\Sigma_0 & \xrightarrow{\;W\;} & \Sigma_1 \\
\chi_0 \downarrow & & \chi_1 \downarrow \\
\mathbb{R}T_0 & \xrightarrow{\;\overline{f}\;} & \mathbb{R}T_1
\end{array}
$$

commutes. Since the morphism W is by definition 0-tame, it maps $\text{sk}_0(\Sigma_0)$ into $\text{sk}_0(\Sigma_1)$; hence $\overline{f}$ restricts to a map $f : T_0 \to T_1$.

To prove that f is a morphism of Λ-pretrees, we consider any non-degenerate presegment $[x, y] \subset T_0$. We may write $x = \chi_0(v)$ and $y = \chi_0(w)$ for some vertices v and w of Σ_0. By completeness, there is a taut path γ joining v to w. We may take γ to have the form $\gamma_1 * \cdots * \gamma_n$, where the γ_i satisfy conditions (i) and (ii) of 5.23. The segment $[x, y] = \chi_0(|\gamma|)$ is the union of the subsegments $\sigma_i = \chi_0(|\gamma_i|)$. If ζ is any subpath of $|\gamma_i|$ then $W \circ \zeta$ is taut, and since W is a morphism, $W \circ \zeta$ has the same length as ζ. Hence $f|\sigma_i$ is an isometry for each i. This shows that f is a morphism.

A similar argument shows that $\overline{f}$ is a morphism of $\mathbb{R}$-trees. Since $\overline{f}$ extends f it follows that $\overline{f} = \mathbb{R}f$. The proves the existence of a morphism f making the diagram in the statement of the proposition commute. The uniqueness of f is clear.

Conversely, suppose that there is a morphism f such that the diagram in the statement of the proposition commutes. Let γ be any taut path in Σ_0. Then $\chi_0 \circ \gamma$ is a monotonic map of $[0, 1]$ onto a segment $[x, y]_{\mathbb{R}T_0}$. Since f is a morphism of Λ-pretrees, we may write $[x, y]_{\mathbb{R}T_0}$ as a union of subsegments $[x_{i-1}, x_i]$ for $i = 1, \cdots, n$, where $x_0 = x$, $x_n = y$, $x_i \in T_0$ for $0 < i < n$, and $f|[x_{i-1}, x_i]$ is an isometric embedding for $1 \leq i \leq n$. Furthermore we may choose the x_i so that $[x_{i-1}, x_i] \cap [x_i, x_{i+1}] = \{x_i\}$ for $0 < i < n$. Let us reparametrize γ in the form $\gamma_1 * \cdots * \gamma_n$ so that $\chi_0(\gamma_i(1)) = x_i$ for $i = 1, \cdots, n$. Then the γ_i are at once seen to satisfy conditions (i) and (ii) of 5.23. Hence W is a morphism.   q.e.d.

The morphism f given by the above proposition will be said to be *induced* by W.

5.25.   A Λ-foliated singular surface with simplicial symmetry, $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$, will be said to be *uniform* if Σ is finite mod Γ.

Proposition.   *Any uniform 1-connected $\mathbb{R}$-foliated singular surface with symmetry is complete.*

5.26.   The main step in proving 5.25 is the

Lemma.   *Let $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ be a uniform, 1-connected $\mathbb{R}$-foliated singular surface with points at infinity. Let x, y and z be points of Σ, and suppose that there exist piecewise-linear quasi-transverse arcs α and β joining z to x and to y respectively. Then there is a piecewise-linear quasi-transverse arc γ joining x to y, with $\mu(\gamma) \leq \mu(\alpha) + \mu(\beta)$.*

Proof of Lemma.   The proof is based on the results on [5]. The first step is to show that the measured foliated singular surface (Σ, μ) has a *good atlas*. We define the notion of a good atlas as in [5, Exposé 5], using standard neighborhoods in place of *domaines de cartes* (cf. 5.7).

It is observed on pp. 73–74 of [5] that every compact measured foliated surface has a good atlas; we shall produce a good atlas on (Σ, μ) by lifting one from an appropriate compact quotient surface.

We may assume that the action of Γ on Σ is effective. Thus we may regard Γ as a planar discontinuous group. It follows from Theorem 4.10.1 of [17] that Γ is residually finite and has a torsion-free subgroup Γ_0 of finite index. The subgroup Γ_0 acts freely on the interior $\overset{\circ}{\Sigma}$ (see 5.4). Set $E = \Sigma - \overset{\circ}{\Sigma}$.

We assert that Γ_0 has a finite-index subgroup Γ_1 such that for every $x \in E$ there are at least two orbits for the action of the stabilizer $\Gamma_x \cap \Gamma_1$ on the set of prongs in a standard $\aleph_0$-prong neighborhood of x. Indeed, let S be a finite complete set of orbit representatives for the action of Γ_0 on E. For each $x \in S$ we choose two distinct prongs P_x and P'_x in a standard neighborhood Z_x. The infinite cyclic or dihedral group Γ_x contains only finitely many elements that map P_x to P'_x. Hence by residual finiteness there is a normal subgroup Γ_1 of Γ_0 such that for any $x \in S$ and any $\gamma \in \Gamma_1 \cap \Gamma_x$ we have $\gamma \cdot P_x \neq P'_x$. It is clear that Γ_1 has the asserted property.

The first barycentric subdivision of Σ induces a triangulation of the quotient set Σ/Γ_1. Since the action of Γ_1 on Σ' is free, Σ/Γ_1 is a compact triangulated surface. The foliation of Σ induces a decomposition of Σ/Γ_1; we claim that this is a foliation of Σ/Γ_1. Indeed, any point of Σ' has the form $p(x)$, where $x \in \Sigma' \cup S$. If $x \in \Sigma'$ then the induced decomposition of a small neighborhood of $p(x)$ is PL homeomorphic to the induced decomposition of a small neighborhood of x. On the other hand, if $x \in S$, the defining property of Γ_1 implies that some neighborhood of $p(x)$ has a standard k-prong decomposition for some $k \geq 2$. Thus the decomposition of Σ/Γ_1 is indeed a foliation.

It is clear that the transverse measure μ induces a transverse measure $\overline{\mu}$ on the induced foliation of Σ/Γ_1. According to [5], $(\Sigma/\Gamma_1, \overline{\mu})$ has a good atlas. Clearly (Σ, μ) inherits one.

Now suppose that we have points x, y and z, and arcs α and β, satisfying the hypotheses of the lemma. If $\overline{\alpha} * \beta$ is quasi-transverse there is nothing to prove; thus we assume that it is not. Similarly we assume that $x \neq y$. After passing to subarcs of α and β, we may assume that $|\alpha| \cap |\beta| = \{z\}$.

We may also assume that z is a regular point of (Σ, μ). Indeed, since $\overline{\alpha} * \beta$ is not quasi-transverse, there are neighborhoods of z in $|\alpha|$ and $|\beta|$ which lie in the same sector D of a standard neighborhood of z in Σ. If

z_1 is an interior point of D sufficiently close to z, then there are quasi-transverse PL arcs α_1 and β_1 joining z_1 to x and to y respectively, and such that $\mu(\alpha_1) < \mu(\alpha)$ and $\mu(\beta_1) < \mu(\beta)$. Replacing z, α and β by z_1, α_1 and β_1, we reduce the argument to the case where z is regular.

Given that z is regular and that $|\alpha| \cap |\beta| = \{z\}$, we may perturb α and β slightly to obtain parametrized arcs α' and β' which are equivalent to α and β respectively, and which satisfy $|\alpha'| \cap |\beta'| = \varnothing$. In particular, $\alpha'(1)$ and $\beta'(1)$ are distinct but lie on the same leaf.

Let A denote the largest subarc of $|\alpha'|$ which has $\alpha'(1)$ as an endpoint and has only regular points of (Σ, μ) in its interior. Let α'' denote the unique PL homeomorphism of $[0, \mu(A)]$ onto A such that $\alpha''(0) = z$ and $\mu(\alpha''([0, t])) = t$ for every t. Let B and β'' be defined similarly. Then α'' and β'' are transverse to the foliation in the sense of [5]. Since (Σ, μ) has a good atlas, the proof of the "stability lemma" of [5, p. 80], shows that there exist a number $t_0 \in [0, \min(\mu(A), \mu(B))]$ and an immersion $H: [0, 1] \times [0, t_0] \to \Sigma$ such that (i) $H(0, t) = \alpha''(t)$ and $H(1, t) = \beta''(t)$ for every $t \in [0, 1]$, (ii) $H([0, 1] \times \{t\})$ is contained in a leaf (cf. 5.8) of (Σ, μ) for every $t \in [0, 1]$, and (iii) either $t_0 = \min(\mu(\alpha''), \mu(\beta''))$, or the arc γ_{t_0} joining $\alpha''(t_0)$ and $\beta''(t_0)$ in L_0 contains a singular point. It is now easy to check that $\alpha''([t_0, 1]) \cup H([0, 1] \times \{t_0\}) \cup \beta''([t_0, 1])$ is the support of a PL quasi-transverse arc from x to y.

5.27. *Proof of 5.25.* Let $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ be a uniform, 1-connected $\mathbb{R}$-foliated singular surface with symmetry. We must show that if r is a real number, and if x and y are points of Σ which are joined by some piecewise-monotonic path of length $\leq r$, then there is a taut path of length $\leq r$ joining x to y.

Any piecewise-monotonic path joining x and y may be written in the form $\zeta = \zeta_1 * \cdots * \zeta_n$, where each ζ_i is a path in a global branch $\mathscr{B}_i$ of Σ. Among all such paths having length $\leq r$, suppose that ζ has been chosen so as to minimize the integer n. Then $\mathscr{B}_i \neq \mathscr{B}_{i+1}$ for $i = 1, \cdots, n-1$. By 5.3, each $\mathscr{B}_i$ is either an arc or a 1-connected surface with points at infinity; and the definition of tautness implies that if γ_i is a path in $\mathscr{B}_i$ which has the same endpoints as ζ_i and is taut with respect to $\mu|\mathscr{B}_i$, then $\gamma = \gamma_1 * \cdots * \gamma_n$ is taut with respect to μ. Hence, without loss of generality, we may assume that Σ is an arc or a surface with points at infinity. When Σ is an arc the result is trivial.

Now suppose that Σ is a surface with points at infinity. A piecewise-monotonic path joining x to y and having length $\leq r$ may be written in the form $\zeta = \eta_1 * \cdots * \eta_n$, where each η_i is taut. We argue by induction on n to show that x and y may be joined by a taut path of length $\leq r$.

For $n = 1$ this is trivial, and the induction step follows immediately from 5.26.   q.e.d.

5.28.   Suppose that $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ is a uniform, 1-connected Λ-foliated singular surface with symmetry. Then according to 5.25 and 5.19–5.20, the leaf space T of (Σ, μ) is defined and is an $\mathbb{R}$-tree. It is also clear that Γ admits a natural action (by isometries) on T. This gives an $\mathbb{R}$-tree with symmetry $\mathscr{T} = (T, \Gamma, \overline{p})$ which we call the *leaf space* of $\mathscr{S}$. The map $\chi_\mu: \Sigma \to T$ is Γ-equivariant. The preleaf space $T_0 \subset T$ is clearly Γ-invariant; hence we have a Λ-pretree with symmetry $\mathscr{T}_0 = (T_0, \Gamma, \overline{p}_0)$. We call it the *preleaf space* of $\mathscr{S}$.

5.29.   Suppose for $i = 0, 1$ that $\mathscr{S}_i$ is a Λ-foliated singular surface with simplicial symmetry which is uniform and 1-connected. Let $\mathscr{T}_i$ denote the preleaf space of $\mathscr{S}_i$. It follows from 5.24 that any morphism $\mathscr{W} = (W, \omega): \mathscr{S}_0 \to \mathscr{S}_1$ induces a morphism of pretrees with symmetry $\mathscr{f} = (f, \omega): \mathscr{T}_1 \to \mathscr{T}_2$.

5.30. Proposition.   *Let $\mathscr{T} = (T, \Gamma, \overline{p})$ be the leaf space of a uniform 1-connected Λ-foliated surface with simplicial symmetry. Then no point of T is fixed by Γ.*

Proof.   Let $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ be the given Λ-foliated surface with symmetry. Since $\mathscr{S}$ is uniform, Γ has a finite generating set S. By Theorem 4.10.1 of [5], the planar discontinuous group Γ has a torsion-free finite-index subgroup Γ_0, and Γ_0 acts freely on Σ. The quotient set $\overline{\Sigma} = \Sigma/\Gamma_0$ has the structure of a closed PL surface, and μ induces a measured foliation $\overline{\mu}$ on $\overline{\Sigma}$.

If Γ fixes a point of T then some leaf L in Σ is invariant under Γ. Let x be a point of L. Let $\gamma_1, \cdots, \gamma_n$ be generators for Γ, and let K be a compact connected subpolyhedron of L containing $\gamma_1 \cdot x, \cdots, \gamma_n \cdot x$. Then the image $\overline{K}$ of K in $\overline{\Sigma}$ carries $\pi_1(\overline{\Sigma})$. Hence for any component C of $\overline{\Sigma} - \overline{K}$, there is a PL map of the disk D^2 onto $\overline{C}$ which maps $\operatorname{int} D^2$ homeomorphically onto C. Since the boundary of C is contained in a leaf of $\overline{\Sigma}$, we can pull back the foliation of Σ to D^2 and use the doubling argument from the proof of 5.16 to obtain a contradiction.   q.e.d.

5.31.   Now let us consider a measured foliation $(\mathscr{F}, \mu)$ on a compact surface F. The universal cover $\widetilde{F}$ of F inherits a measured foliation $\widetilde{\mu}$; thus if ρ is the inclusion of the covering group $\Gamma \approx \pi_1(F)$ in $\operatorname{Aut}(\widetilde{\Sigma}, \mu)$, then $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ is a uniform, 1-connected $\mathbb{R}$-foliated singular surface with symmetry. Its leaf space is an $\mathbb{R}$-tree T on which Γ acts in a natural way; by 5.30, the action is nontrivial. We call T the *dual tree* of the measured foliation μ.

The stabilizer of any segment in T is a cyclic subgroup of Γ; thus the action of Γ on T is "small" in the sense of the introduction. (Indeed, the stabilizer of any segment is contained in the stabilizer of some nonsingular leaf of $\widetilde{F}$; the latter group is easily seen to be cyclic.)

Let $\tilde{v}$ be a base point in $\widetilde{F}$, and let x and v denote the images of $\tilde{v}$ in T and Σ respectively. For any $g \in \Gamma$, the definition of the distance on T implies that $d(x, g \cdot x) = \min \mu(\alpha)$, where α ranges over all based loops representing the element of $\pi_1(\Sigma, v)$ corresponding to g. Hence the length function determined by the action of Γ on T is given by $l(g) = \min \mu(\alpha)$, where α ranges over all closed curves in the free homotopy class σ corresponding to the conjugacy class of g. Thus in the terminology of [5], $l(g)$ is the geometric intersection number $I(\mathscr{F}, \mu; \sigma)$. We call l the *length function determined by* μ and denote it l_μ. (When μ is defined by the construction of [5, Exposé 5, §III] from a compact two-sided 1-manifold $C \subset F$ with no homotopically trivial components, we have $l_\mu = l_C$ in the notation of 2.10.)

5.32. We conclude this section with a few remarks about groups that act on measured foliated surfaces, or on measured foliated surfaces with points at infinity.

For most planar discontinuous groups Γ, it follows from the classification in [17] that Γ splits (2.7) over a cyclic or dihedral subgroup. The exceptional cases arise when Γ has a presentation of the form

(i) $\langle a, b: a^l = b^m = (ab)^n = 1 \rangle$ or
(ii) $\langle x, y, z: x^2 = y^2 = z^2 = (xy)^l = (yz)^m = 1 \rangle$.

(Groups of type (i) are *triangle groups* if l, m, $n > 1$; otherwise they are finite cyclic groups. Groups of type (ii) contain groups of type (i) as subgroups of index 2.)

On the other hand, a group of type (i) or (ii) does not split at all; this is a consequence of [15, Chapter I, Theorem 13] and the following result.

Proposition. *Let Γ be a group which admits a presentation of the form* (i) *or* (ii). *Let Λ be any ordered abelian group and let Γ act without inversions on a Λ-tree T. Then Γ has a fixed point in T.*

Proof. First suppose that Γ has a presentation of the form (i). Then it follows from Theorem II.2.3 of [12] that each of the elements a, b and ab has a fixed point in T. Hence by the proof of [12, Proposition II.2.15], there is a point of T that is fixed by all of Γ. Now suppose that Γ has a presentation of the form (ii). Then some subgroup Γ' of index 2 in Γ has a presentation of the form (i). Hence Γ' has a fixed point x in T. It follows that Γ has a fixed point as well. q.e.d.

We shall call a planar discontinuous group Γ *splittable* if it does not have a presentation of the form (i) or (ii).

5.33.   The following result can be proved by a direct geometric argument, but it is interesting as well as convenient to derive it from Proposition 5.32.

Proposition. *Let (Σ, μ) be a nondegenerate, 1-connected Λ-foliated singular surface with points at infinity. Let Γ be a group that acts on (Σ, μ) by automorphisms, and suppose that Σ is finite $\mathrm{mod}\,\Gamma$. Then Γ is a splittable planar discontinuous group.*

Proof.   By definition Γ is a planar discontinuous group.

Recall from 5.4 that $\overset{\circ}{\Sigma}$ is a 1-connected surface. If $\overset{\circ}{\Sigma} \neq \Sigma$, then Γ acts on $\overset{\circ}{\Sigma}$ with a noncompact fundamental domain; it follows from the classification given in [17] that Γ is splittable in this case.

Now suppose that $\overset{\circ}{\Sigma} = \Sigma$, so that Σ is a surface. By 5.30, Γ admits an action on an $\mathbb{R}$-tree T such that no point of T is fixed by Γ. Hence by 5.32, Γ is splittable.

6. The rank-2 case: statements of results

In this section we state the structure theorem (Theorem 6.2) for certain group actions on Λ-trees in the case where $\Lambda \subset \mathbb{R}$ has $\mathbb{Q}$-rank 2. This theorem was explained informally in the introduction. We also give formal statements of Theorems A, B, C and D of the introduction for the rank-2 case, a simplified version of the structure theorem in the finitely presented case (6.4), and a partial reinterpretation (6.5) of the structure theorem in terms of graphs of groups analogous to 2.6. In this section we show that the latter results follow from the structure theorem; the theorem itself will be proved in §11.

6.1.   As we indicated in the introduction, the structure theorem applies to a countable Λ-tree T and an action of a group Γ on T that satisfies the following condition:

Ascending Chain Condition for segment stabilizers.   *Let*

$$[x_0, y_0] \supset [x_1, y_1] \supset \cdots$$

be a monotone decreasing sequence of segments in T. Suppose that the segments all have the same midpoint in $\mathbb{R}T$, or equivalently that $d(x_i, x_{i+1}) = d(y_i, y_{i+1})$ for all i. Suppose also that $\lim_{i \to \infty} d(x_i, y_i) = 0$. Then there exists an $m \geq 0$ such that $\Gamma_{[x_i, y_i]} = \Gamma_{[x_m, y_m]}$ for all $i \geq m$.

6.2. Theorem. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2, and let $\mathscr{T} = (T, \Gamma, \rho)$ be a countable Λ-gree with symmetry. Suppose that the action of Γ on T satisfies the Ascending Chain Condition 6.1. Then $\mathscr{T}$ is the limit of a strongly convergent direct system $(\mathscr{T}_i : /_{ij})$ in the category of Λ-pretrees with symmetry, indexed by the positive integers, where each $\mathscr{T}_i$ is the preleaf space (5.28) of a 1-connected, uniform L_i-foliated singular surface with symmetry $\mathscr{S}_i$ for some rank-2 free abelian subgroup L_i of Λ. Furthermore, we have $L_i \subset L_j$ whenever $i \le j$; and the morphisms $/_{ij}$ are induced (5.25) by morphisms $\mathscr{W}_{ij} = (W_{ij}, \omega_{ij}): \mathscr{S}_i \to \mathscr{S}_j$ of Λ-foliated singular surfaces with symmetry. Finally, for each i, the kernel of the canonical homomorphism $\omega_i: \Gamma_i \to \Gamma$ acts freely on Σ_i.*

In the above statement it is understood (see 1.29) that the direct system $(\mathscr{T}_i ; /_{ij})$ is indexed by the natural numbers. The theorem would remain true if T were not assumed to be countable, except that the direct system would in general be indexed by some filtered preordered set. This generalization would require only minor changes in the proof.

For the remainder of this section, Theorem 6.2 will be assumed.

6.3.   The following result is the rank-2 case of Theorem A of the introduction.

Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2. Let Γ be a finitely generated group which acts freely, without inversions, on some Λ-tree T. Then Γ is a free product of infinite cyclic groups and surface groups.*

Furthermore, T is the limit of a strongly convergent system (T_i, f_{ij}) in the category of Λ-pretrees, where each T_i is the leaf space of a complete, 1-connected, Λ-foliated singular surface (Σ_i, μ_i); there exist free actions without inversions of Γ on the T_i such that f_{ij} and the canonical morphisms $f_i: T_i \to T$ are Γ-equivariant. The actions of Γ on the T_i are induced by free actions by simplicial automorphisms on the (Σ_i, μ_i), and the Σ_i are finite $\mathrm{mod}\,\Gamma$. Finally, each μ_i takes its values in a rank-2 free abelian subgroup L_i of Λ, so that T_i is an L_i-pretree.

Proof. Since Γ and Λ are countable, T has some countable, Γ-invariant subset which is a Λ-tree. Hence we may assume that T is countable.

Since Γ acts freely and without inversions on T, it acts freely on $\mathbb{R}T$ by 1.18. Also, since the action of Γ on T is free, the Ascending Chain Condition 6.1 trivially holds.

Let $(\mathscr{T}_i, /_{ij})$ be the direct system given by Theorem 6.2; set $\mathscr{T}_i = (T_i, \Gamma_i, \rho_i)$ and $/_{ij} = (f_{ij}, \omega_{ij})$. Let $\mathscr{S}_i = (\Sigma_i, \mu_i, \Gamma_i, \rho_i)$ be the singular surfaces with symmetry given by 6.2. For $i = 1, 2, \ldots$ there is

a canonical morphism $\mathscr{f}_i = (f_i, \omega_i): \mathscr{T}_i \to \mathscr{T}$ of Λ-pretrees with symmetry. By 1.17 there is a corresponding morphism $\mathbb{R}\mathscr{f}_i: \mathbb{R}\mathscr{T}_i \to \mathbb{R}\mathscr{T}$ of $\mathbb{R}$-trees with symmetry. Since $\chi_i = \chi_{\mu_i}: \Sigma_i \to \mathbb{R}T_i$ is Γ_i-equivariant, it follows that $\mathscr{k}_i = (\mathbb{R}f_i \circ \chi_i, \omega_i)$ is a morphism of sets with symmetry. Since the action of Γ on $\mathbb{R}T$ is free it follows that any element of Γ_i that fixes a point of Σ_i lies in the kernel of ω_i. But by 6.2, the kernel of ω_i acts freely on Σ_i. Therefore the action of Γ_i on Σ_i is free.

Any subgroup of Γ_i acts freely and simplicially on the singular surface Σ_i and is therefore isomorphic to the fundamental group of a singular surface. Hence any subgroup of Γ_i is a free product of infinite cyclic groups and surface groups.

It follows for example from Theorem 1 of [16] that a finitely generated free product of surface groups and infinite cyclic groups is isomorphic to a subgroup of $GL_2(\mathbb{C})$. Thus every finitely generated subgroup of Γ_i is isomorphic to a subgroup of $GL_2(\mathbb{C})$.

By the definition of the limit of a strongly convergent system, Γ is the direct limit of the direct system of groups (Γ_i, ω_{ij}). Hence we can apply Lemma 2.3 to conclude that Γ is isomorphic to a subgroup of one of the Γ_i and is therefore itself a free product of surface groups and infinite cyclic groups.

The last assertion of the proposition is proved in the same way as the last assertion of 2.2.

6.4. The following result is a simplified version of the structure theorem in the case where Γ is finitely presented.

Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2, and let Γ be a finitely presented group that acts on a Λ-tree T. Suppose that the action of Γ on T satisfies the Ascending Chain Condition 6.1. Then T is the limit of a strongly convergent system (T_i, f_{ij}) in the category of Λ-pretrees, where each T_i is the leaf space of a complete, 1-connected, Λ-foliated singular surface (Σ_i, μ_i); there exist actions of Γ on the T_i such that the f_{ij} and the canonical morphisms $f_i: T_i \to T$ are Γ-equivariant. The actions of Γ on the T_i are induced by actions by simplicial automorphisms on the (Σ_i, μ_i), and the Σ_i are finite mod Γ. Finally, each μ_i takes its values in a rank-2 free abelian subgroup Λ_i of Λ, so that T_i is a Λ_i-pretree.*

Proof. As in the proof of 6.3, we may assume that T is countable.

Let $(\mathscr{T}_i; \mathscr{f}_{ij})$ be the direct system given by Theorem 6.2. Let us write $\mathscr{T}_i = (T_i, \Gamma_i, \rho_i)$ and $\mathscr{f}_{ij} = (f_{ij}, \omega_{ij})$. By the definition of the limit of a strongly convergent direct system, Γ is the direct limit of the direct system of groups (Γ_i, ω_{ij}). Since Γ is finitely presented, we can apply 2.3

to show that for all sufficiently large i, say for $i \geq m$, there is a subgroup Γ_i^* of some Γ_i such that the natural homomorphism $\omega_i \colon \Gamma_i \to \Gamma$ maps Γ_i^* isomorphically onto Γ. As in the last step of the proof of 2.2 we can then conclude that there is an action of Γ on T_m such that f_m is Γ-equivariant. The result follows at once.   q.e.d.

6.5.   The next result is the analogue of 2.6 for the rank-2 case.

Proposition.   *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2, and let Γ be a finitely presented group that acts without inversions on a Λ-tree T. Suppose that the action of Γ on T satisfies the Ascending Chain Condition 6.1. Then Γ can be identified isomorphically with the fundamental group of a graph of groups $\mathscr{G}$ whose vertex set can be expressed as a disjoint union $V_1 \amalg V_2$ in such a way that the following conditions hold.*

(i)   *For each $v \in V_1$ the vertex group $\mathscr{G}(v)$ is contained in the stabilizer of a point in T.*

(ii)   *For each $v \in V_2$ there is an exact sequence*

$$1 \to N_v \to \mathscr{G}(v) \overset{p_v}{\to} Q_v \to 1,$$

where Q_v is a splittable planar discontinuous group, and N_v is a subgroup of Γ which fixes some nondegenerate segment in T.

(iii)   *For each edge e of $\mathscr{G}$ whose endpoints are both in V_1, the edge group $\mathscr{G}(e)$ fixes some nondegenerate segment in T.*

(iv)   *For each edge e of $\mathscr{G}$ that has at least one endpoint v in V_2, we have $\mathscr{G}(e) \subset p_v^{-1}(\mathscr{X})$, where $\mathscr{X} \subset Q_v$ is an elliptic or parabolic subgroup.*

6.6.   The proof of 6.5 uses the following general construction.

Let Σ be any 1-connected singular surface. We can associate to Σ an abstract simplicial complex K of dimension ≤ 1 in the following way. The vertex set of K is a disjoint union $S \amalg S'$, where S is a set in 1-1 correspondence with the set of global branches of X, and S' is in 1-1 correspondence with the set of all vertices of X which lie in at least two distinct global branches of X. A 1-simplex of K is a pair $\{v, w\}$, where $v \in S$, $w \in S'$, and the vertex of X corresponding to w lies in the global branch corresponding to v.

Since X is 1-connected it is easy to show that K is also 1-connected; that is, K is a simplicial tree.

6.7.   *Proof of Proposition 6.5.*   It follows from 6.4 that Γ acts by automorphisms on a simply connected Λ-foliated singular surface (Σ, μ) so that Σ is finite mod Γ, and that there is a Γ-equivariant morphism f of

Λ-pretrees from the preleaf space T_0 of (Σ, μ) to T. Set $\chi = \chi_\mu \colon \Sigma \to \mathbb{R}T_0$. Since f and χ are Γ-equivariant, so is $k = \mathbb{R}f \circ \chi \colon \Sigma \to \mathbb{R}T$. By definition we have $\chi(\mathrm{sk}_0(\Sigma)) = T_0$; hence $k(\mathrm{sk}_0(\Sigma)) \subset T$.

Let K be the simplicial tree associated to Σ be the construction of 6.6. In the notation of 6.6, the vertex set of K is a disjoint union $S \amalg S'$. An element of the set S corresponds by definition to a global branch $\mathscr{B}$ of Σ. By 5.3, each global branch is either an isolated 1-simplex or a surface with points at infinity. We write $S = S_1 \amalg S_2$, where every vertex in S_1 corresponds to an isolated 1-simplex and every vertex in S_2 corresponds to a surface with points at infinity.

The simplicial action of Γ on Σ induces a simplicial action on K. Moreover, the sets S and S' are invariant under Γ; in view of the definition of the 1-simplices of K, it follows that Γ acts on K without inversions in the sense of [15]. Hence by [15, Chapter I, Theorem 13], there is a natural isomorphic identification of Γ with the fundamental group of a graph of groups $\mathscr{G}$ whose underlying graph is $G = K/\Gamma$. For each cell c of G, the subgroup $\mathscr{G}(c)$ of Γ, which is defined up to conjugacy, is the stabilizer of a simplex of K which maps to c under the projection $p \colon K \to G$.

The vertex set of K is $S \amalg S' = (S_1 \amalg S') \amalg S_2$. Hence the vertex set of G is the disjoint union of $V_1 = p(S_1 \amalg S')$ and $V_2 = p(S_2)$. We claim that conditions (i)–(iv) of the statement of Proposition 6.5 hold with these choices of V_1 and V_2.

Given $v \in V_1$, we have either $v \in p(S')$ or $v \in p(S_1)$. If $v \in p(S')$ then $\mathscr{G}(v)$ is the stabilizer of a vertex $s \in S'$. By the definition of S' this means that $\mathscr{G}(v)$ is the stabilizer of a vertex x of Σ. Since $k \colon \Sigma \to \mathbb{R}T$ is Γ-equivariant, $G(v)$ is contained in the stabilizer of $k(x)$; but $k(x) \in T$ since $k(\Sigma^0) \subset T$. This proves (i) in the case $v \in p(S')$.

Now suppose that $v \in p(S_1)$. Then $\mathscr{G}(s)$ is the stabilizer of an isolated 1-simplex τ of Σ. Let x and y denote the endpoints of τ. Since k is Γ-equivariant, the set $\{k(x), k(y)\}$ is invariant under $\mathscr{G}(s)$. We have $k(x), k(y) \in T$ since $k(\Sigma^0) \subset T$, and the segment $[k(x), k(y)] \subset T$ must be invariant under $\mathscr{G}(s)$. Since Γ acts on T without inversions it follows that $G(s)$ has a fixed point in T. This completes the proof of (i).

To prove (ii) we must consider a vertex $v \in V_2 = p(S_2)$. The group $\mathscr{G}(v)$ is the stabilizer of a global branch $\mathscr{B}$ of Σ which is a surface with points at infinity. Since Σ is simply connected, so is $\mathscr{B}$. Clearly $\mathscr{B}$ is invariant under $\mathscr{G}(v)$. Let N_v denote the normal subgroup of $\mathscr{G}(v)$ consisting of elements that fix $\mathscr{B}$ (pointwise). Then $Q_v = \mathscr{G}(v)/N_v$ acts

effectively on $\mathscr{B}$. Hence we may regard Q_v as a planar discontinuous group (5.4); the elliptic and parabolic subgroups of Q_v are the stabilizers of simplices of $\mathscr{B}$. Since the Λ-foliation $\mu|\mathscr{B}$ is invariant under the action of Q_v on $\mathscr{B}$, it follows from 5.33 that Q_v is splittable.

By definition, N_v fixes $\mathscr{B}$ and hence fixes $\chi(\mathscr{B}) \subset \mathbb{R}T'$. Since $\chi(\mathscr{B})$ contains a nondegenerate segment, and since f is a morphism of Λ-pretrees, $f(\chi(\mathscr{B}))$ also contains a nondegenerate segment, which is again fixed by N_v. This proves (ii).

To prove (iii) we must consider an edge e of $\mathscr{G}$ whose endpoints are both in $V_1 = p(S' \amalg S_1)$. Thus $e = p(\tilde{e})$, where $\tilde{e}$ is an edge of K whose endpoints are in $S' \amalg S_1'$. The definition of the complex K implies that one endpoint of $\tilde{e}$, say $\tilde{v}$, must lie in S', and the other, say $\tilde{v}'$, must lie in S_1. Thus $\tilde{v}$ corresponds to a vertex w of Σ and $\tilde{v}'$ corresponds to an isolated edge τ of Σ. Furthermore, w must be an endpoint of τ. The edge group $\mathscr{G}(e)$ is conjugate to the subgroup of Γ consisting of all elements that fix w and leave τ invariant; this is the same as the subgroup consisting of all elements that fix τ pointwise. Since f is a morphism of Λ-pretrees, $\mathbb{R}f(\tau)$ contains a nondegenerate segment $[x, y] \subset T$. Since f is Γ-equivariant, $\mathscr{G}(e)$ is conjugate to a subgroup of Γ that fixes $[x, y]$. This proves (iii).

Similarly, if an edge e of $\mathscr{G}$ has at least one endpoint v in V_2, then $v = p(\tilde{v})$, where $\tilde{v}$ corresponds to a global branch $\mathscr{B}$ of Σ which is a singular surface. Furthermore, the other endpoint of e has the form $p(\tilde{u})$, where $\tilde{u}$ corresponds to a vertex w of $\mathscr{B}$; and after a conjugation we may take $\mathscr{G}(e)$ to be the subgroup of Γ consisting of all elements that fix w and leave $\mathscr{B}$ invariant. Hence in the notation of the proof of (ii) we have $\mathscr{G}(e) \subset p_v^{-1}(\mathscr{Z})$, where $\mathscr{Z}$ is the stabilizer of w in Q_v. This proves (iv).   q.e.d.

6.8. We next turn to the rank-2 version of Theorem D of the introduction.

Proposition. *Let Λ be a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2. Suppose that a finitely presented group Γ admits a nontrivial action, without inversions, on a Λ-tree, and that the action satisfies the Ascending Chain Condition 6.1. Then Γ splits over some subgroup.*

Proof. By 6.5, we can identify Γ with the fundamental group of a graph of groups $\mathscr{G}$ in such a way that each vertex group is either contained in the stabilizer of a point of T or admits a homomorphism onto a splittable planar discontinuous group. If all the vertex groups are proper subgroups, it follows from 2.8 that Γ splits. If one of the vertex groups is

equal to Γ, then Γ either has a fixed point in T—a contradiction to the hypothesis that the action is nontrivial—or admits a homomorphism onto a splittable planar discontinuous group Q. In the latter case, Q splits by the discussion in 5.32, and hence so does Γ.   q.e.d.

6.9.   The following result is the rank-2 case of Theorem C of the introduction.

Proposition.   *Let Λ be a subgroup of $\mathbb{R}$ whose ℚ-rank is 2. Let Γ be a finitely presented group with the property that every small subgroup of Γ is finitely generated. Suppose that Γ admits a nontrivial action, without inversions, on a Λ-tree T and that the stabilizer of every nondegenerate segment in T is a small subgroup of Γ. Then Γ splits over a small subgroup.*

Proof.   First note that the action of Γ on T satisfies the Ascending Chain Condition 6.1; for a monotone union of small subgroups of Γ is small and hence finitely generated by the hypothesis.

Let us apply Proposition 6.5 to the action of Γ on T. Thus we identify Γ with the fundamental group of a graph of groups with vertex set $V_1 \amalg V_2$ in such a way that conditions (i)–(iv) of 6.5 hold.

We claim that the edge groups in $\Gamma = \pi_1(\mathscr{G})$ are all small subgroups. If e is an edge whose endpoints both belong to V_1, then by (iii), $\mathscr{G}(e)$ is contained in the stabilizer of a nondegenerate segment in T and is therefore small by hypothesis. If at least one of the endpoints of e, say v, belongs to V_2, then by (iv), $\mathscr{G}(e)$ is an extension of a cyclic or dihedral group by the group N_v. By (ii), N_v fixes a nondegenerate segment and is therefore small. Hence $\mathscr{G}(e)$ is small in this case as well.

If the vertex groups in $\pi_1(\mathscr{G})$ are all proper, the conclusion of the proposition now follows from 2.8. If one of the vertex groups $\mathscr{G}(v)$ is equal to Γ, then since the action of Γ on T is nontrivial we must have $v \in V_2$. Thus Γ admits a homomorphism onto a splittable planar discontinuous group Q_v. As we observed above, the kernel N_v is small. By the discussion in 5.32, Q_v splits over a small subgroup. Now Γ inherits a splitting from Q_v; and the group over which Γ splits is an extension of a small group by the small group N_v and is therefore small.   q.e.d.

Finally, here is the rank-2 case of Theorem B of the introduction.

6.10.   Proposition.   *Let F be a compact surface without boundary, let Λ be a subgroup of $\mathbb{R}$ whose ℚ-rank is at most 2, and let $\pi_1(F)$ act on a Λ-tree T in such a way that the stabilizer of every nondegenerate segment in T is a cyclic subgroup of Λ. Then the length function defined by the action has the form l_μ for some measured foliation μ on F.*

6.11.   The proof of 6.10 will use both 6.2 and the following general fact.

Proposition. *Let* F *be a compact surface without boundary and let* $\mathscr{S} = (\Sigma, \mu, \Gamma, \rho)$ *be a uniform, 1-connected measured foliated singular surface with symmetry, where* $\Gamma = \pi_1(F)$. *Let* $\mathscr{T} = (T, \Gamma, \bar{\rho})$ *denote the leaf of* $\mathscr{S}$, *and suppose that the stabilizer of every nondegenerate segment in* T *is a cyclic subgroup of* Λ. *Then the length function defined by* $\mathscr{T}_0$ *has the form* l_μ *for some measured foliation* μ *on* F.

6.12. *Proof that 6.2 and 6.11 imply 6.10.* The action of Γ on T satisfies the Ascending Chain Condition 6.1, since any ascending chain of cyclic subgroups in the surface group Γ must terminate. Thus the action satisfies the hypothesis of 6.2; let (T_i, f_{ij}) be the strongly convergent system given by 6.2. Let l and l_i denote the length function defined by the actions of Γ on T and T_i respectively. It follows from 6.11 that l_i is the length function associated to some measured foliation μ_i on F. Since Γ acts nontrivially on T, the function l is not identically zero. By 1.28, the l_i converge strongly to l.

Recall from 5.31 that l_i assigns to each $g \in \Gamma$ the intersection numbers $I(\mu_i, \sigma)$, where σ is the homotopy class of curves corresponding to the conjugacy class of g. It follows from [5, Exposé 6, §IV.5 and Appendix], that there are a finite number of simple closed curves $\sigma_1, \cdots, \sigma_N$ of Γ such that any two measured foliations μ and μ' which satisfy $I(\mu, \sigma_j) = I(\mu', \sigma_j)$ for $i = 1, \cdots, n$ are equivalent via Whitehead moves [5, Exposé 5]; this implies that μ and μ' satisfy $I(\mu, \sigma) = I(\mu', \sigma)$ for every closed curve σ in Σ. Hence if i and i' are positive integers such that l_i and $l_{i'}$ agree on elements $g_1, \cdots, g_N$ of Γ whose conjugacy classes correspond to $\sigma_1, \cdots, \sigma_N$, then $l_i = l_{i'}$. Since the l_i converge strongly to l, it follows that $l_i = l$ for some i.   q.e.d.

The rest of this section is devoted to the proof of 6.11.

The proof will involve the notion of a measured foliation of a (PL) surface with boundary. This may be defined as in [5], except that as usual we take the coordinate charts to be PL rather than smooth. The definition of a taut path (5.14) goes through without change for the case of a measured foliated surface with boundary.

6.13. Let (Σ_0, μ_0) be a measured foliated surface with boundary and let (Σ_1, μ_1) be a measured foliated singular surface. By a *taut map* from (Σ_0, μ_0) to (Σ_1, μ_1) we shall mean a continuous map $W: \Sigma_0 \to \Sigma_1$ which maps each leaf in Σ_0 into a leaf in Σ_1, and such that for every taut path γ in Σ_0, the path $W \circ \gamma$ is also taut and has the same length as γ.

6.14. For the rest of this section, the hypotheses of 6.11 will be understood to hold. Let K be the simplicial tree associated to Σ by the

construction of 6.6. In the notation of 6.6, the vertex set of K is a disjoint union $S \amalg S'$.

By [1] or [4], K has a minimal Γ-invariant subtree K_0. Let Σ_0 denote the union of all global branches of Σ corresponding to vertices in $K_0 \cap S$. Then $T_0 = \chi_\mu(\Sigma_0)$ is a Γ-invariant subtree of T; by [1] or [4], the actions of Γ on T and T_0 define the same length function. Hence we may assume that K is already minimal.

As in the proof of 6.5, we write $S = S_1 \amalg S_2$, where every vertex in S_1 corresponds to an isolated 1-simplex and every vertex in S_2 corresponds to a surface with points at infinity. It follows from the proof of 6.5 and 6.9 that the stabilizers of the edges in K are small subgroups of Γ; since Γ is a surface group, and since the edge stabilizers have infinite index in Γ, they are in fact cyclic.

It now follows from [12, Theorem III.2.6] that the action of Γ on the simplicial tree K is "F-geometric". This means that there is a compact two-sided 1-manifold $C \subset F$, with no homotopically trivial components, such that if $\widetilde{C}$ denotes the preimage of C in the universal cover $\widetilde{F}$ of F, then the vertices of K are in 1-1 correspondence with the components of $\widetilde{F} - \widetilde{C}$, the edges of K are in 1-1 correspondence with the components of $\widetilde{C}$, and an edge is incident to a vertex if and only if the corresponding component of $\widetilde{C}$ is contained in the closure of the corresponding component of $\widetilde{F} - \widetilde{C}$. The action of Γ on T is induced by its natural action on $\widetilde{F}$. In particular, the stabilizer of every edge of K is an *infinite* cyclic subgroup of Γ.

For each edge e of K, we let $\widetilde{C}_e$ denote the corresponding component of $\widetilde{C}$. For every vertex v of K, we let R_v denote the closure of the corresponding component of $\widetilde{F} - \widetilde{C}$. Note that each R_v is 1-connected and that its boundary components are homeomorphic to $\mathbb{R}$. We set $\widetilde{F}_0 = \bigcup_{v \in S} R_v$. For each $v \in S$, we let $\mathscr{B}_v$ denote the global branch of Σ corresponding to v. For each $w \in S'$, we let X_w denote the vertex of Σ corresponding to w.

6.15. For each $v \in S$, the stabilizer Γ_v is also the stabilizer of $\mathscr{B}_v$.

Lemma. *For each vertex $v \in S_2$, the group Γ_v is noncyclic and torsion-free, and the action of Γ_v on $\mathscr{B}_v$ is effective.*

Proof. Let Q denote the normal subgroup of Γ_v consisting of all elements that fix $\mathscr{B}_v$ (pointwise). Then Γ_v/Q acts effectively on $\mathscr{B}_v$. Since $\mathscr{S}$ is uniform, $\mathscr{B}$ is finite $\mathrm{mod}\,\Gamma$; hence it follows from 5.33 that Γ_v/Q is a splittable planar discontinuous group. In particular it is noncyclic. Hence Γ_v is noncyclic.

Since Γ_v is a subgroup of $\Gamma = \pi_1(F)$, it is itself the fundamental group of a surface. Any surface except $\mathbb{R}P^2$ has a torsion-free fundamental group. Since Γ_v is noncyclic, it must be torsion-free.

To show that Γ_v acts effectively on $\mathscr{B}_v$ we must show that $Q = \{1\}$. Assume that $Q \neq \{1\}$. Note that Q fixes $\chi_\mu(\mathscr{B}_v)$. In particular, Q fixes a nondegenerate segment in T and is therefore cyclic. But Γ_v is the fundamental group of a surface; since it contains a nontrivial cyclic normal subgroup, Γ_v must itself be either cyclic or free abelian of rank 2. Hence Γ_v/Q is cyclic. This is a contradiction.

6.16. **Lemma.** *For any $v \in S_2$ the action of Γ_v on the interior $\overset{\circ}{\mathscr{B}}_v$ of $\mathscr{B}_v$ is free. Furthermore, $\overset{\circ}{\mathscr{B}}_v$ is contractible.*

Proof. It follows from 6.15 that Γ_v is torsion-free and acts effectively on $\overset{\circ}{\mathscr{B}}_v$. But any effective, properly discontinuous action of a torsion-free group on a surface is free.

It also follows from 6.15 that Γ_v is infinite. Since Γ_v acts properly discontinuously on $\overset{\circ}{\mathscr{B}}_v$ it follows that $\overset{\circ}{\mathscr{B}}_v$ is noncompact. But any noncompact, 1-connected surface is contractible.

6.17. **Lemma.** *Suppose that $v \in S_2$ and $w \in S'$ are joined by an edge in K. Then x_w is a point at infinity in $\mathscr{B}_v$.*

Proof. By 6.15, Γ_v acts effectively and simplicially on $\mathscr{B}_v$; hence it acts properly discontinuously on $\overset{\circ}{\mathscr{B}}_v$. But the stabilizer of the point x_w in Γ_v is $\Gamma_v \cap \Gamma_w$, which is the stabilizer of an edge in K and is therefore infinite cyclic according to 6.14. Therefore $x_w \notin \overset{\circ}{\mathscr{B}}_v$.

6.18. **Lemma.** *There exist a Γ-invariant measured foliation $\tilde{\nu}_0$ on $\tilde{F}_0$ and a Γ-equivariant taut map $f: (\tilde{F}_0, \tilde{\nu}_0) \to (\Sigma, \mu)$. Furthermore, for every $v \in S$ we have $f(R_v) = \mathscr{B}_v$, and if e is any edge of K with endpoints $v \in S$ and $w \in S'$, then $f_v(\tilde{C}_e) = \{x_w\}$.*

Proof. We shall show that for every $v \in S$ there exist a Γ_v-invariant measured foliation ν_v of R_v and a Γ_v-equivariant taut map $f_v: (R_v, \nu_v) \to (\mathscr{B}_v, \mu|\mathscr{B}_v)$ such that if e is any edge of K having v as one endpoint, and if $w \in S'$ denotes the other endpoint of e, then $f_v(\tilde{C}_e) = \{x_v\}$. This will imply the lemma; for if S^* is a complete set of orbit representatives for the action of Γ on S we may then define f and $\tilde{\nu}_0$ by setting $f|R_v = f_v$ and $\tilde{\nu}_0|R_v = \nu_v$ for each $v \in S^*$ and extending Γ-equivariantly. The Γ_v-equivariance of f_v and ν_v guarantees that these extensions are well defined.

To define f_v we first consider the case where $v \in S_1$. Then $\mathcal{B}_v$ is an arc with endpoints x_{w_0} and x_{w_1}, where the w_i are vertices in S'; the edges incident to v are e_0 and e_1, where $e_i = \{v, w_i\}$. Hence the boundary components of R_v are $\widetilde{C}_0$ and $\widetilde{C}_1$, where $\widetilde{C}_i = \widetilde{C}_{e_i}$. Thus R_v/Γ_v is a compact surface with boundary whose universal cover, R_v, has exactly two boundary components; hence R_v/Γ_v is PL homeomorphic to an annulus or a Möbius band. It follows that R_v admits a trivial fibration over $[0, 1]$ which is invariant under the infinite cyclic covering group Γ_v and has fiber $\mathbb{R}$. We may take the projection $p: R_v \to [0, 1]$ to map $\widetilde{C}_i$ to i for $i = 0, 1$.

If R_v/Γ_v is an annulus then each fiber in R_v is Γ_v-invariant. Hence if $h: [0, 1] \to \mathcal{B}_v$ is any PL homeomorphism such that $h(\widetilde{C}_i) = x_{w_i}$ for $i = 0, 1$, then $f_v = h \circ p$ is Γ_v-equivariant. Furthermore, since f_v is a trivial fibration of R_v over $\mathcal{B}_v$, it is clear that there is a Γ_v-invariant measured foliation $(\mathcal{G}_v, \nu_v)$ that makes f taut: the leaves of $\mathcal{G}_v$ are the fibers of f_v, and ν_v is defined by $\nu_v(\gamma) = \mu(f_v \circ \gamma)$ for any piecewise-monotone path γ.

If R_v/Γ_v is a Möbius band, and if g is a generator of Γ_v, we may choose the fibration p so that $p(g \cdot x) = 1 - p(x)$ for every $x \in R_v$. In particular g interchanges the two boundary components of R_v. Hence g interchanges e_0 and e_1, and therefore also interchanges x_{w_0} and x_{w_1}. Hence there is a PL homeomorphism $h: [0, 1] \to \mathcal{B}_v$ such that $h(\widetilde{C}_i) = x_{w_i}$ for $i = 0, 1$, and such that $h(1 - t) = g \cdot h(t)$ for every $t \in [0, 1]$. We again set $f_v = h \circ p$; once again, f_v is a trivial fibration and is Γ-equivariant, and we proceed as in the case of an annulus.

Now consider the case $v \in S_2$. Since the action of Γ_v on R_v is free and properly discontinuous, there is a Γ_v-invariant boundary collar H in R_v. We set $R'_v = \overline{R_v - H}$. Let M denote the link of v in K. For each $w \in M$, we write $e(w)$ for the 1-simplex $\{v, w\}$, and H_w denotes the component of H containing $\widetilde{C}_{e(w)}$; we set $Y_w = H_w \cap R'_v$. Thus R'_v is a surface with boundary, and the Y_w are the components of $\partial R'_v$.

If w is any vertex in M, then by 6.17, x_w is a point at infinity in $\mathcal{B}_v$. We choose a standard neighborhood Z_w of x_w in $\mathcal{B}_v$, and let X_w denote the frontier of Z_w in $\mathcal{B}_v$. Since Γ_v acts simplicially on $\mathcal{B}_v$, and acts freely on $\overset{\circ}{\mathcal{B}}_v$, we may choose the Z_w for $w \in M$ so that $Z = \bigcup_{w \in P} Z_w$ is Γ-invariant, and so that when $w \neq w'$ we have either $Z_w = Z_{w'}$ or $Z_w \cap Z_{w'} = \varnothing$. We set $\mathcal{B}'_v = (\mathcal{B}_v - Z) \cap \overset{\circ}{\mathcal{B}}_v$; clearly $\mathcal{B}'_v$ is a Γ_v-invariant

surface with boundary, and the X_w are its boundary components. The action of Γ_v on $\mathscr{B}_v'$ is free and properly discontinuous.

Now let M^* be a complete set of orbit representatives for the action of Γ_v on M. For each $w \in M^*$, the infinite cyclic group $\Gamma_{e(w)} = \Gamma_v \cap \Gamma_w$ acts freely and properly discontinuously on each of the connected 1-manifolds Y_w and X_w. Hence there is a $\Gamma_{e(w)}$-equivariant PL homeomorphism $h_w : Y_w \to X_w$. Define a Γ_v-equivariant map $h_\partial : \partial R_v' \to \partial \mathscr{B}_v'$ by setting $h_\partial | Y_w = h_w$ for each $w \in M^*$ and extending equivariantly. For each $w \in M$ we have $h_\partial(Y_w) \subset X_w$.

Since Γ_v acts freely and properly discontinuously on R_v', and since $\mathscr{B}_v'$ is contractible by 6.16, we may extend h_∂ to a Γ_v-equivariant PL map $h : R_v' \to \mathscr{B}_v'$. Consider the induced map $\overline{h} : \overline{R}_v' \to \overline{\mathscr{B}}_v'$, where $\overline{R}_v' = R_v'/\Gamma_v$ and $\overline{\mathscr{B}}_v' = \mathscr{B}_v'/\Gamma_v$. Since the action of Γ_v on $\mathscr{B}_v$ is also properly discontinuous, and is free by 6.16, it follows from covering space theory that $\overline{h}$ induces an isomorphism of fundamental groups. By construction, $\overline{h}$ is boundary-preserving, i.e., $\overline{h}(\partial \overline{R}_v') \subset \partial \overline{\mathscr{B}}_v'$.

It follows from Theorem 13.1 of [7] that if a boundary-preserving map between surfaces with boundary induces an isomorphism of fundamental groups, and if the domain surface is compact and has noncyclic fundamental group, then the map is homotopic, through a family of boundary-preserving maps, to a PL homeomorphism. Since $\overline{R}_v'$ may be identified with a subsurface of F bounded by certain components of C, it is compact; and by 6.15, $\pi_1(\overline{R}_v') \cong \Gamma_v$ is noncyclic. Hence $\overline{h}$ is homotopic by a boundary-preserving homotopy to a PL homeomorphism $\overline{h}' : \overline{R}_v' \to \overline{\mathscr{B}}_v'$. By the covering homotopy property for covering spaces, $\overline{h}'$ is covered by a Γ_v-equivariant PL homeomorphism $h' : R_v' \to \mathscr{B}_v'$ such that $h_\partial(Y_w) = X_w$ for each $w \in M$.

We wish to extend h' to a map $f_v : R_v \to \mathscr{B}_v$. For this purpose we consider an arbitrary vertex w in M^*, and fix a generator g_w of $\Gamma_{e(w)}$. Let A be an arc which is a fundamental domain for the action of $\Gamma_{e(w)}$ on Y_w. We may suppose A to be chosen so that its endpoints belong to prongs in Z_w. Then A is the intersection of Z_w with unions of n sectors, where n is some positive integer.

We identify H_w via a PL homeomorphism with the product $R \times [0, 1]$ so that $\mathbb{R} \times \{0\} = \tilde{C}_{e(w)}$ and $\mathbb{R} \times \{1\} = Y_w$, and so that $g_w | H_w$ is the translation $(s, t) \mapsto (s + n, t)$. We make this identification in such a way that for $1 \leq i \leq n$ we have $h'([i-1, i]) \times \{1\}) = D_i \cap X_w$ for some sector D_i in Z.

Let P_i denote the prong in Z_w containing $h'(i, 1)$. Note that D_i is a PL disk whose boundary consists of the arcs P_{i-1}, P_i and $D_i \cap X_w$. Hence we may extend $h'|Y_w$ to a $\Gamma_{e(w)}$-equivariant continuous map f_w from $H_w = \mathbb{R} \times [0, 1]$ to Z_w, which maps $\mathbb{R} \times \{0\}$ to $\{x_w\}$, maps $\{i\} \times [0, 1]$ by a PL homeomorphism onto P_i, and maps $(i-1, i) \times (0, 1)$ homeomorphically onto $\text{int} D_i$, for each $i \in \mathbb{Z}$.

For each $i \in \mathbb{Z}$, the map f_v defines a decomposition of $[i-1, i] \times [0, 1]$ whose elements are the preimages of the elements of the decomposition of D_i induced by the underlying foliation of μ. While f_v cannot be PL, it follows from the construction that the induced decomposition of $[i - 1, i] \times [0, 1]$ is a product decomposition in the topological category. Hence after precomposing f_v with a $\Gamma_{e(w)}$-equivariant self-homeomorphism of H_w, we may assume that the decomposition of $[i - 1, i] \times [0, 1]$ defined by f_v is a PL product decomposition.

Now define a map $f_H \colon H \to Z$ by setting $f_H|H_w = f_w$ for each $w \in M$ and extending Γ_v-equivariantly. We define $f_v \colon R_v \to \mathscr{B}_v$ to agree with h' on R'_v and with f_H on H. It is now straightforward to verify that the preimages under f_v of the leaves of $\mu|\mathscr{B}_v$ form a (PL) foliation $\mathscr{G}$ of R_v. The diagram depicts $\mathscr{G}$ in the neighborhood of a component of ∂R_v.

It is also straightforward to show that for each path γ in R_v which is piecewise-monotonic with respect to $\mathscr{G}$, the path $f_v \circ \gamma$ is piecewise-monotonic; and that the function ν_v defined by $\nu_v(\gamma) = \mu(f_v \circ \gamma)$ is a transverse measure for $\mathscr{G}$. Thus $f \colon (R_v, \nu_v) \to (\Sigma, \mu)$ is taut. q.e.d.

6.19. Let $\tilde{\nu}_0$ be the foliation, and f the map, given by the preceding lemma. Since $\tilde{\nu}_0$ is Γ-invariant it induces a measured foliation ν_0 on the surface $F_0 = \tilde{F}_0/\Gamma$. We regard the surface-with-boundary F_0 as a subsurface of F whose boundary is C. It follows from the proof of the lemma that every component of C is a union of classical leaves and singular points of the foliation $\tilde{\nu}_0$. Thus C is an invariant set of ν_0 in the sense of [5].

As in [5, Exposé 5, §III.1], using the measured foliation ν_0 on F_0, we can obtain a measured foliation ν on F by *élargissement*. Recall that ν is constructed by using a surjective (PL) immersion $j \colon F_0 \to F$

which maps ∂F_0 onto a spine $\mathfrak{G}$ of $\overline{F - F_0}$, maps the interior of F_0 homeomorphically onto $F - \mathfrak{G}$, and is the identity outside a boundary collar in F_0. One defines a measured foliation on $F - \mathfrak{G}$ by pushing $\nu_0 | \operatorname{int} F_0$ forward via the homeomorphism $j | \operatorname{int} F_0$; this pushed-forward measured foliation admits a unique extension to F, which we denote by ν.

Since j is homotopic to the inclusion map $F_0 \to F$, the covering homotopy property for covering spaces gives a Γ-equivariant map $\tilde{j}: \tilde{F}_0 \to \tilde{F}$ covering j. Furthermore, $\tilde{j}_{\operatorname{int}} = \tilde{j} | \operatorname{int} \tilde{F}_0$ is a homeomorphism of $\operatorname{int} \tilde{F}_0$ onto $\tilde{F} - \tilde{\mathscr{G}}$, where $\tilde{G}$ denotes the preimage of $\mathfrak{G}$ in $\tilde{F}$.

We define a map $W: \tilde{F} \to \Sigma$ as follows. We set $W | \tilde{F} - \tilde{\mathfrak{G}} = f \circ (\tilde{j}_{\operatorname{int}})^{-1}$. To define $W | \tilde{\mathfrak{G}}$, note that each component of $\tilde{F} = \tilde{F}_0$ contains a unique component of $\tilde{\mathfrak{G}}$, and that any component of $\tilde{F} - \tilde{F}_0$ has the form R_w for some $w \in S'$. We let $\tilde{\mathfrak{G}}_w$ denote the component of $\tilde{\mathfrak{G}}$ contained in R_w, and define $W | \tilde{\mathfrak{G}}$ by setting $W(\tilde{\mathfrak{G}}_w) = \{x_w\}$.

The properties of f stated in Lemma 6.18 imply that W is continuous. It is clear that W is Γ-equivariant. The continuity of W implies that $W \circ \tilde{j} = f$, since $W \circ \tilde{j}$ and f obviously agree on the dense subset $\tilde{F} - \tilde{\mathfrak{G}}$ of $\tilde{F}$.

6.20. Lemma. *The map W is a taut map from $(\tilde{F}, \tilde{\nu})$ to (Σ, μ).*

Proof. Since $W \circ \tilde{j} = f: (\tilde{F}_0, \tilde{\nu}_0) \to (\Sigma, \mu)$ is taut, it is clear that W maps leaves into leaves. We must show that if γ is any taut path in $\tilde{F}$ then $W \circ \gamma$ is taut and has the same length as γ. By 5.16 we may assume that γ is quasi-transverse. After reparametrization we may further assume that γ has the form $\gamma_1 * \cdots * \gamma_n$, where each γ_i is a path supported in the closure of some component of $\tilde{F} - \tilde{\mathfrak{G}}$. Thus for each i we have $|\gamma_i| \subset \tilde{j}(R_i)$, where $R_i = R_{v_i}$ for some $v_i \in S$. Among all reparametrizations of this type we choose one for which n has the smallest possible value; this guarantees that $v_i \neq v_{i+1}$ for $1 \leq i < n$.

Since f is a taut map from $(\tilde{F}_0, \tilde{\nu}_0)$ to (Σ, μ), it follows from the definition of $\tilde{\nu}$ that $W | \tilde{j}(R_v): (\tilde{j}(R_v), \tilde{\nu} | \tilde{j}(R_v)) \to (\Sigma, \mu)$ is taut for every $v \in S$. Taking $v = v_i$ we conclude that for each i, the path $W \circ \gamma_i$ is taut and has the same length as γ_i. In particular, $W \circ \gamma$ is piecewise-monotonic and has the same length as γ.

It remains to show that $W \circ \gamma$ is taut. Set $\chi = \chi_\mu: \Sigma \to \mathbb{R}T$. We must show that $\delta = \chi \circ W \circ \gamma$ is a monotonic map of $[0, 1]$ onto a segment in $\mathbb{R}T$. We have $\delta = \delta_1 * \cdots * \delta_n$, where $\delta_i = \chi \circ W \circ \gamma_i$. For each i, since each $W \circ \gamma_i$ is taut, δ_i is a monotonic map of $[0, 1]$ onto a segment in

$\mathbb{R}T$. Since $\mathbb{R}T$ is an $\mathbb{R}$-tree, we need only show that $|\delta_i| \cap |\delta_{i+1}| = \delta_i(1)$ for $1 \leq i < n$.

Note that for $1 \leq i \leq n$, the path $W \circ \gamma_i$ is supported in the global branch $\mathcal{B}_i = \mathcal{B}_{v_i}$ of Σ, so that $|\delta_i| \subset \mathcal{B}_i$. For $1 \leq i < n$, we have $v_i \neq v_{i+1}$ and therefore $\mathcal{B}_i \neq \mathcal{B}_{i+1}$. The point $z = W(\gamma_i(1))$ lies in both $\mathcal{B}_i$ and $\mathcal{B}_{i+1}$. Since Σ is 1-connected, we must have $\mathcal{B}_i \cap \mathcal{B}_{i+1} = \{z\}$, and any path in Σ that meets both $\mathcal{B}_i$ and $\mathcal{B}_{i+1}$ must contain z. Hence any leaf of μ that meets $\mathcal{B}_i$ and $\mathcal{B}_{i+1}$ must contain z; thus $\chi(\mathcal{B}_i) \cap \chi(\mathcal{B}_{i+1}) = \{z\}$. In particular we have $|\delta_i| \cap |\delta_{i+1}| = \{z\}$, as required.   q.e.d.

Proof of 6.11, *concluded.* Let T' denote the dual tree of ν. Since W is taut, it maps each leaf in $\tilde{F}$ into a leaf in Σ; hence there is a unique map ι of the leaf space T' of $(\tilde{F}, \tilde{\nu})$ into the leaf space $\mathbb{R}T$ of (Σ, μ) such that $\iota \circ \chi_{\tilde{\nu}} = x_\mu \circ W$. The definition of a taut map implies that ι is an isometric embedding of $\mathbb{R}$-trees. Since W is Γ-equivariant, so is ι. In particular, the minimal Γ-invariant subtrees of $\mathbb{R}T$ and T' are Γ-equivariantly isometric. Hence the actions of Γ on $\mathbb{R}T$ and on T' define the same length function.   q.e.d.

Proposition 6.10 has a relative version (cf. [12, III.1.6]) for the case where Γ is the fundamental group of a surface with boundary. We leave the details to the reader.

7. T-complexes as foliated objects

In §5 we showed how a Λ-foliated singular surface can define a Λ-tree. In this section we shall show that if T is a Λ-tree then certain T-complexes can be regarded as Λ-foliated singular surfaces.

7.1. Let X be a complex, and let l be a function that assigns a positive real number to each 1-simplex of X. We shall call l a *length system* for X if for every 2-simplex δ, the edges of δ can be labelled τ, τ', and τ'' in such a way that $l(\tau) = l(\tau') + l(\tau'')$. The edge τ, which is uniquely determined by l and δ, is called the *long edge* of δ.

Let l be a length system on a complex X. Let Δ be any simplex of X. There is an affine map $f: \Delta \to \mathbb{R}$ which maps each edge τ of Δ onto an interval of length $l(\tau)$; such a map will be called an *l-decomposing function* on Δ. An l-decomposing function is unique up to composition with an isometry of $\mathbb{R}$. In particular, the sets $f^{-1}(s)$ for $s \in f(\Delta)$ constitute a decomposition of Δ which is independent of the choice of l-decomposing function f; we call it the *decomposition of Δ defined by l*.

7.2. Now let us consider a length system l on a singular surface Σ. If β is a 2-dimensional local branch at a vertex v of Σ, we define the *order* of β to the cardinality of the set of all 2-simplices $\delta \subset \beta$ such that x is not incident to the longest edge of β. We shall say that l is *nondegenerate* if the following conditions hold for every vertex v of Σ:

(i) each 2-dimensional local branch at v has order at least 2; and

(ii) each noncompact 2-dimensional local branch at v has infinite order.

7.3. Proposition. *Let l be a nondegenerate length system on a singular surface Σ. Then there is a unique 0-tame measured foliation $(\mathscr{F}, \mu)$ on Σ such that for each simplex Δ of Σ,*

(i) *the decomposition of Δ induced by $\mathscr{F}$ is the decomposition defined by l, and*

(ii) *the length $\mu(\gamma)$ of any monotonic path γ in Δ is the length of the interval $f(|\gamma|)$, where f is an l-decomposing function on Δ.*

Proof. We shall describe the measured foliation $(\mathscr{F}, \mu)$ by local data as in 5.11. For every point x of Σ and every local branch β at x we must define a neighborhood U_β of x in β, a standard decomposition $\mathscr{F}_\beta$ of U_β, and a standard measure μ_D on each sector D of U_β.

If β is 1-dimensional then it consists of a single 1-simplex τ; in this case we set $U_\beta = \tau$. An l-decomposing function on τ is a homeomorphism of τ onto an interval in $\mathbb{R}$, and hence defines a standard measure on τ.

If x is an interior point of a 2-simplex δ of Σ then the decomposition of δ defined by l induces a product decomposition $\mathscr{F}_\beta$ of some neighborhood U_β of x in δ. Furthermore, an l-decomposing function on δ induces a homeomorphism of the quotient of this product decomposition onto an interval in $\mathbb{R}$, and hence defines a standard measure on each sector of U_β.

Now suppose that x is an interior point of a 1-simplex τ which has valence 2 in Σ. Let δ_1 and δ_2 denote the 2-simplices incident to τ, let f_i be l-decomposing functions on the δ_i such that $f_1|\tau = f_2|\tau$, and define a map $f: \delta_1 \cup \delta_2 \to \mathbb{R}$ by $f|\delta_i = f_i$. Then the fibers of f constitute a decomposition of $\delta_1 \cup \delta_2$, which induces a product decomposition $\mathscr{F}_\beta$ of some neighborhood U_β of x in δ. Again, f induces a homeomorphism of the quotient of this product decomposition onto an interval in $\mathbb{R}$, and hence defines a standard measure on each of the two sectors of U_β.

Finally, suppose that x is a vertex of Σ. For each 2-simplex δ in β, let f_δ be an l-decomposing function on δ such that $f_\delta(x) = 0$ and

$f_\delta(\delta) \subset [0, \infty)$. Then there is a map $f: \beta \to \mathbb{R}$ such that $f|\delta = f_\delta$ for each 2-simplex δ in β.

If $\delta \subset \beta$ is a 2-simplex whose longest edge is incident to x, then $f_\delta^{-1}(0) = \{x\}$; on the other hand, if the longest edge τ of δ is not incident to x, then $f_\delta^{-1}(0)$ is a line segment having one endpoint at x and the other in the interior of τ. Hence, if we regard β as the cone over a component L of the link of x, then $f^{-1}(0)$ is the cone over a discrete subset K of L; the number of points in K is the order k of β with respect to l. By condition (i) of 7.2 we have $k \geq 2$; by condition (ii), each component of $L - K$ is relatively compact.

For each component E of $\beta - f^{-1}(0)$, we choose a number $\varepsilon_E > 0$ such that each 1-simplex contained in $\overline{E}$ has length $< \varepsilon$. We set $D_E = \overline{E} \cap f^{-1}([0, \varepsilon])$. We define U_β to be the union of the D_E as E ranges over the components of $\beta - f^{-1}(0)$. The fibers of $f|U_\beta$ constitute a decomposition of U_β. It is easily seen that for each component E of $\beta - f^{-1}(0)$, the decomposition $\mathscr{F}_\beta$ induces a product decomposition of D_E. Hence $\mathscr{F}_\beta$ is a standard k-prong decomposition of U_β, and the sets D_E are the sectors of U_β. Furthermore, for each sector D_E of U_β, the map f induces a homeomorphism of the quotient of the product decomposition of D_E onto an interval in $\mathbb{R}$, and hence defines a standard measure on D_E.

The compatibility conditions of 5.11 are straightforward to verify; thus we have defined a measured foliation $(\mathscr{F}, \mu)$. It is immediate from the construction that $(\mathscr{F}, \mu)$ is tame and satisfies conditions (i) and (ii) of the proposition. The uniqueness assertion is also clear.   q.e.d.

7.4. Proposition. *Let Σ be a singular surface, let Λ be a subgroup of $\mathbb{R}$, let l be a nondegenerate length system for Σ taking values in Λ, and let μ denote the measured foliation defined by l. Then (Σ, μ) is a Λ-foliated singular surface.*

Proof. We must show that for any taut path γ in Σ whose endpoints are vertices we have $\mu(\gamma) \in \Lambda$. By 5.16 we may assume that γ is a piecewise-linear quasi-transverse arc. We may assume that γ has the form $\gamma = \gamma_1 * \cdots * \gamma_n$, where each γ_i is a monotonic arc in a 2-simplex δ_i of Σ having its endpoints in the boundary of δ_i, and its interior in the interior of δ_i. We argue by induction on the integer n.

If $n = 1$ then $\mu(\gamma) = l(\tau)$, where τ is the edge of δ_1 with endpoints $\gamma(0)$ and $\gamma(1)$. Hence in this case we have $\mu(\gamma) \in \Lambda$. If $n > 1$, and if there is some j with $1 \leq j \leq n$ such that $\gamma_j(1)$ is a vertex, then the

induction hypothesis implies that $\mu(\gamma_1 * \cdots * \gamma_j)$ and $\mu(\gamma_{j+1} * \cdots * \gamma_n)$ both belong to Λ; hence their sum $\mu(\gamma_1 * \cdots * \gamma_n)$ also belongs to Λ.

Now suppose that $n > 1$ and that none of the points $\gamma_1(1), \cdots, \gamma_{n-1}(1)$ is a vertex. We may write $\delta_1 = \Delta(x, y, z)$, where $x = \gamma_1(0)$. Then $\gamma_1(1)$ is interior to one of the edges $\Delta(x, y)$, $\Delta(x, z)$, $\Delta(y, z)$. If $\gamma(1)$ is interior to $\Delta(x, y)$ or $\Delta(x, z)$ then $\gamma_1 * \gamma_2$ is equivalent to a monotonic path γ' in δ_2; and we have $\mu(\gamma) = \mu(\gamma' * \gamma_3 * \cdots * \gamma_n) \in \Lambda$.

Now suppose that $\gamma(1)$ is interior to $\Delta(y, z)$. By symmetry we may assume that the long edge of δ_1 is either $\Delta(x, y)$ or $\Delta(y, z)$.

If $\Delta(x, y)$ is the long edge of δ_1, then γ_1 is equivalent to $\gamma_1' * \gamma_1''$, where γ_1' is a parametrization of $\Delta(x, z)$, and γ_1'' is a parametrization of the subarc of $\Delta(y, z)$ having endpoints z and $\gamma_1(1)$. We have $\mu(\gamma_1') = l(\Delta(y, z)) \in \Lambda$. On the other hand, it is clear in this case that $\gamma_1'' * \cdots * \gamma_n$ is taut; using the induction hypothesis we can conclude that $\mu(\gamma_1'' * \cdots * \gamma_n) \in \Lambda$. Hence $\mu(\gamma) \in \Lambda$.

If $\Delta(y, z)$ is the long edge of δ_1, then x belongs to a nondegenerate element λ of the decomposition of δ_1; furthermore, λ is a line segment having x as one endpoint. The other endpoint w of λ is an interior point w of $\Delta(y, z)$, which divides $\Delta(y, z)$ into two subarcs A and B, where A has endpoints y and w, and B has endpoints z and w. The point $\delta_1(1)$ lies in either A or B. By symmetry we may assume that $\delta_1(1) \in A$.

Let γ' be a parametrized path in $\Delta(y, z)$ having initial point z and terminal point w. Then the path $\gamma' * \gamma_2 * \cdots * \gamma_n$ is taut, and the induction hypothesis implies that $\mu(\gamma') + \mu(\gamma_2 * \cdots * \gamma_n) \in \Lambda$. On the other hand we have $\mu(\gamma') - \mu(\gamma_1) = l(\Delta(x, z)) \in \Lambda$, so that $\mu(\gamma) = \mu(\gamma_1) + \mu(\gamma_2 * \cdots * \gamma_n) \in \Lambda$. q.e.d.

7.5. Let T be a Λ-tree, and let (X, ϕ) be a T-complex of dimension ≤ 2. For each 1-simplex τ of X let us write $l(\tau)$ for the length of τ in the sense of 3.6.

Let δ be any closed 2-simplex of X. The definition of a geometric T-complex implies that $\phi|\delta$ is an affine map of δ onto the segment $|\delta| \subset \mathbb{R}T$, and that ϕ maps two of the vertices of δ, say x and z, to the endpoints of $|\delta|$, and the third, say y, to an interior point. The edges of δ are $\tau = \Delta(x, z)$, $\tau_1 = \Delta(x, y)$ and $\tau_2 = \Delta(y, z)$. Their lengths and supports satisfy $l(\tau) = l(\tau_1) + l(\tau_2)$ and $|\tau| = |\tau_1| \cup |\tau_2| = |\delta|$. It follows that l is a Λ-valued length system. We call l the *natural length function* for the T-complex X.

Let δ be a 2-simplex of X. According to 7.1, the length system l defines a decomposition of δ. This decomposition is readily described

in terms of T: the map $\phi \colon X \to \mathbb{R}T$ restricts to a map of δ onto the segment $|\delta| \subset \mathbb{R}T$, and the elements of the decomposition are the fibers over points in $|\delta|$.

It follows that the decomposition of X determined by l also has a simple description in terms of the tree T: the elements of the decomposition are the connected components of the fibres of the map $\phi \colon X \to \mathbb{R}T$.

If γ is a parametrization of a segment $[x, y] \subset \mathbb{R}T$, and if $\hat{\gamma}$ is a piecewise-monotonic lift of γ to X, then we have $l(\gamma) = d(x, y)$.

7.6. Proposition. *Let T be a Λ-tree and let (Σ, ϕ) be a T-complex which is a 1-connected singular surface. Suppose that the natural length system for Σ is nondegenerate and hence defines a Λ-foliation μ on Σ. Suppose also that (Σ, μ) is complete. Let T_0 denote the preleaf space of (Σ, μ) and set $\chi = \chi_\mu | \operatorname{sk}_0(\Sigma)$. Then there is a unique morphism of Λ-pretrees $f \colon T_0 \to T$ such that $f \circ \chi = \phi$.*

Proof. By 7.5, ϕ maps each leaf of (Σ, l) to a point of $\mathbb{R}T$. Since, by the definition of a T-map, f maps vertices of Σ to points of T, each 0-tame leaf of (Σ, l) is mapped to a point of T. Thus there is a unique map of sets $f \colon T_0 \to T$ such that $f \circ \chi = \phi$.

To show that f is a morphism of Λ-pretrees we consider two arbitrary points $x, y \in T_0$. We may write $x = \chi(v)$ and $y = \chi(w)$, where v and w are vertices of Σ_0. We may join v to w by a path $\gamma = \gamma_1 * \cdots * \gamma_n$, where each γ_j is a parametrization of a 1-simplex τ_j of Σ_0. It follows from 5.16 that γ_j is a taut path in Σ_0; hence $\chi(\tau_j)$ is a parametrized segment in $\mathbb{R}T_0$. Since the endpoints of τ_j are vertices, the endpoints of $\chi(\tau_j)$ lie in T_0; thus $\sigma_j = \chi(\tau_j) \cap T_0$ is a presegment in T_0. Furthermore, it follows from the definition of the metric on the leaf space $\mathbb{R}T_0$ that $f | \chi(\tau_j)$ is an isometric embedding of $\chi(\tau_j)$ in $\mathbb{R}T$ for each j. In particular $f | \sigma_j$ is an isometric embedding of σ_j in T. By 1.8 it now follows that f is a morphism.   q.e.d.

7.7. Remark. It is easy to show that f maps T_0 isometrically onto a subpretree of T if and only if the fibers of ϕ are all connected. In general this need not be the case.

7.8.   Now let $\mathscr{T}$ be a Λ-tree with symmetry and let $\mathscr{S}$ be a $\mathscr{T}$-complex which is a uniform, 1-connected singular surface with symmetry. Suppose that the natural length system of $\mathscr{S}$ is nondegenerate, so that $\mathscr{S}$ is a Λ-foliated singular surface with symmetry, and let $\mathscr{T}_0$ denote its preleaf space (5.28).

Then 7.6 gives a natural morphism $\mathscr{f} \colon \mathscr{T}_0 \to \mathscr{T}$ of Λ-pretrees with symmetry. We call $\mathscr{f}$ the *natural morphism* from $\mathscr{T}_0$ to $\mathscr{T}$.

Proposition. *Let T be a Λ-tree, and let Σ_0 and Σ_1 be T-complexes which are singular surfaces. Suppose that the natural length systems for the Σ_i are nondegenerate, so that each Σ_i has a natural Λ-foliation μ_i. Then any T-map from Σ_0 to Σ_1 is a morphism between the Λ-foliated singular surfaces (Σ_0, μ_0) and (Σ_1, μ_1). Furthermore, if T_i denotes the leaf space of (Σ_1, μ_i), if $f: T_0 \to T_1$ denotes the morphism induced by W and if $f_i: T_i \to T$ denotes the natural morphism, then we have $f_1 \circ f = f_0$.*

Proof. Let $f: \Sigma_0 \to \Sigma_1$ be a T-map. Let $\phi_i: \Sigma_i \to \mathbb{R}T$ be the map defining the structure of a T-complex on Σ_i. Then f maps the fibers of ϕ_0 into fibers of ϕ_1; hence by 7.5 it maps the leaves of (Σ_0, μ_0) into leaves of (Σ_1, μ_1). Thus if T_i denotes the preleaf space of (Σ_i, μ_i), and if we set $\chi_i = \chi_{\mu_i}$, there is a map of sets $\overline{f}: \mathbb{R}T_0 \to \mathbb{R}T_1$ making the diagram

$$
\begin{array}{ccc}
\Sigma_0 & \xrightarrow{\ W\ } & \Sigma_1 \\
\chi_0 \downarrow & & \downarrow \chi_1 \\
\mathbb{R}T_0 & \xrightarrow{\ \overline{f}\ } & \mathbb{R}T_1
\end{array}
$$

commute. It is clear from 7.6 that $\mathbb{R}f_1 \circ \overline{f} = \mathbb{R}f_0$, where $f_i: T_i \to T$ denotes the natural morphism.

Since W is 0-tame, $\overline{f}$ restricts to a map $f: T_0 \to T_1$. We claim that f is a morphism.

Consider any two points $x, y \in T_0$. We may write $x = \chi_0(v)$ and $y = \chi_0(w)$, where v and w are vertices of Σ_0. We may join v to w by a path $\gamma = \gamma_1 * \cdots * \gamma_n$, where each γ_j is a parametrization of a 1-simplex τ_j of Σ_0. It follows from 5.16 that γ_j is a taut path in Σ_0; hence $\chi_0(\tau_j)$ is a parametrized segment in $\mathbb{R}T_0$. Since the endpoints of τ_j are vertices, the endpoints of $\chi(\tau_j)$ lie in T_0. According to 1.8, in order to show that f is a morphism, it is enough to show that $f|\chi(\tau_j)$ is an isometric embedding of $\chi(\tau_j)$ in $\mathbb{R}T_1$ for each j. This amounts to showing that $W \circ \gamma_j$ is a taut path.

Since Σ_0 is a T-complex, ϕ_0 maps τ_j homeomorphically onto a segment in $\mathbb{R}T$. In particular, $W|\delta_j$ is 1-1. Hence $W \circ \gamma_j$ is a parametrized arc which meets each leaf of (Σ_1, μ_1) in at most one point. It now follows from the definitions that $W \circ \gamma_j$ is taut. This shows that f is a morphism.

A similar argument shows that $\overline{f}$ is a morphism. Since $\overline{f}$ extends f it follows that $\overline{f} = \mathbb{R}f$. Hence by the above commutative diagram and

5.24, W is a morphism and f is the morphism of Λ-pretrees induced by W. Finally, since $\mathbb{R}f_1 \circ \overline{f} = \mathbb{R}f_0$, we have $f_1 \circ f = f_0$.

8. Continued fractions

8.1. Let L be a rank-two free abelian subgroup of $\mathbb{R}$. A basis of L (as a free abelian group) is called *positive* if both its elements are positive in $\mathbb{R}$. Let $\mathfrak{B} = \mathfrak{B}(L)$ denote the set of all unordered positive bases of L. If B is any element of $\mathscr{B}$, we may write B uniquely in the form $\{\alpha, \beta\}$, where $\alpha > \beta$. We shall often indicate this by writing $B = \{\alpha > \beta\}$.

If $B = \{\alpha > \beta\} \in \mathscr{B}$, then $B' = \{\alpha - \beta, \beta\}$ is again an element of $\mathscr{B}$; of course we may have $\alpha - \beta < \beta$ or $\alpha - \beta > \beta$. In any case, B' is uniquely determined by B; we shall call B' the *successor* of B.

We define a partial order on the set $\mathfrak{B}$ by writing $B' \leq B$ if there exists a finite sequence $B = B_0, B_1, \cdots, B_n = B'$ of elements of $\mathfrak{B}$ such that B_i is the successor of B_{i-1} for $i = 1, \cdots, n$. Reflexivity and transitivity are obvious; antisymmetry follows from the observation that if $B > B'$ then the greater element of B is greater than either element of B'.

8.2. The main purpose of this section is to prove the following slight refinement of a standard result [6, Theorem 175] on continued fractions, which will be needed in §9.

Proposition. *Given any* B, $B' \in \mathfrak{B}$, *there exists* $B'' \in \mathfrak{B}$ *such that* $B'' \leq B$ *and* $B'' \leq B'$.

8.3. Our proof of 8.2 will be self-contained and elementary. Lemmas 8.4–8.6 below are needed for the proof.

Since $\Lambda \subset \mathbb{R}$, any ordered basis $\{\alpha, \beta\}$ of Λ determines a vector $v = \binom{\alpha}{\beta} \in \mathbb{R}^2$.

For any integer n we shall let $Q(n)$ denote the matrix $\binom{n\ 1}{1\ 0} \in \mathrm{GL}_2(\mathbb{Z})$.

If $B = \{\alpha > \beta\}$ is any element of $\mathscr{B}$, there is a unique positive integer n such that $0 < \alpha - n\beta < \beta$. If we set $\alpha' = \beta$ and $\beta' = \alpha - n\beta$, then we have $B' = \{\alpha' > \beta'\} \in \mathscr{B}$, and it is clear from the definition of the partial order on $\mathscr{B}$ that $B' \leq B$. We call B' the *canonical descendant* of B. The vectors $v = \binom{\alpha}{\beta}$ and $v' = \binom{\alpha'}{\beta'}$ satisfy $v = Q(n) \cdot v'$.

On the other hand, if $B' = \{\alpha' > \beta'\} \in \mathscr{B}$ is given, and if n is any positive integer, and if we set $v' = \binom{\alpha'}{\beta'}$ and $v = \binom{\alpha}{\beta} = Q(n) \cdot v'$, then it is clear that $B = \{\alpha > \beta\} \in \mathscr{B}$ and that B' is the canonical descendant of B.

8.4. The following result is somewhat more subtle.

Lemma. *Let* $B = \{\alpha > \beta\} \in \mathscr{B}$ *and* $B' = \{\alpha' > \beta'\} \in \mathscr{B}$ *be given. Set* $v = \binom{\alpha}{\beta}$ *and* $v' = \binom{\alpha'}{\beta'}$, *and suppose that* $v = Q(n) \cdot v'$ *for some* $n \in \mathbb{Z}$. *Then* $n > 0$, *and hence* v' *is the canonical descendant of* v.

Proof. Setting $x = \alpha/\beta$ and $x' = \alpha'/\beta'$, we have $x = n + (1/x')$. Since x and x' are > 1 we must have $n > 0$. q.e.d.

The next two lemmas are elementary facts about $GL_2(\mathbb{Z})$, and do not involve the order structure of Λ.

8.5. Lemma. *Let $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ be an element of $GL_2(\mathbb{Z})$ such that $c > d > 0$. Then M may be written in the form $Q(n_1) \cdots Q(n_k)$, where k is a nonnegative integer, $n_1, \cdots, n_k$ are integers, and $n_i > 0$ for $1 < i \leq k$.*

Proof. We argue by induction on the positive integer d. We first consider the case $d = 1$. Since $\det M = \pm 1$ we have $a = bc \pm 1$. If $a = bc + 1$ then $M = Q(b) \cdot Q(c)$, and if $a - bc = -1$ then $M = Q(b - 1) \cdot Q(1) \cdot Q(c - 1)$. Furthermore, we have $c > 1$ by hypothesis. Thus the lemma is verified in this case.

Now suppose that $d > 1$. Let us write $c = nd + r$, where n is a positive integer and $0 < r < d$. Then $M = M' \cdot Q(n)$, where M' is the matrix $\begin{pmatrix} b & a-nb \\ d & r \end{pmatrix}$. By the induction hypothesis we may write M' in the form $Q(n_1) \cdots Q(n_{k'})$, where $n_i > 0$ for $1 < i \leq k'$. We thus obtain the conclusion by setting $k = k' + 1$ and $n_k = n$. q.e.d.

8.6. Lemma. *Let M be any element of $GL_2(\mathbb{Z})$, and let $(n_i)_{i \geq 0}$ be a sequence of (strictly) positive integers. For each $i \geq 0$ set $M_i = M \cdot Q(n_1) \cdots Q(n_i)$, and write*

$$M_i = \begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix}.$$

Then we have either $c_i > d_i > 0$ for all sufficiently large i, or $c_i < d_i < 0$ for all sufficiently large i.

Proof. For each i we have $c_{i+1} = n_i c_i + d_i$ and $d_{i+1} = c_i$. It follows that if for a given i the integers c_i and d_i are both positive, then $c_{i+1} > d_{i+1} = 0$. Likewise, if c_i and d_i are both negative, then $c_{i+1} < d_{i+1} < 0$. Hence to prove the lemma we need only show that for some i the integers c_i and d_i are both positive or both negative.

We first consider the case that for some i we have $d_i = 0$. Since $M_i \in GL_2(\mathbb{Z})$ it follows that $c_i = \pm 1$. Since $n_i > 0$, the integers $c_{i+1} = n_i c_i$ and $d_{i+1} = c_i$ are either both positive or both negative.

Next suppose that for some i we have $c_i = 0$. Then $d_{i+1} = 0$, and hence by the argument just given, c_{i+2} and d_{i+2} are both positive or both negative.

Thus we may assume that $c_i \neq 0 \neq d_i$ for every i. Hence $x_i = c_i/d_i$ is always a nonzero rational number. We must show that $x_i > 0$ for some i.

For each i we have $x_{i+1} = n_i + x_i^{-1}$. If for some i we have $|x_i| \geq 1$, then $x_{i+1} \geq n_i - 1 \geq 0$; since $x_{i+1} \neq 0$ we must have $x_{i+1} > 0$, and we are done. It remains to consider the case that $|x_i| < 1$ for every i. In this case we have $|c_i| < |d_i| = |c_{i-1}|$ for every $i \geq 1$. This means that $(|c_i|)_{i \geq 1}$ is a strictly monotone decreasing sequence of positive integers. Hence this case cannot occur.

8.7.   *Proof of* 8.2.   Let us write $B = \{\gamma > \delta\}$ and $B' = \{\alpha > \beta\}$. Set $v = \binom{\alpha}{\beta}$ and $w = \binom{\gamma}{\delta} \in \mathbb{R}^2$. Since B and B' are both bases for Λ, there is a matrix $M \in \mathrm{GL}_2(\mathbb{Z})$ such that $M \cdot w = v$.

We recursively define a sequence $B_0, B_1, \ldots$ of elements of $\mathscr{B}$ by setting $B_0 = B$, and defining B_{i+1} to be the canonical descendant of B_i. Thus $B_i \leq B$ for every $i \geq 0$. We write $B_i = \{\gamma_i > \delta_i\}$.

For each $i \geq 0$ we set $w_i = \binom{\gamma_i}{\delta_i} \in \mathbb{R}^2$. Thus $w_0 = w$, and according to 8.3, for each $i \geq 0$ there is a positive integer n_{i+1} such that $w_i = Q(n_{i+1}) \cdot w_{i+1}$. Hence for each i we have $v = M \cdot w = M_i \cdot w_i$, where $M_i = M \cdot Q(n_1) \cdots Q(n_i)$.

Let us write

$$M_i = \begin{pmatrix} a_i & b_i \\ c_i & d_i \end{pmatrix}.$$

Then by 8.6 we have either $c_i > d_i > 0$ for all sufficiently large i, or $c_i < d_i < 0$ for all sufficiently large i. Since $M_i \cdot w_i = v$, and since v and w_i have positive entries, the alternative $c_i < d_i < 0$ cannot occur. Thus we may fix an index l such that $c_l > d_l > 0$. We shall complete the proof of 8.2 by showing that $B_l \leq B$.

Set $P = M_l$. By 8.5 we may write $P = Q(n'_1) \cdots Q(n'_k)$, where the n'_i are integers and $n'_2, \cdots, n'_k$ are positive. For $i = 0, \cdots, k$ we set $M'_i = Q(n'_1) \cdots Q(n'_i)$, so that M'_0 is the identity and $M'_k = P$. Since $M'_i \in \mathrm{GL}_2(\mathbb{Z})$, there is a unique vector

$$w'_i = \begin{pmatrix} \gamma'_i \\ \delta'_i \end{pmatrix} \in \mathbb{R}^2$$

such that $v = M'_i \cdot w'_i$. Thus $w'_0 = v$, and $w'_i = Q(n'_{i+1}) \cdot w'_{i+1}$ for $0 \leq i < k$. Since $v = M_l \cdot w_l = P \cdot w_l$, we have $w'_k = w_l$, i.e., $\gamma'_k = \gamma_l$ and $\delta'_k = \delta_l$. In particular we have $0 < \delta'_k < \gamma'_k$.

For $0 \leq i \leq k$, we have a basis $B'_i = \{\gamma'_i, \delta'_i\}$ of Λ. Since $w'_i = Q(n'_{i+1}) \cdot w'_{i+1}$, and since $n'_{i+1} > 0$ for $1 \leq i < k$, it follows from 8.3, using induction on $k - i$, that $0 < \delta'_i < \gamma'_i$ for $1 \leq i \leq k$, and that B'_{i+1} is the canonical descendant of B'_i whenever $1 \leq i < k$.

On the other hand, by the hypothesis of the proposition, $B_0' = B'$ is also a positive basis, and $\gamma_0' = \alpha > \beta = \delta_0'$. Hence by Lemma 8.4 we have $n_1 > 0$, and B_1' is the canonical descendant of B_0'. It follows that $B_k' \leq B_0' = B'$, as required.

8.8. We conclude this brief section with an easy lemma that will also be needed in §9.

Lemma. *Let L be a rank-2 free abelian subgroup of $\mathbb{R}$. Let α be a positive unimodular element of L. Then there is a unique element β of L such that* (i) $0 < \beta < \alpha/2$ *and* (ii) α *and* β *form a basis of L.*

Proof. If β_0 is an element of L such that α and β_0 form a basis, then any β for which α and β form a basis has the form $\varepsilon\beta_0 + n\alpha$ for some $n \in \mathbb{Z}$ and $\varepsilon = \pm 1$. There is a unique way of choosing n and ε so that $0 < \varepsilon\beta_0 + n\alpha < \alpha/2$.

9. A natural T-complex for the rank-2 case

In this section Λ will denote a subgroup of $\mathbb{R}$, and T will denote a Λ-tree. From 9.15 onward it will be assumed that Λ has $\mathbb{Q}$-rank 2.

9.1. Let L be a rank-2 free abelian subgroup of Λ. We shall associate to L a natural abstract T-complex (3.6) which we denote $(\mathfrak{X}(L), \Phi^L)$.

The complex $\mathfrak{X}(L)$ will have dimension at most 2, and its 1-skeleton will be the complex $U(L)$ defined in 3.10. Thus a vertex of $\mathfrak{X}$ is a point of T, and a 1-simplex is a two-element subset $\{x, y\}$ of T such that $d(x, y)$ is a unimodular element of L.

It remains to define the 2-simplices of $\mathfrak{X}(L)$. A three-element subset of T is called a 2-simplex if it has the form $\{x, y, z\}$, where (i) $y \in [x, z]$ and (ii) the distances $d(x, y)$ and $d(y, z)$ form a basis of L. (It follows that $d(x, y)$ and $d(x, z)$ also form a basis of L, as do $d(x, z)$ and $d(y, z)$. In particular, $\{x, y\}$, $\{x, z\}$ and $\{y, z\}$ are 1-simplices. Thus $\mathfrak{X}(L)$ is indeed an abstract simplicial complex.)

We define Φ^L to be the identity map of T. It is clear that $(\mathfrak{X}(L), \Phi^L)$ is a natural abstract T-complex. Thus the geometric realization of $(\mathfrak{X}(L), \Phi^L)$ is a natural geometric T-complex $(X(L), \phi^L)$.

Recall that by 7.5, $X(L)$ has a natural length system l, and l defines a decomposition of each closed 2-simplex of $X(L)$. By the definition of $X(L)$, the function l takes its values in L.

9.2. The complex $X(L)$ has the following fundamental property.

Proposition. *For any 1-simplex τ of $X(L)$, there are exactly two 2-simplices of which τ is the longest edge.*

Proof. Let τ be the 1-simplex $\Delta(x, z)$, let $\alpha = d(x, z)$ be the length of τ, and let h denote the unique isometry of $[0, \alpha] \subset \Lambda$ onto the segment $[x, y] \subset T$ which maps 0 to x and α to y. A 2-simplex δ has τ as its longest edge if and only if $\delta = \Delta(x, y, z)$, where (i) $y \in [x, z]$ and (ii) α and $d(x, y)$ form a basis of L. The only possible choices of y are $h(\beta)$ and $h(\alpha - \beta)$, where β is given by Lemma 8.8.   q.e.d.

Note that the assumption that L has rank 2 is crucial in the above proof.

In §10, Proposition 9.2 will be the starting point for a construction for producing subcomplexes of the complexes $X(L)$, for free abelian subgroups L of Λ, that are singular surfaces. The singular surfaces appearing in the statement of Theorem 6.2 will be constructed as covering spaces of these singular surfaces in the complexes $X(L)$.

The main goal of this section is to develop analogues of 3.12, 4.3 and 4.5 for the complex $X(L)$. These results depend on a number-theoretic investigation of homotopies between 1-tame lifts to $X(L)$ of parametrized segments in T.

In 9.3–9.10, L will denote a fixed rank-2 abelian subgroup of Λ. We shall write $(X, \phi) = (X(L), \phi^L)$.

9.3.   Let γ be a parametrized segment in $\mathbb{R}T$. Since $U = \mathrm{sk}_1(X)$, a 1-tame lift of γ to X is the same thing as a 1-tame lift of γ to U. In particular, it follows from 3.11 that γ admits a 1-tame lift to X if and only if the endpoints x and y of γ belong to T and $d(x, y) \in L$. On the other hand, homotopy (or T-homotopy) of 1-tame lifts in X is a weaker notion than homotopy (or T-homotopy) in U.

Proposition.   *Any two 1-tame lifts of γ to X are tamely T-homotopic* (3.9) *in X.*

9.4.   The next four lemmas are needed for the proof of Proposition 9.3. In these lemmas it is understood that L is a rank-two free abelian subgroup of Λ.

Let $\delta = \Delta(x, y, z)$ be a 2-simplex of X whose longest edge is $\Delta(x, z)$. The definition of X implies that ϕ maps each of the topological intervals $\Delta(x, z)$ and $\Delta(x, y) \cup \Delta(y, z)$ homeomorphically onto $[x, z]$. Hence a parametrized segment in T joining x to z admits exactly two 1-tame lifts to X whose images are contained in the boundary of δ.

9.5. Lemma.   *Let $\delta = \Delta(x, y, z)$ be a 2-simplex of X whose longest edge is $\Delta(x, z)$. Let γ be a parametrized segment in T joining x to z. Then the two tame lifts of γ whose images are contained in the boundary of δ are tamely T-homotopic.*

Proof. Set $\alpha = d(x, y)$ and $\beta = d(y, z)$. Then the edges $\tau = \Delta(x, z)$, $\tau' = \Delta(x, y)$ and $\tau'' = \Delta(y, z)$ of δ have respectively lengths α, β, and $\alpha + \beta$.

Write $\hat{\gamma}$ for the lift of γ with image τ, and $\hat{\gamma}'$ for the lift with image $\tau' \cup \tau''$. If $t \in [0, 1]$ is given, $\hat{\gamma}(t)$ is the unique point of τ which is mapped by ϕ to $\gamma(t)$. Likewise, $\hat{\gamma}'(t)$ is the unique point of $\tau' \cup \tau''$ which is mapped by ϕ to $\gamma(t)$. Thus $\hat{\gamma}(t)$ and $\hat{\gamma}'(t)$ are the endpoints of a line segment λ_t in the decomposition of δ defined by l.

We define a homotopy $\Xi : [0, 1] \times [0, 1] \to \delta$ as follows. For each $t \in [0, 1]$ we define $\Xi|\{t\} \times [0, 1]$ to be the unique affine map of $\{t\} \times [0, 1]$ onto λ_t that maps $(t, 0)$ to $\hat{\gamma}(t)$ and $(t, 1)$ to $\hat{\gamma}'(t)$. We have $\Xi_0 = \hat{\gamma}$, $\Xi_1 = \hat{\gamma}'$, and $\phi \circ \Xi_t = \gamma$ for all $t \in [0, 1]$. Furthermore, $\Xi([0, 1] \times [0, 1])$ is the closed simplex δ, a subcomplex of $X(L)$.   q.e.d.

9.6. As in §8, we denote the partially ordered set of all positive bases of L by $\mathfrak{B}(L)$.

Let γ be a parametrized segment in $\mathbb{R}T$, with endpoints x and y. Recall that γ admits a 1-tame lift to U—or, equivalently, a 1-tame lift to X—if and only if x, $y \in T$ and $d(x, y) \in L$. If these conditions do hold, then by 3.11, the lifts of γ are in 1-1 correspondence with the ordered partitions of $d(x, y)$ into positive unimodular elements of L. Recall how the correspondence is defined: if $\hat{\gamma}$ is a lift of γ, let $0 = a_0 < \cdots < a_n = 1$ be the points of $[0, 1]$ that are mapped by $\hat{\gamma}$ to vertices of X. Set $x_i = \hat{\gamma}(a_i)$ and $\alpha_i = d(x_{i-1}, x_i)$. Then $d(x, y) = \alpha_1 + \cdots + \alpha_n$ is the partition corresponding to $\hat{\gamma}$.

If the terms in the partition corresponding to a given lift $\hat{\gamma}$ of γ are all elements of a fixed basis $B \in \mathfrak{B}(L)$, we shall say that $\hat{\gamma}$ is *defined in terms of* B.

9.7. Lemma. *Let γ be an parametrized segment in T, let $\hat{\gamma}$ be a tame lift of γ to X which is defined in terms of a basis $B \in \mathfrak{B}(L)$, and let B' be an element of $\mathfrak{B}(L)$ such that $B' \leq B$. Then $\hat{\gamma}$ is tamely T-homotopic to a subdivision of $[x, y]$ defined in terms of B'.*

Proof. It is enough to prove the lemma in the case where B' is the successor (8.1) of B. Write $B = \{\alpha, \beta\}$ where $\alpha > \beta$; thus $B' = \{\alpha - \beta, \beta\}$.

Define a_i and x_i as in 9.6. Then $\alpha_i = d(x_{i-1}, x_i) \in B$ for $1 \leq i \leq n$. We need only show that, for $i = 1, \cdots, n$, the lift $\hat{\gamma}|[a_{i-1}, a_i]$ of $\gamma|[a_{i-1}, a_i]$ is tamely T-homotopic to a lift of $\gamma|[a_{i-1}, a_i]$ defined in terms of B'.

If $d(x_{i-1}, x_i) = \beta$, then $\gamma|[a_{i-1}, a_i]$ is itself defined in terms of B'. Now suppose that $d(x_{i-1}, x_i) = \alpha$. Let x' denote the unique point

of $\mathbb{R}T$ such that $d(x_{i-1}, x') = \beta$ and $d(x', x_i) = \alpha - \beta$. By definition, $\delta = \Delta(x_{i-1}, x', x_i)$ is a 2-simplex of X. Hence by Lemma 9.5, $\hat{\gamma}|[a_{i-1}, a_i]$ is tamely T-homotopic to the lift of $\gamma|[a_{i-1}, a_i]$ whose image is $\Delta(x_{i-1}, x') \cup \Delta(x', x_i)$. But this lift corresponds to the partition $\alpha = \beta + (\alpha - \beta)$ and is therefore defined in terms of B'.

9.8.   Let γ be a parametrized segment in T with endpoints x and y, and let $\hat{\gamma}$ be a 1-tame lift of γ to X. Let $d(x, y) = \alpha_1 + \cdots + \alpha_n$ be the ordered partition corresponding to $\hat{\gamma}$. A basis $B \in \mathfrak{B}(L)$ will be called a *mesh* for the lift $\hat{\gamma}$ if there exist $B_1, \cdots, B_n \in \mathfrak{B}(L)$ such that $\alpha_i \in B_i$ and $B \le B_i$ for $i = 1, \cdots, n$. It follows from 8.2 that every 1-tame lift has a mesh.

Lemma.   *Let $\hat{\gamma}$ be a tame lift to X of a parametrized segment γ in T. If B is a mesh for $\hat{\gamma}$ then $\hat{\gamma}$ is tamely T-homotopic to some tame lift of γ which is defined in terms of B.*

Proof.   Define a_i and x_i as in 9.6. It is enough to prove that for $1 \le i \le n$, the lift $\hat{\gamma}|[a_{i-1}, a_i]$ is tamely T-homotopic to a lift of $\gamma|[a_{i-1}, a_i]$ defined in terms of B. Choose $B_i \in \mathfrak{B}(L)$ such that $\alpha_i = d(x_{i-1}, x_i) \in B_i$ and $B \le B_i$. In particular the lift $\hat{\gamma}|[a_{i-1}, a_i]$ is defined in terms of B_i. Hence it follows from 9.7 that $\hat{\gamma}|[a_{i-1}, a_i]$ is tamely T-homotopic to a lift defined in terms of B.

9.9.   Lemma.   *Let γ be a parametrized segment in T, and let $\hat{\gamma}$ and $\hat{\gamma}'$ be two lifts of γ defined in terms of the same basis $B \in \mathfrak{B}(L)$. Then $\hat{\gamma}$ and $\hat{\gamma}'$ are tamely T-homotopic.*

Proof.   As in 9.6, let $0 = a_0 < \cdots < a_n = 1$ be the points of $[0, 1]$ that are mapped by $\hat{\gamma}$ to vertices of X. Set $x_i = \hat{\gamma}(a_i)$ and $\alpha_i = d(x_{i-1}, x_i)$. Similarly, let $0 = b_0 < \cdots < b_{n'} = 1$ be the points of $[0, 1]$ that are mapped by $\hat{\gamma}'$ to vertices of X; set $y_i = \hat{\gamma}'(b_i)$ and $\beta_i = d(y_{i-1}, y_i)$. Then the α_i and β_j belong to B, and $d(x, y) = \alpha_1 + \cdots + \alpha_n = \beta_1 + \cdots + \beta_{n'}$.

Since B is a basis we must have $n = n'$ and $\beta_j = \alpha_{\pi(j)}$ for $1 \le j \le n$, where π is some permutation of $\{1, \cdots, n\}$. In proving that $\hat{\gamma}$ and $\hat{\gamma}'$ are tamely T-homotopic, we may assume that π is a transposition of the form $(k, k + 1)$ with $1 \le k \le n$, since transpositions of this form generate the symmetric group on n letters. In this case, for $i \ne k$ we have $a_i = b_i$ and $x_i = y_i$, and we need only prove that the lifts $\hat{\gamma}|[a_{i-1}, a_{i+1}]$ and $\hat{\gamma}'|[a_{i-1}, a_{i+1}]$ of $\gamma|[a_{i-1}, a_{i+1}]$ are tamely T-homotopic.

If $\alpha_i = \alpha_{i+1}$, then $x_k = y_k$, and the two lifts coincide. If $\alpha_i \ne \alpha_{i+1}$ then $B = \{\alpha_i, \alpha_{i+1}\}$, and therefore $\Delta(x_{k-1}, x_k, x_{k+1})$ and $\Delta(x_{k-1}, y_k, x_{k+1})$ are 2-simplices of X. Thus by 9.5, the lifts $\hat{\gamma}|[a_{i-1}, a_{i+1}]$

and $\hat{\gamma}'|[a_{i-1}, a_{i+1}]$ are both tamely T-homotopic to the lift of γ whose image is $\Delta(x_{k-1}, x_{k+1})$.

9.10. *Proof of Proposition* 9.3. Let $\hat{\gamma}$ and $\hat{\gamma}'$ be two 1-tame lifts of γ. It follows from 8.2 that $\hat{\gamma}$ and $\hat{\gamma}'$ have a common mesh B. Hence by 9.8, they are tamely T-homotopic to 1-tame lifts $\hat{\gamma}_1$ and $\hat{\gamma}'_1$ defined in terms of B. By 9.9, the paths $\hat{\gamma}_1$ and $\hat{\gamma}'_1$ are tamely T-homotopic. q.e.d.

9.11. In the above discussion, L has been a fixed rank-two free abelian subgroup of Λ. Now let L and L' be rank-two free abelian subgroups of Λ with $L \subset L'$. Recall from 3.12 that there is a canonical T-map $H = H_{L,L'}$ from $U(L) = \mathrm{sk}_1(X(L))$ to $U(L') = \mathrm{sk}_1(X(L'))$, and that H is $\mathrm{Aut}(T)$-equivariant.

Proposition. *The map* $H_{L,L'}$ *may be extended to a 2-tame* $\mathrm{Aut}(T)$*-equivariant* T-map $J: X(L) \to X(L')$.

The next two lemmas are needed for the proof of this proposition. We shall write $H = H_{L,L'}$.

9.12. Lemma. *Let* δ *be a closed 2-simplex of* $X(L)$. *The map* $H|\partial\delta$ *may be extended to a continuous map* $J^\delta: \delta \to X(L')$ *such that* $\phi^{L'} \circ J^\delta = \phi^L|\delta$, *and such that* $J^\delta(\delta)$ *is a subcomplex of* $X(L')$.

Proof. By definition, δ has the form $\Delta(x, y, z)$, where $y \in [x, z]$. The longest edge of δ is $\tau = \Delta(x, z)$; the others are $\tau' = \Delta(x, y)$ and $\tau'' = \Delta(y, z)$. Let γ be any homeomorphism of $[0, 1]$ onto $|\tau| = [x, z]_{\mathbb{R}T}$. Then γ has a unique 1-tame lift $\hat{\gamma}_1$ to $X(L)$ whose image is τ, and a unique 1-tame lift $\hat{\gamma}_2$ whose image is $\tau' \cup \tau''$. Since H is a 1-tame T-map, $H \circ \hat{\gamma}_1$ and $H \circ \gamma_2$ are 1-tame lifts of γ to $X(L')$.

By Proposition 9.3 there is a tame T-homotopy Ω from $\hat{\gamma}$ to $\hat{\gamma}'$. We shall use Ω to construct the map J^δ.

For any $t \in [0, 1]$, the points $\hat{\gamma}_1(t)$ and $\hat{\gamma}_2(t)$ are the endpoints of a line segment λ_t in the decomposition of δ defined by l. There is a unique affine map $I_t: [0, 1] \to \lambda_t$ such that $I_t(0) = \hat{\gamma}_1(t)$ and $I_t(1) = \gamma_2(t)$. Note that I_t is an affine isomorphism for $t \neq 0, 1$.

For any point $p \in \delta$, there is a unique $t = t(p) \in [01,]$ such that $p \in \lambda_t$. For $p \neq x, z$ we have $t(p) \neq 0, 1$. For all $p \in \delta$ we define $J^\delta(p) = \Omega(t(p), I_{t(p)}^{-1}(p))$. Since the T-homotopy Ω is by definition constant on 0 and 1, the map J^δ is well defined. The definition of a T-homotopy also implies that $\phi^{L'} \circ J^\delta = \phi^L|\delta$. Furthermore, J^δ agrees with H on $\partial\delta$ since $\Omega_0 = \hat{\gamma}$ and $\Omega_1 = \hat{\gamma}'$. Since Ω is a tame homotopy, $J^\delta(\delta)$ is a subcomplex of $X(L')$. q.e.d.

Note that one could conclude immediately from Lemma 9.12 that $H_{L,L'}$ extends to a 2-tame T-map from $X(L)$ to $X(L')$. However, property (iii) is not immediate because the construction of J^δ in the proof of Lemma 9.12 is not canonical. (Indeed, the proof of 9.3 does not provide a canonical homotopy Ω.) The next lemma will allow us to deal with this point.

9.13.   Lemma.   *Let* $g \in \operatorname{Aut}(T)$ *be given, and let* δ *be a 2-simplex of* $X(L)$ *which is invariant under* g. *Then* g *fixes each point of* δ.

Proof.   Let $\delta = \Delta(x,y,z)$, where $\Delta(x,z)$ is the longest edge. Clearly $\Delta(x,z)$ must be invariant under g, and hence $g \cdot y = y$. If g were to interchange x and z, we would have $d(x,y) = d(y,z)$; this is impossible, since $d(x,y)$ and $d(y,z)$ form a basis of L by the definition of $X(L)$.

9.14.   *Proof of Proposition* 9.11.   We choose a set Σ of orbit representatives for the action of $\operatorname{Aut}(T)$ on the set of 2-simplices of X. (Thus each 2-simplex has the form $g \cdot \delta$ for a unique $\delta \in \Sigma$ and for some— not necessarily unique— $g \in \operatorname{Aut}(T)$.) For each $\delta \in \Sigma$, using Lemma 9.12, we fix a map $J^\delta : \delta \to X(L')$, agreeing with H on $\partial\delta$, such that $\phi^{L'} \circ J^\delta = \phi^L | \delta$, and such that $J^\delta(\delta)$ is a subcomplex of $X(L')$.

In order to extend H to a map $J: X(L) \to X(L')$, we note that any point of $X(L) - U(L)$ can be written in the form $g \cdot p$, where p is an interior point of some $\delta \in \Sigma$ and g is an element of $\operatorname{Aut}(T)$. We wish to define $J(g \cdot p) = g \cdot J^\delta(p)$. We must check that J is well defined.

Suppose that we are given elements g, g' of $\operatorname{Aut}(T)$ and interior points p, p' of δ, $\delta' \in \Sigma$, such that $g \cdot p = g' \cdot p'$. Then $g \cdot \delta = g' \cdot \delta'$, and hence $\delta = \delta'$. Lemma 9.13 implies that g and g' agree on δ. Hence $p = p'$. On the other hand, since $\phi^{L'} \circ J^\delta = \phi^L | \delta$, we have $\phi^{L'}(J^\delta(\delta)) = \phi^L(\delta) = |\delta|$. Thus the subcomplex $J^\delta(\delta)$ of $X(L')$ is a union of closed simplices of $X(L')$ whose supports are contained in $|\delta|$. In particular, g and g' agree on the vertices of these simplices, and therefore agree on $J^\delta(\delta)$. Hence $g \cdot J^\delta(p) = g' \cdot J^\delta(p) = g' \cdot J^{\delta'}(p')$. This shows that J is well defined. It is clear that J is a T-map. Since H is 1-tame and since $J^\delta(\delta)$ is a subcomplex of $X(L')$ for each $\delta \in \Sigma$, the map J is 2-tame. q.e.d.

9.15.   Now let us suppose that the Q-rank of Λ is 2. Then we can write $\Lambda = \bigcup_{i=1}^\infty L_i$, where L_i is a rank-2 free abelian subgroup of Λ and $L_i \subset L_{i+1}$. (The inclusions need not be proper.) We fix such a sequence (L_i), and for each $i \geq 1$ we set $(X_i, \phi_i) = (X(L_i), \phi^{L_i})$. Using Proposition 9.11, we fix an $\operatorname{Aut}(T)$-equivariant T-map $J_{i,i+1} : X_i \to X_{i+1}$.

If $i \leq j$ are positive integers we set $J_{ij} = J_{j-1,j} \circ \cdots \circ J_{i,i+1} : X_i \to X_j$. Thus we have a direct system $((X_i, \phi_i); J_{ij})$ of natural T-complexes and 2-tame $\text{Aut}(T)$-equivariant T-maps. We choose consistent base vertices for this system (see 4.1); we can then form the direct limit $\varinjlim \pi_n(X_i)$ for any $N \geq 0$.

Theorem. *For* $n = 0, 1$, *we have* $\varinjlim \pi_n(X_i) = 0$.

Proof. For $n = 0$ the result follows from 4.3, since $U_i = \text{sk}_i(X_i)$. For $n = 1$ we shall apply Proposition 4.4. To check condition (1) of 4.4, note that if γ is a parametrized segment in $\mathbb{R}T$ whose endpoints x and y lie in T, then $d(x, y) \in \Lambda$ and hence $d(x, y) \in L_i$ for some i; thus γ admits a 1-tame lift to X_i. Condition (2) of 4.4 holds by virtue of Proposition 9.3.

10. Edge-shortening

In this section we consider a Λ-tree T, where Λ is a subgroup of $\mathbb{R}$, a group Γ acting on T by isometries, and a rank-2 free abelian subgroup L of Λ. For any subset S of T, we denote the stabilizer of S in Γ by Γ_S. We shall suppose that the action of Γ on T satisfies the Ascending Chain Condition 6.1.

As in §9, we consider the geometric T-complex $(X, \phi) = (X(L), \phi^L)$. Recall that X is a natural T-complex, so that in particular Γ acts in a natural way on X and ϕ is Γ-equivariant.

If Y is any subcomplex of X, then $(Y, \phi|Y)$ is a geometric T-complex. In general when a subcomplex Y of X is given we shall refer to Y as a geometric T-complex, the map $\phi|Y$ being understood. In particular, if Y and Y' are subcomplexes of X, it makes sense to speak of a T-map from Y to Y'.

As in §9, we shall denote by l the natural L-valued length system on the 2-dimensional geometric T-complex X. If Y is any subcomplex of X then $(Y, \phi|Y)$ is also a geometric T-complex. Its natural length system is the restriction of l to the set of 1-simplices of Y; we denote it $l|Y$.

10.1. Proposition. *Let* Y *be a* Γ-*invariant subcomplex of* X *which is finite* $\text{mod}\,\Gamma$. *Then there exists a* Γ-*equivariant* T-*map* $\mathfrak{G} : Y \to X$ *such that*

(i) $\mathfrak{G}|\text{sk}_1(Y)$ *is homotopic rel* $\text{sk}_0(Y)$ *to the inclusion map* $\text{sk}_1(Y) \to X$;

(ii) $\Sigma = \mathfrak{G}(Y)$ *is a singular surface; and*

(iii) *the* L-*valued length system* $l|\Sigma$ *is nondegenerate* (7.2).

To interpret condition (i) note that since X is natural, any 0-tame T-map from Y to X must fix every vertex of Y.

The proof of this proposition will be given at the end of this section. Lemmas 10.3—10.9 involve the basic constructions used in the proof.

10.2. Let Y be any subcomplex Y of X which is Γ-invariant and finite $\mathrm{mod}\,\Gamma$, and which has dimension ≤ 2. We shall write $\mathscr{E}(Y)$ for the set of all 1-simplices of Y which are nonisolated, i.e., which are of strictly positive valence in Y. The number of Γ-orbits in $\mathscr{E}(Y)$ will be denoted $n(Y)$. Thus $n(Y) > 0$ if and only if Y has dimension 2.

The elements of $\mathscr{E}(Y)$ in a given orbit all have the same length. In particular, if Y is 2-dimensional, then $\mathscr{E}(Y)$ contains an element of maximal length; that is, there is some $\omega \in \mathscr{E}(Y)$ such that $l(\omega) \geq l(\tau)$ for all $\tau \in \mathscr{E}(Y)$. The set of all elements of maximal length in $\mathscr{E}(Y)$, which is a finite union of Γ-orbits, will be denoted $\mathscr{E}^*(Y)$.

10.3. Lemma. *Let* Y *be a* Γ-*invariant* 2-*dimensional subcomplex of* X *which is finite* $\mathrm{mod}\,\Gamma$. *Then every* 1-*simplex in* $\mathscr{E}^*(Y)$ *has valence* 1 *or* 2 *in* Y.

Proof. Let ω be a 1-simplex in $\mathscr{E}^*(Y)$. Since $\omega \in \mathscr{E}(Y)$ we have $\mathrm{val}_Y(\omega) > 0$. On the other hand, since ω has maximal length among the 1-simplices in $\mathscr{E}(Y)$, any 2-simplex of Y having ω as an edge has ω as its longest edge. By 9.2 this means that $\mathrm{val}_Y(\omega) \leq 2$.   q.e.d.

For $c = 1, 2$, we shall write $\mathscr{E}^*_c(Y)$ for the set of all elements of $\mathscr{E}^*(Y)$ that have valence c in Y. Thus $\mathscr{E}^*(Y) = \mathscr{E}^*_1(Y) \cup \mathscr{E}^*_2(Y)$.

Given a Γ-invariant subcomplex Y of X, finite $\mathrm{mod}\,\Gamma$, and a 1-simplex $\omega \in \mathscr{E}^*(Y)$, we shall see that there is a canonical Γ-equivariant T-map of Y into X determined by ω. Proposition 10.1 will be proved by composing maps of this canonical type. The construction of the canonical map depends on whether ω has valence 1 or 2. The case of valence 1 is relatively easy. It is contained in the following "Collapsing Lemma".

10.4. Let Z be a subcomplex of a T-complex Y. By a *T-deformation retraction* of Y onto Z we mean a retraction $c: Y \to Z$ which is also a T-map, and such that c, regarded as a map of Y into itself, is homotopic to the identity. Now let Y be a Γ-invariant subcomplex of X, and let Z be a subcomplex of Y. We say that Z is an *equivariant T-deformation retract* of Y if there exists a Γ-equivariant T-deformation retraction of Y onto Z. (Any equivariant T-retract of Y is in particular itself Γ-invariant.)

Collapsing Lemma. *Let* Y *be a* Γ-*invariant* 2-*dimensional subcomplex of* X, *and let* ω *be a* 1-*simplex in* $\mathscr{E}^*_1(Y)$. *Then there exists an equivariant*

T-deformation retract Z of Y, such that (i) Z contains all the vertices of Y, and (ii) $\mathscr{E}(Z) \subset \mathscr{E}(Y) - \Gamma \cdot \omega$. In particular we have $n(Z) < n(Y)$.

Proof. For each closed 2-simplex δ of X we shall define a natural continuous map $K^{\delta}: \delta \to \delta$. Let τ denote the long edge of δ, and let τ' and τ'' denote the other two edges. Any point $p \in \delta$ lies in a unique line segment λ in the decomposition of δ defined by l. We define $K^{\delta}(p)$ to be the unique endpoint of λ that lies in $\tau \cup \tau''$.

Since the closed simplex δ is a subcomplex of X, it may be regarded as a geometric T-complex; and K^{δ} is then a T-map. Note that K^{δ} is a retraction of δ onto $\tau \cup \tau''$.

We construct a map $c: Y \to Y$ as follows. On $\mathrm{sk}_0(Y)$ we define c to be the identity. Now let τ be any 1-simplex of Y. If τ is in the orbit of ω then $\tau \in \mathscr{E}^*(Y)$; thus τ is an edge of a unique 2-simplex δ, and is in fact the longest edge of δ. We define c to agree with K^{δ} on τ. If τ is not in the orbit of ω we define c to be the identity on τ. In either case c restricts to the identity on $\partial \tau$.

It remains to extend c to the 2-simplices of Y. Let δ be any 2-simplex of Y. Let τ denote the longest edge of δ, and let τ' and τ'' denote its other two edges. Then τ' and τ'' are not in $\mathscr{E}^*(Y)$, and hence are not in the orbit of ω. As for τ, it may or may not be in the orbit of ω.

If τ is in the orbit of ω we define $c|\delta$ to be K^{δ}. Since K^{δ} is a retraction of δ onto $\tau' \cup \tau''$, it is clear that this definition of c agrees with the earlier one on $\partial \delta$. If τ is not in the orbit of ω we define c to be the identity on δ. Again this agrees with the earlier definition on $\partial \delta$. Thus c is well defined on Y.

Since $K^{\delta}: \delta \to \delta$ is a T-map for each 2-simplex δ, one sees that c is a T-map. It is clear that c is a Γ-equivariant deformation retraction of Y onto Z, where Z is a subcomplex of Y having properties (i) and (ii) in the statement of the lemma. q.e.d.

10.5. Now let Y be a Γ-invariant 2-dimensional subcomplex of X, and let ω be a 1-simplex in $\mathscr{E}_2^*(Y)$. We shall describe here a Γ-invariant subcomplex Z of X, determined by Y, X and ω via a process we call "edge-shortening". Lemma 10.9 below will provide a Γ-equivariant T-map of Y onto Z.

According to 9.2, there are exactly two 2-simplices of X, say δ_0 and δ_1, of which ω is the longest edge. Since $\omega \in \mathscr{E}_2^*(Y)$, both δ_0 and δ_1 must be simplices of Y, and they must be the only 2-simplices of Y having ω as an edge.

Let x and z denote the vertices of ω. Then $\delta_0 = \Delta(x, y_0, z)$ and $\delta_1 = \Delta(x, y_1, z)$ for some vertices y_0 and y_1. We may regard x, y_0,

y_1 and z as points of T. Set $\alpha = d(x, z)$, and let h denote the unique isometry of $[0, \alpha] \cap \Lambda$ onto $[x, y]$ which maps 0 to x and α to y. According to the proof of 9.2, there exists $\beta \in \Lambda$, with $0 < \beta < \alpha/2$, such that $\{\alpha, \beta\}$ is a basis for Λ, and the points y_0 and y_1 are $h(\beta)$ and $h(\alpha - \beta)$ in some order. After relabelling δ_0 and δ_1 if necessary, we may assume that $y_0 = h(\beta)$ and that $y_1 = h(\alpha - \beta)$.

In the tree T we have $y_0 \in [x, y_1]$. Furthermore, $d(x, y_0) = \beta$ and $d(y_0, y_1) = (\alpha - \beta) - \beta = \alpha - 2\beta$. Since $\{\alpha, \beta\}$ is a basis for Λ, so is $\{\beta, \alpha - 2\beta\}$. Hence the definition of X implies that $\rho_- = \Delta(x, y_0, y_1)$ is a 2-simplex of X. Similarly, $\rho_+ = \Delta(y_0, y_1, z)$ is a 2-simplex of X. The 2-simplices ρ_- and ρ_+ have the common edge $\psi = \Delta(y_0, y_1)$. The 1-simplex ψ may or may not lie in the complex Y. However, the other edges of ρ_- and ρ_+ do lie in Y.

Since δ_0 and δ_1 are the only 2-simplices of Y having ω as an edge, the only 2-simplices having an edge in the orbit $\Gamma \cdot \omega$ are those in the orbits $\Gamma \cdot \delta_0$ and $\Gamma \cdot \delta_1$. Hence deleting the simplices in the set $\Gamma \cdot \omega \cup \Gamma \cdot \delta_0 \cup \Gamma \cdot \delta_1$ from Y gives a subcomplex W. Note that W contains all the 1-simplices of Y that have length $< \alpha$. Clearly W is Γ-invariant.

Now consider the set of 2-simplices $\Gamma \cdot \rho_- \cup \Gamma \cdot \rho_+$. All the edges of the simplices in this set either belong to the orbit $\Gamma \cdot \psi$, or are 1-simplices of Y which have length $< \alpha$ and are therefore in the subcomplex W. Hence we can form a new subcomplex Z of X by adjoining to Y the simplices of the set $\Gamma \cdot \psi \cup \Gamma \cdot \rho_- \cup \Gamma \cdot \rho_+$. Again, Z is clearly Γ-invariant.

The complex Z is said to be *obtained from Y via the edge-shortening determined by ω*. The 1-simplex ψ is called *the short edge replacing the long edge ω*. As we have mentioned, ψ may or may not belong to $\mathscr{E}(Y)$.

10.6. Note that $\mathscr{E}(Z) = (\mathscr{E}(Y) - \Gamma \cdot \omega) \cup \Gamma \cdot \psi$. Hence we have $n(Z) \leq n(Y)$, and the inequality is strict if and only if $\psi \in \mathscr{E}(Y)$.

Note that we have $l(\psi) = l(w) - 2l(\chi)$ for some $\chi \in \mathscr{E}(Y)$. (In fact we have $\chi = \Delta(x, y_0)$, where x and y_0 are defined as in 10.5.)

Note also that by the construction, the segment $|\psi|$ is a subsegment of $|\omega|$ which has the same midpoint in T as $|\omega|$.

The following three lemmas express further properties of the edge-shortening operation. In these lemmas, it is understood that Y is a Γ-invariant 2-dimensional subcomplex of X, and that $\omega \in \mathscr{E}_2^*(Y)$. Furthermore, Z denotes the subcomplex of X obtained from Y via the edge-shortening determined by ω, and ψ denotes the short edge replacing the long edge ω.

10.7. Lemma. *For every 1-simplex τ of Z that is not in the Γ-orbit of ψ we have $\mathrm{val}_Z(\tau) \leq \mathrm{val}_Y(\tau)$.*

Proof. We use the notation of 10.5.

Let $\mathscr{D}_Y$ (resp. $\mathscr{D}_Z$) denote the set of all 2-simplices of Y (resp. Z) that have τ as an edge. Thus $\mathrm{val}_Z(\tau)$ and $\mathrm{val}_Y(\tau)$ are the cardinalities of $\mathscr{D}_Y$ and $\mathscr{D}_Z$.

Consider first the case where the length of τ is not equal to β or $\alpha - \beta$. In this case τ is not in the orbit of any of the 1-simplices $\Delta(x, y_0), \Delta(y_0, z), \Delta(x, y_1)$ or $\Delta(y_1, z)$. By hypothesis it is not in the orbit of $\psi = \Delta(y_0, y_1)$. Hence by the definition of edge-shortening we have $\mathscr{D}_Z = \mathscr{D}_Y$ in this case, and hence $\mathrm{val}_Z(\tau) = \mathrm{val}_Y(\tau)$.

Next consider the case $l(\tau) = \beta$. In this case we shall see that $\mathrm{val}_Z(\tau) \leq \mathrm{val}_Y(\tau)$ by defining a surjective map $\mathscr{G} : \mathscr{D}_Y \to \mathscr{D}_Z$.

Let δ be any 2-simplex in $\mathscr{D}_Y$. If $\delta \notin \Gamma \cdot \delta_0 \cup \Gamma \cdot \delta_1$, we set $\mathscr{G}(\delta) = \delta$. Now suppose that $\delta \in \Gamma \cdot \delta_0 \cup \Gamma \cdot \delta_1$. We wish to set $\mathscr{G}(\delta) = g \cdot \rho_-$ if $\delta = g \cdot \delta_0$ with $g \in \Gamma$, and $\mathscr{G}(\delta) = g \cdot \rho_+$ if $\delta = g \cdot \delta_1$ with $g \in \Gamma$. It is necessary to check that $\mathscr{G}(\delta)$ is well defined; a priori, the definition of $\mathscr{G}(\delta)$ depends on the choice of g, and—since δ_0 and δ_1 may happen to lie in the same Γ-orbit—on the choice of i as well.

To show that $\mathscr{G}(\delta)$ is independent of the choice of g, we must show that if an element of Γ which leaves δ_0 (resp. δ_1) invariant also leaves ρ_- (resp. ρ_+) invariant. But if $g \in \Gamma$ leaves δ_0 invariant then the segment $[x, z] = |\delta_0| \subset T$ must be pointwise fixed by g. In particular, $[x, y_1] = |\rho_-| \subset T$ must be pointwise fixed by g, and hence g must leave ρ_- invariant. Similarly, if g leaves δ_1 invariant then it leaves ρ_+ invariant.

To show that $\mathscr{G}(\delta)$ is independent of the choice of i, we must show that if g is an element of Γ such that $g \cdot \delta_0 = \delta_1$, then $g \cdot \rho_- = \rho_+$. If $g \cdot \delta_0 = \delta_1$, then g must leave the segment $[x, z] = |\delta_0| = |\delta_1|$ invariant and map y_0 to y_1. Since $y_0 \neq y_1$, the automorphism g cannot fix the endpoints of $[x, z]$; it must therefore interchange x and z. Thus g induces an involution of $[x, z]$, and since $g \cdot y_0 = y_1$ it follows that $g \cdot y_1 = g \cdot y_0$. Hence g transforms $\rho_- = \Delta(x, y_0, y_1)$ onto $\Delta(z, y_1, y_0) = \rho_+$.

Thus $\mathscr{G}(\delta)$ is well defined for all $\delta \in \mathscr{D}_Y$. Furthermore, we have $\mathscr{G}(\delta) \in \mathscr{D}_Z$. This is obvious if $\delta \notin \Gamma \cdot \delta_0 \cup \gamma \cdot \delta_1$, since in this case $\mathscr{G}(\delta) = \delta$. If $\delta = g \cdot \delta_0$ for some $g \in \Gamma$, then by definition we have $\mathscr{G}(\delta) = g \cdot \rho_-$. Now since τ has length β, and since $\Delta(x, y_0)$ is the only edge of $\delta_0 = \Delta(x, y_0, z)$ whose length is β, we must have $\tau = g \cdot \Delta(x, y_0)$. Hence τ is an edge of $g \cdot \Delta(x, y_0, y_1) = \mathscr{G}(\delta)$, i.e., $\mathscr{G}(\delta) \in \mathscr{D}_Z$. A precisely similar argument shows that if $\delta = g \cdot \delta_0$ for some $g \in \Gamma$ then $\mathscr{G}(\delta) \in \mathscr{D}_Z$.

To see that $\mathcal{G}: \mathcal{D}_Y \to \mathcal{D}_Z$ is surjective, let δ be any 2-simplex in $\mathcal{D}_Z$. If $\delta \notin \Gamma \cdot \rho_- \cup \Gamma \cdot \rho_+$, then clearly $\delta \in \mathcal{D}_Y$ and $\delta \notin \Gamma \cdot \delta_0 \cup \Gamma \cdot \delta_1$. By the definition of $\mathcal{G}$ we have $\delta = \mathcal{G}(\delta) \in \mathcal{G}(\mathcal{D}_Y)$. If $\delta = g \cdot \rho_-$ for some $g \in \Gamma$, then since $l(\tau) = \beta$, and since $\Delta(x, y_0)$ is the only edge of $\rho_- = \Delta(x, y_0, y_1)$ whose length is β, we must have $\tau = g \cdot \Delta(x, y_0)$. Hence the 2-simplex $\delta' = g \cdot \delta_0$ belongs to $\mathcal{D}_Y$. By the definition of $\mathcal{G}$ we have $\delta = \mathcal{G}(\delta') \in \mathcal{G}(\mathcal{D}_Y)$. A precisely similar argument shows that if $\delta = g \cdot \rho_+$ for some $g \in \Gamma$ then $\delta \in \mathcal{G}(\mathcal{D}_Y)$. Thus $\mathcal{G}$ is surjective.

This proves that $\mathrm{val}_Z(\tau) \leq \mathrm{val}_Y(\tau)$ in the case where $l(\tau) = \beta$. In the remaining case, where $l(\tau) = \alpha - \beta$, we define a surjection $\mathcal{G}: \mathcal{D}_Y \to \mathcal{D}_Z$ as follows. Let $\delta \in \mathcal{D}_Y$ be given. If $\delta \notin \Gamma \cdot \delta_0 \cup \Gamma \cdot \delta_1$, we set $\mathcal{G}(\delta) = \delta$. If $\delta = g \cdot \delta_0$ for some $g \in \Gamma$ we set $\mathcal{G}(\delta) = g \cdot \rho_+$. If $\delta = g \cdot \delta_1$ we set $\mathcal{G}(\delta) = g \cdot \rho_-$. The proof that $\mathcal{G}$ is a well-defined surjection from $\mathcal{D}_Y$ to $\mathcal{D}_Z$ is precisely parallel to the previous argument for the case $l(\tau) = \beta$. q.e.d.

10.8. Lemma. *Suppose that* $\Gamma_{|\psi|} = \Gamma_{|\omega|}$ *and that* $n(Z) = n(Y)$. *Then* ψ *has valence* 2 *in* Z.

Proof. We use the notation of 10.5. The 2-simplices ρ_- and ρ_+ are distinct and both have ψ as an edge. To prove the lemma we must show that ρ_- and ρ_+ are the only 2-simplices in Z having ψ as an edge.

Let ρ be any 2-simplex in Z having ψ as an edge. Either ρ is a 2-simplex of Y or it is in the orbit of ρ_- or ρ_+. If ρ is a 2-simplex of Y then we have $\psi \in \mathcal{E}(Y)$. But by 10.6 this implies $n(Z) < n(Y)$, in contradiction to the hypothesis of the lemma. Hence ρ must lie in the orbit of ρ_- or ρ_+. By symmetry we may assume that ρ is in the orbit of ρ_-, so that $\rho = g \cdot \rho_-$ for some $g \in \Gamma$. Since ψ is the unique edge of ρ_- having length $\alpha - 2\beta$, we must have $g \cdot \psi = \psi$. Thus $g \in \Gamma_{|\psi|}$, and by hypothesis we therefore have $g \in \Gamma_{|\omega|}$. Thus g either fixes $|\omega|$ pointwise or interchanges the endpoints of $|\omega|$.

Suppose that g fixes $|\omega|$ pointwise. Since $|\omega| = [x, z] \supset [x, y_1] = |\rho_-|$, it follows that g fixes $|\rho_-|$ pointwise. Hence $\rho = g \cdot \rho_- = \rho_-$.

Now suppose that g interchanges the endpoints of $|\omega| = [x, z]$. Thus $g \cdot x = z$ and $g \cdot z = x$. Since $d(x, y_0) = \beta = d(z, y_1)$ we must have $g \cdot y_0 = y_1$ and $g \cdot y_1 = y_0$. Hence g transforms $\rho_- = \Delta(x, y_0, y_1)$ onto $\rho_+ = \Delta(z, y_1, y_0)$, so that in this case $\rho = \rho_+$. q.e.d.

10.9. Lemma. *There exists a* Γ-*equivariant* T-*map* $\mathfrak{G}_\omega: Y \to Z$. *Furthermore,* $\mathfrak{G}_\omega | \mathrm{sk}_1(Y)$ *is homotopic rel* $\mathrm{sk}_0(Y)$ *to the inclusion map* $\mathrm{sk}_1(Y) \to X$.

Proof. We use the notation of 10.5. The map $\mathfrak{G}_\omega$ will be constructed by extending the inclusion map $W \to Z$ to a map of Y into Z. This map

will be defined so as to be simplicial with respect to appropriate simplicial subdivisions of Y and Z.

For each 1-simplex τ of X we shall write m_τ for the midpoint of τ.

We define a subdivision Y' of the complex Y as follows. The vertices of Y' are the vertices of Y and the points m_τ for $\tau \in \Gamma \cdot \omega$. Every simplex of W is a simplex of Y'. Each 1-simplex $\tau \in \Gamma \cdot \omega$ is subdivided at m_τ into two 1-simplices of Y'. If δ is a 2-simplex in $\Gamma \cdot \delta_1 \cup \Gamma \cdot \delta_2$, and if $\tau \in \Gamma \cdot \omega$ is the longest edge of δ, then the segment in δ joining τ to the opposite vertex subdivides δ into two 2-simplices of Y'.

Similarly we can define a subdivision Z' of Z. The vertices of Z' are the vertices of Z and the points m_τ for $\tau \in \Gamma \cdot \psi$. Every simplex of W is a simplex of Z'. Each 1-simplex $\tau \in \Gamma \cdot \psi$ is subdivided at m_τ into two 1-simplices of Z'. If δ is a 2-simplex in $\Gamma \cdot \delta_1 \cup \Gamma \cdot \delta_2$, and if $\tau \in \Gamma \cdot \omega$ is the unique edge of δ having length $\alpha - 2\beta$, then the long segment in δ joining τ to the opposite vertex subdivides δ into two 2-simplices of Y'.

We define a map $\mathfrak{G}_0 \colon \mathrm{sk}_0(Y') \to \mathrm{sk}_0(Z')$ as follows. For any vertex v of Y we set $\mathfrak{G}_0(v) = v$. For any vertex v of Y' which is not a vertex of Y, we may write $v = g \cdot m_\omega$ for some $g \in \Gamma$. We wish to define $\mathfrak{G}_0(v) = g \cdot m_\psi$. In order to show that $\mathfrak{G}_0$ is well defined, we must check that if an element g of Γ fixes m_ω then it fixes m_ψ. But if g fixes m_ω, then since Γ acts simplicially on X, the 1-simplex ω must be invariant under g. Hence $|\omega| = [x, z] \subset T$ is invariant under g. Thus $|\psi| = [y_0, y_1]$, which is a subsegment of $[x, z]$ having the same midpoint as $[x, z]$, must also be invariant under g. Since X is a natural T-complex it follows that ψ is invariant under g; hence g fixes m_ψ. Thus $\mathfrak{G}_0$ is well defined.

It is clear from the definition that $\mathfrak{G}_0$ is Γ-equivariant.

We claim that $\mathfrak{G}_0$ can be extended in a unique way to a simplicial map $\mathfrak{G}_\omega \colon Y' \to Z'$. In order to show this we must show that for any simplex Δ of Y', the vertices of Δ are mapped by $\mathfrak{G}_0$ to the vertices of a simplex of Z'. If Δ is a simplex of W this is obvious. Thus we may assume that Δ is in the Γ-orbit of one of the simplices of Y' contained in δ_0 or δ_1. By the equivariance of S_0 we may in fact assume that Δ is itself contained in δ_i for $i = 0$ or 1. Hence Δ is a face of either $\delta_i^- = \Delta(x, y_i, m_\omega)$ or $\delta_i^+ = \{z, y_i, m_\omega\}$. But then $\mathfrak{G}_0$ maps the vertices of Δ into either $\{x, y_i, m_\psi\}$ or $\{z, y_i, m_\psi\}$; by definition, each of the latter sets spans a 2-simplex of Z'.

The map $\mathfrak{S}_\omega$ is clearly Γ-equivariant. It is also clear that $S_\omega|W$ is the identity; in particular, $\mathfrak{S}_\omega$ fixes the vertices of Y, so that we may regard it as a 0-tame map from Y to Z. To show that $\mathfrak{S}_\omega$ is a T-map, it suffices—in view of equivariance—to show that for $i = 0$, 1, the maps $\phi \circ \mathfrak{S}_\omega$ and ϕ agree on $\delta_i = \delta_i^- \cup \delta_i^+$ for $i = 0$, 1. To see this, note that $\phi \circ \mathfrak{S}_\omega|\delta_i^-$ and $\phi|\delta_i^-$ are affine maps of δ_i^- into $[x, z]$ that respectively map x, y_i and m_ω to x, y_i and m, where m denotes the midpoint of $[x, z]$. Hence $\phi \circ \mathfrak{S}_\omega|\delta_i^- = \phi|\delta_i^-$. Similarly we have $\phi \circ \mathfrak{S}_\omega|\delta_i^+ = \phi|\delta_i^+$.

Finally, we must show that $\mathfrak{S}_\omega|\operatorname{sk}_1(Y)$ is homotopic rel $\operatorname{sk}_0(Y)$ to the inclusion map $\operatorname{sk}_1(Y) \to X$. Since $\mathfrak{S}_\omega$ is Γ-equivariant and agrees with the identity on W, this amounts to showing that $\mathfrak{S}_\omega|\omega$ is homotopic rel the endpoints of ω to the inclusion of ω in X. But ω and $\mathfrak{S}_\omega(\omega)$ are both contained in the subcomplex $B = \delta_0 \cup \delta_1 \cup \rho_- \cup \rho_+$ of X. The complex B is isomorphic to the boundary of a 3-simplex and is therefore simply connected; thus any two paths in B having the same endpoints are homotopic. q.e.d.

For the remainder of this section the hypothesis of Proposition 10.1 will be understood to hold. Thus we are given a Γ-invariant subcomplex Y of X which is finite $\operatorname{mod}\Gamma$.

10.10 Lemma. *There exist a sequence* $(Y_i)_{i \geq 0}$ *of Γ-invariant subcomplexes of X, all finite $\operatorname{mod}\Gamma$, and a sequence $(\mathfrak{S}_i \colon Y_i \to Y_{i+1})_{i \geq 0}$ of Γ-equivariant T-maps, such that*

(i) $Y_0 = Y$;

(ii) $\mathfrak{S}_i|\operatorname{sk}_1(Y_i)$ *is homotopic rel $\operatorname{sk}_0(Y_i)$ to the inclusion map $\operatorname{sk}_1(Y_i) \to X$;*

(iii) $n(Y_i) = n(Y_{i+1})$ *for all sufficiently large i; and*

(iv) *for all sufficiently large i, either Y_i has dimension at most 1, or Y_{i+1} is obtained from Y_i via the edge-shortening determined by some $\omega \in \mathscr{E}_2^*(Y_i)$.*

Proof. We define the Y_i and $\mathfrak{S}_i$ recursively. We set $Y_0 = Y$. Now suppose that $i \geq 0$ and that Y_i has been defined and is a Γ-invariant subcomplex of X, finite $\operatorname{mod}\Gamma$. Then either Y_i has dimension at most 1 or $\mathscr{E}^*(Y) \neq \phi$. If $\dim Y_i \leq 1$ we set $Y_{i+1} = Y_i$ and define $\mathfrak{S}_i$ to be the identity map. Otherwise, we choose a 1-simplex $\omega_i \in \mathscr{E}^*(Y_i) = \mathscr{E}_1^*(Y_i) \cup \mathscr{E}_2^*(Y_i)$.

If $\omega_i \in \mathscr{E}_1^*(Y_i)$, then according to the Collapsing Lemma 10.4, there exists an equivariant T-deformation retract Y_{i+1} of Y_i, such that Y_{i+1} contains all the vertices of Y_i, and $n(Y_{i+1}) < n(Y_i)$. We take S_i to be a Γ-equivariant T-deformation retraction of Y_i onto Y_{i+1}.

If $\omega_i \in \mathscr{E}_2^*(Y_i)$, we define Y_{i+1} to be the subcomplex of X obtained from Y_i via the edge-shortening determined by ω_i. According to Lemma 10.9, there exists a 0-tame Γ-equivariant T-map $\mathfrak{G}_i$ of Y_i onto Y_{i+1}. Furthermore, $\mathfrak{G}_i | \mathrm{sk}_1(Y_i)$ is homotopic rel $\mathrm{sk}_0(Y_i)$ to the inclusion map $\mathrm{sk}_1(Y_i) \to X$.

This completes the construction of the Γ-invariant subcomplexes Y_i, finite mod Γ, and the Γ-equivariant T-maps $\mathfrak{G}_i$. Conditions (i) and (ii) of the lemma are immediate from the construction. If one of the Y_i, say Y_{i_0}, has dimension at most 1, it follows from the construction that for all $i \geq i_0$ we have $Y_i = Y_{i_0}$, and hence (iii) and (iv) hold.

Now suppose that all the Y_i are 2-dimensional. By the construction of the Y_i it follows that for each i, either $n(Y_{i+1}) < n(Y_i)$, or Y_{i+1} is obtained from Y_i via the edge-shortening determined by some $\omega_i \in \mathscr{E}_2^*(Y_i)$. Using 10.6 we conclude that $n(Y_{i+1}) \leq n(Y_i)$ for all $i \geq i_0$. Hence for sufficiently large i, say for $i \geq i_0$, we must have $n(Y_{i+1}) = n(Y_i)$. In particular, for all $i \geq i_0$, the complex Y_{i+1} is obtained from Y_i via the edge-shortening determined by a 1-simplex $\omega_i \in \mathscr{E}_2^*(Y_i)$. Thus conditions (iii) and (iv) hold in this case as well.

10.11.   For the rest of this section we fix a sequence of subcomplexes Y_i, and a sequence of maps $\mathfrak{G}_i$, having the properties stated in Lemma 10.10. In order to prove Proposition 10.1, it is sufficient to prove that for some $i \geq 0$, the complex Y_i is a singular surface and the length system $l | Y_i$ is nondegenerate.

A complex of dimension ≤ 1 is obviously a singular surface. Furthermore, any length system on such a complex is automatically nondegenerate. Hence we may assume that all the Y_i are 2-dimensional. According to condition (iv) of 10.10, this means that for all sufficiently large i, the complex Y_{i+1} is obtained from Y_i via the edge-shortening determined by some 1-simplex in $\mathscr{E}_2^*(Y_i)$. After passing to a subsequence we may assume that this is the case for all $i \geq 1$. Likewise, by virtue of condition (iii), we may assume that $n(Y_i)$ has a constant value n for all $i \geq 1$.

Let us fix representatives $\tau_1, \tau_2, \cdots, \tau_n$ of the n distinct Γ-orbits of 1-simplices in $\mathscr{E}(Y_1)$. In terms of the indexed family $(\tau_1, \tau_2, \cdots, \tau_n)$ we shall define recursively, for $i \geq 1$, an indexed family $(\tau_{1i}, \tau_{2i}, \cdots, \tau_{ni})$ of representatives of the n distinct Γ-orbits of 1-simplices in $\mathscr{E}(Y_i)$.

For $j = 1, \cdots, n$, we set $\tau_{j1} = \tau_j$. Now suppose that $i \geq 1$ and that the τ_{ji} have been defined. Since the τ_{ji} form a complete system of representatives of the Γ-orbits of 1-simplices of Y_i, one of them, say $\omega_i = \tau_{j(i),i}$, is a valence-2 maximal nonisolated 1-simplex of Y_i, such that

Y_{i+1} is obtained from Y_i via the edge-shortening determined by ω_i. Let ψ_{i+1} denote the short edge in Y_{i+1} replacing the long edge τ_{ji} (see 10.5).

We set $\tau_{j(i), i+1} = \psi_{i+1}$, and $\tau_{j, i+1} = \tau_{j, i}$ for $j \neq j(i)$. The definition of edge-shortening implies that every Γ-orbit of 1-simplices of Y_{i+1} contains one of the simplices $\tau_{1, i+1}, \cdots, \tau_{n, i+1}$. Since $n(Y_{i+1}) = n$ it follows that $\tau_{1, i+1}, \cdots, \tau_{n, i+1}$ represent distinct Γ-orbits.

For $i \geq 1$ and $1 \leq j \leq n$ we set $l_{ji} = l(\tau_{ji})$. In particular we have $l_{j(i), i} = l(\omega_i)$ and $l_{j(i), i+1} = l(\psi_{i+1})$.

10.12. Lemma. *For each j with $1 \leq j \leq n$, we have $\lim_{i \to \infty} l_{ji} = 0$, Furthermore, for each j there are infinitely many $i \geq 1$ for which $j(i) = j$. For all such values of i, the 1-simplex τ_{ji} has valence 2 in Y_i.*

Proof. Let $i \geq 1$ be given. For all $j \neq j(i)$, it follows from the definition of the τ_{ji} that $\tau_{j, i+1} = \tau_{ji}$, and hence that $l_{j, i+1} = l_{ji}$. On the other hand, according to 10.6, we have

(1) $l_{j(i), i+1} = l_{j(i), i} - 2l_{j'(i), i}$ for some $j'(i) \neq j(i)$.

In particular, it follows that

(2) $l_{j, i+1} \leq l_{ji}$ whenever $i \geq 1$ and $1 \leq j \leq n$.

Equality holds unless $j = j(i)$.

For each $i \geq 1$ let us set $m_i = \min_{1 \leq i \leq n} l_{ji}$ and $\mathcal{M}_i = \max_{1 \leq i \leq n} l_{ji}$. Then it follows from (1) that

(3) $l_{j(i), i+1} \leq l_{j(i), i} - 2m_i.$

Let us also observe that since $\omega_i \in \mathcal{E}^*(Y_i)$, we have

(4) $l_{j(i), i} = \mathcal{M}_i.$

If we set $S_i = \sum_{j=1}^n l_{ji}$, then (2) and (3) imply that

(5) $S_{i+1} \leq S_i - 2m_i.$

Since $S_i \geq 0$ for all i, it follows form (5) that $m_i \to 0$ as $i \to \infty$. This means that for some j we have $l_{ji} \to 0$ as $i \to \infty$. After reindexing $\tau_1, \cdots, \tau_n$ if necessary we may assume that $l_{1i} \to 0$ as $i \to \infty$. In particular there are infinitely many values of $i \geq 1$ for which $l_{1, i+1} < l_{1i}$. For any such i it follows from (2) that $j(i) = 1$, and hence from (4) that $l_{1i} = \mathcal{M}_i$.

Thus $\mathcal{M}_i$ takes arbitrarily small values as $i \to \infty$. On the other hand, it follows from (2) that $\mathcal{M}_i$ decreases monotonically. Hence $\lim_{i \to \infty} \mathcal{M}_i = 0$. This means that for each j we have $l_{ji} \to 0$ as $i \to \infty$. This is the first assertion of the lemma.

In particular, given any j with $1 \le j \le n$, there are infinitely many values of $i \ge 1$ for which $l_{j,i+1} < l_{ji}$. For any such i it follows from (2) that $j(i) = j$, which by the definition of $j(i)$ implies that τ_{ji} has valence 2 in Y_i. This is the second assertion of the Lemma.   q.e.d.

10.13. Lemma. *For each j with $1 \le j \le n$ and each $i \ge 1$, the segment $|\tau_{j,i+1}| \subset T$ is a subsegment of $|\tau_{ji}|$ having the same midpoint as $|\tau_{ji}|$.*

Proof. If $j \ne j(i)$ then $\tau_{ji} = \tau_{j,i+1}$ and the assertion is trivial. If $j = j(i)$, then since Y_{i+1} is obtained from Y_i via the edge-shortening determined by τ_{ji}, and $\tau_{j,i+1}$ is the short edge replacing the long edge τ_{ji}, the assertion follows from 10.6.   q.e.d.

10.14. Lemma. *For each j with $1 \le j \le n$ there is an integer $i_0 \ge 1$ such that $\Gamma_{\tau_{ij}}$ is independent of i for $i \ge i_0$. Furthermore, for all $i \ge i_0$, the valence of $\tau_{j,i+1}$ in Y_{i+1} is less than or equal to the valence of τ_{ji} in Y_i.*

Proof. Given j, consider the sequence $|\tau_{j1}|$, $|\tau_{j2}|$, ... of segments of T. According to Lemma 10.13 this is a monotone decreasing sequence of segments all having the same midpoint in $\mathbb{R}T$. According to Lemma 10.12, the lengths of these segments tend to 0. Thus the Ascending Chain Condition 6.1 asserts that there is a positive integer i_0 such that $\Gamma_{|\tau_{ji}|}$ is independent of i for $i \ge i_0$.

Now let $i \ge i_0$ be given. We must show that $\mathrm{val}_{Y_{i+1}}(\tau_{j,i+1}) \le \mathrm{val}_{Y_i}(\tau_{ji})$. If $j \ne j(i)$ then $\tau_{j,i+1}$ is not in the orbit of ψ_{i+1}; hence it follows from Lemma 10.7 that $\mathrm{val}_{Y_{i+1}}(\tau_{j,i+1}) \le \mathrm{val}_{Y_i}(\tau_{ji})$. On the other hand, suppose that $j = j(i)$. Then since $n(Y_i) = n = n(Y_{i+1})$ and $\Gamma_{|\tau_{ji}|} = \Gamma_{|\tau_{j,i+1}|}$, it follows from Lemma 10.8 that $\mathrm{val}_{Y_{i+1}}(\tau_{j,i+1}) = 2 = \mathrm{val}_{Y_i}(\tau_{ji})$.

10.15. According to Lemma 10.14, we may assume after replacing the sequence (Y_i) be a subsequence that, for every j with $1 \le j \le n$, the group $\Gamma_{|\tau_{ji}|}$ is independent of i for $i \ge 1$, and $\mathrm{val}_{Y_{i+1}}(\tau_{j,i+1}) \le \mathrm{val}_{Y_i}(\tau_{ji})$ whenever $i \ge 1$.

10.16. Lemma. *For all sufficiently large i, the complex Y_i is a singular surface.*

Proof. Since each Γ-orbit in $\mathscr{E}(Y_i)$ contains one of the τ_{ji}, the complex Y_i is a singular surface if and only if τ_{ji} has valence 2 in Y_i for $j = 1, \cdots, n$. Thus we need only show that for each j with $1 \le j \le n$, the 1-simplex τ_{ji} has valence 2 in Y_i for all sufficiently large i.

Let j be given. For each $i \ge 1$, let c_i denote the valence of τ_{ji} in Y_i. By 10.15 we have $c_{i+1} \le c_i$ for all i. On the other hand, by Lemma 10.12,

we have $c_i = 2$ for infinitely many i. Hence $c_i = 2$ for all sufficiently large i.

10.17. According to Lemma 10.16, we may assume, after replacing the sequence (Y_i) by a subsequence, that Y_i is a singular surface for all $i \geq 1$. In order to complete the proof of Proposition 10.1 it is enough to show that $l|Y_1$ is nondegenerate. The following elementary fact about complexes will be useful.

10.18. Lemma. *Let a group Γ act simplicially on a complex K. Suppose that K is finite* $\mod \Gamma$. *Let x be any vertex of K. Then the link of x in K is finite modulo the stabilizer Γ_x of x in Γ.*

Proof. Define a *based simplex* of K to be a pair (Δ, y), where Δ is a simplex of K and y is a vertex of Δ. Since K is finite $\mod \Gamma$, and since a simplex has only finitely many vertices, there are only finitely many Γ-orbits of based simplices in K. On the other hand, two based simplices with base vertex x lie in the same Γ-orbit only if they lie in the same Γ_x orbit. Hence there are only finitely many Γ_x-orbits of based simplices with the vertex x. The conclusion follows. q.e.d.

10.19. To show that $l|Y_1$ is nondegenerate we must verify conditions (i) and (ii) of 7.2. It follows from 10.18 that every local branch β is finite modulo its stabilizer; this immediately implies (ii). It remains to verify (i): that the order of $l|Y_1$ in any 2-dimensional local branch of Y_1 is at least 2.

In the following lemmas, 10.20–10.26, x will denote a vertex of Y_1 and β_1 will denote a 2-dimensional local branch at x. Recall that β_1 is a subcomplex of the star of x, and is the cone over a component B_1 of the link of x. Since Y_1 is a singular surface, B_1 is homeomorphic to either S^1 or $\mathbb{R}$. We shall write o for the order of $l|Y_i$ in β_1. Lemma 10.22 will assert that $o \neq 1$, and Lemma 10.26 will imply that $o \neq 0$. When these have been proved, the proof of Proposition 10.1 will be complete.

10.20. Lemma. *If B_1 is homeomorphic to $\mathbb{R}$ then o is either 0 or $\aleph_0$.*

Proof. Let Γ_1 denote the subgroup of Γ consisting of all elements which fix x and leave B_1 invariant. It follows from 10.18 that B_1 is finite $\mod \Gamma_1$. Hence if B_1 is homeomorphic to $\mathbb{R}$ then each vertex of B_1 must have an infinite orbit under Γ_1. The assertion follows. q.e.d.

Since X is natural, we may identify $sk_0(X)$ with T. We define an equivalence relation $\cong$ on the set of vertices of B_1 by writing $y \cong y'$ if and only if $[x, y]_T \cap [x, y']_T \neq \{x\}$. It follows from the tree axioms that this is indeed an equivalence relation.

10.21. Lemma. *Let τ be a 1-simplex of B_1, and let y and y' denote the endpoints of τ. Then $y \not\cong y'$ if and only if τ is the longest edge of the 2-simplex $\delta = \Delta(x, y, y')$.*

Proof. Suppose that τ is the longest edge of δ. Then it follows from the definition of the complex X that $x \in [y, y']$; hence $[x, y] \cap [x, y'] = \{x\}$, so that $y \not\cong y'$. Now suppose that τ is not the longest edge of δ; by symmetry we may suppose that $\Delta(x, y)$ is the longest edge of δ. By the definition of X we have $y' \in [x, y]$. Hence $[x, y'] \subset [x, y]$, so that $y \cong y'$.

10.22. Lemma. *We cannot have $o = 1$.*

Proof. If B_1 is homeomorphic to $\mathbb{R}$ this follows from Lemma 10.20. Now suppose that B_1 is homeomorphic to S^1. Then the valence c of x in β_1 is finite. Let the vertices of B_1 be labelled y_r, where r ranges over the integers $\bmod\, c$, in such a way that y_r and y_s are the endpoints of a 1-simplex of B_1 if and only if $r - s = \pm 1$. If $o = 1$, there is a unique r, say $r = 0$, such that the longest edge of $\Delta(x, y_r, y_{r+1})$ does not contain x. It then follows from Lemma 10.21 that $y_r \cong y_{r+1}$ if and only if $r \neq 0$. Thus we have $y_1 \cong y_2 \cong \cdots \cong y_{c-1} \cong y_0 \not\cong y_1$, an absurdity.   q.e.d.

We now turn to the proof that $o \neq 0$. Recall from 10.11 that for each $i \geq 1$, we obtain Y_{i+1} from Y_i by the edge-shortening determined by the 1-simplex ω_i. In particular, the subcomplexes Y_0, Y_1, ... of X all have the same 0-skeleton. Thus x is a vertex of Y_i for any $i \geq 1$.

In the following lemmas, 10.23–10.26, i will denote an arbitrary integer ≥ 1, and β will denote an arbitrary local branch of Y_i at x. We shall denote the order of x in β by $o_i(\beta)$. Lemma 10.26 will assert that $o_i(\beta) \neq 0$; if we take $i = 1$ and $\beta = \beta_1$ it will follow that $o \neq 0$.

We shall denote by $S_i(\beta)$ the set of all integers j with $1 \leq j \leq n$ such that the Γ-orbit of τ_{ji} in Y_i contains a 1-simplex in β having x as an endpoint. We shall write $m_i(\beta)$ for the cardinality of the set $S_i(\beta)$.

10.23 Lemma. *Suppose that $o_i(\beta) = 0$ and that $j(i) \notin S_i(\beta)$. Then β is contained in Y_{i+1} and is in fact a local branch of Y_{i+1} at x. Furthermore, we have $m_{i+1}(\beta) = m_i(\beta)$ and $o_{i+1}(\beta) = 0$.*

Proof. The local branch β is the cone over a component B of the link of x in Y_i. If τ is any 1-simplex in B, there is a unique 2-simplex δ in β having τ as an edge. Since $o_i(\beta) = 0$, the long edge of δ cannot be τ; in particular, $\tau \notin \mathscr{E}^*(Y_i)$, and so τ cannot be in the orbit of ω_i. Thus B contains no 1-simplices in the orbit of ω_i. On the other hand, by hypothesis we have $j(i) \notin S_i(\beta)$, which means that no 1-simplex in β having x as an endpoint belongs to the orbit of ω_i. In short, no 1-simplex in β belongs to the orbit of ω_i. Since Y_{i+1} is obtained from Y_i

by the edge-shortening determined by ω_i, it now follows that $\beta_{i+1} \subset Y$.

Hence B is contained in the link L of x in Y_{i+1}. Let B' denote the component of L containing B. Since Y_i and Y_{i+1} are singular surfaces, each of the complexes B and B' is a connected combinatorial 1-manifold. Since $B \subset B'$ we must have $B = B'$. This shows that β is a local branch of Y_{i+1}.

It now follows from the definitions that $m_{i+1}(\beta) = m_i(\beta)$ and $o_{i+1}(\beta) = 0$.

10.24. Lemma. *Suppose that $o_i(\beta) = 0$ and that $j(i) \in S_i(\beta)$. Then there is a local branch β' of Y_{i+1} at x such that $o_{i+1}(\beta') = 0$ and $m_{i+1}(\beta') < m_i(\beta)$.*

Proof. The local branch β is the cone over a component B of the link of x in Y_i. As in the proof of 10.23, the hypothesis $o_i(\beta) = 0$ implies that B contains no 1-simplices in the orbit of ω_i. On the other hand, the hypothesis that $j(i) \in S_i(\beta)$ means that for some vertex y of B, the 1-simplex $\Delta(x, y)$ is in the orbit of ω_i.

Let Q denote the set of all vertices y of B such that $\Delta(x, y)$ is in the orbit of ω_i. Then for any $y \in Q$, we have $\Delta(x, y) \in \mathscr{E}_2^*(Y_i)$, and we may regard Y_{i+1} as being obtained from Y_i via the edge-shortening determined by $\Delta(x, y)$. The link of y in B consists of two vertices z and w, and the star of y in B is the arc $A_y = \Delta(y, z) \cup \Delta(y, w)$. The 2-simplices of Y_i incident to $\Delta(x, y)$ are $\Delta(x, y, z)$ and $\Delta(x, y, w)$. Hence $\Delta(x, z)$ and $\Delta(x, w)$ are strictly shorter than $\Lambda(x, y)$, so that the endpoints z and w of A_y do not belong to Q, and the definition of edge-shortening implies that $\Delta(x, z, w)$ is a simplex of Y_{i+1}. In particular, $e_y = \Delta(z, w)$ is a 1-simplex in the link of x in Y_{i+1}. For each $y \in Q$ we fix a homeomorphism h_y of A_y onto e_y which fixes the endpoints of A_y.

If $\Delta(u, v)$ is a 1-simplex in $B - \bigcup_{y \in Q} A_y$, then no edge of $\Delta(x, u, v)$ is in the orbit of ω_i; hence by the definition of edge-shortening, $\Delta(x, u, v)$ is a simplex of Y_{i+1}, and thus $\Delta(u, v)$ lies in the link of x in Y_{i+1}.

We define a map h of B into the link of x in Y_{i+1} by setting $h|A_y = h_y$ for every $y \in Q$, and defining h to be the identity outside $\bigcup_{y \in Q} A_y$. (To see that h is well defined we must know that $A_y \cap A_{y'} = \varnothing$ for any distinct elements y and y' of Q; this follows from the above observation that for any $y \in Q$, the endpoints z and w of A_y lie outside Q.)

We shall show that h is a homeomorphism of B onto a component B' of the link of x in Y_{i+1}, and that the conclusions of the lemma hold for the local branch β' corresponding to B'.

We must first show that h is 1-1. To prove this, it is enough to show (i) that $h(A_y)$ is disjoint from $B - (\bigcup_{y \in Q} \text{int } A_y)$ for every $y \in Q$, and (ii) that $h(\text{int } A_y) \cap h(\text{int } A_{y'}) = \varnothing$ for any distinct elements y and y' of Q. To prove (i), note that $h(A_y) = e_y$ is, by construction, the edge replacing the long edge $\Delta(x, y)$ in the edge-shortening. Hence by 10.11, e_y is in the Γ-orbit of ψ_{i+1}. On the other hand, each 1-simplex of $B - (\bigcup_{y \in Q} \text{int } A_y)$ is the Γ-orbit of $\tau_{j, i+1}$ for some $j \neq j(i)$, and by 10.11, the orbits of $\tau_{1, i+1}, \cdots, \tau_{n, i+1}$ are distinct. This proves (i).

To prove (ii), we must show that the 1-simplices e_y and $e_{y'}$ are distinct when $y, y' \in Q$ are distinct. We may write $\Delta(x, y) = g \cdot \omega_i$ and $\Delta(x, y') = g' \cdot \omega_i$ for some $g, g' \in \Gamma$. Then we have $e_y = g \cdot \psi_{i+1}$ and $e_{y'} = g' \cdot \psi_{i+1}$. Thus $e_y = e_{y'}$ would imply $g^{-1} g' \in \Gamma_{|\psi_{i+1}|}$; by 10.15 this would in turn imply $g^{-1} g' \in \Gamma_{|\omega_i|}$, and hence $\Delta(x, y) = \Delta(x, y')$. This is impossible if $y \neq y'$. Thus (ii) is proved, and the map h has been shown to be 1-1.

In particular $h(B)$ is a 1-manifold (without boundary). It is clear from the definition of h that $h(B)$ is a subcomplex of the link of x in Y_{i+1}. Since Y_{i+1} is a singular surface it follows that $B' = h(B)$ is a connected component of the link of x in Y_{i+1}. Thus B' determines a local branch β' of Y_{i+1} at x. We must verify the conclusions of the lemma with this choice of β'.

To show that $o_{i+1}(\beta') = 0$, we must show that for any 1-simplex $\Delta(z, w)$ of B', the longest edge of $\Delta(x, z, w)$ is $\Delta(x, z)$ or $\Delta(x, w)$ and not $\Delta(z, w)$. By the definition of B', either $\Delta(z, w)$ is a 1-simplex of B, or $\Delta(z, w) = e_y$ for some $y \in Q$. If $\Delta(z, w)$ is a 1-simplex of B, then since $o_i(\beta) = 0$, the longest edge of $\Delta(x, z, w)$ cannot be $\Delta(z, w)$. Now suppose that $\Delta(z, w) = e_y$ for some $y \in Q$. Then we may denote the lengths of $\Delta(x, y)$, $\Delta(x, z)$ and $\Delta(x, w)$ respectively by a, b and $a - b$; and after possibly changing notation so as to reverse the roles of z and w, we may assume that $b < 2a$. Then the length of $\Delta(z, w)$ is $a - 2b < a - b$. Hence $\Delta(z, w)$ is not the longest edge of $\Delta(x, z, w)$.

It remains to show that $m_{i+1}(\beta') < m_{i+1}(\beta)$. It follows from the definition of B' that for any 1-simplex τ in X having x as an endpoint, τ lies in β' if and only if it lies in β and is not in the orbit of ω_i. Thus we have $S_{i+1}(\beta') = S_i(\beta) - \{j(i)\}$, and hence $m_{i+1}(\beta') = m_{i+1}(\beta) - 1$.

10.25. Lemma. *Suppose that $o_i(\beta) = 0$. Then there exist an integer $i' > i$ and a local branch β' of $Y_{i'}$ at x such that $o_{i'}(\beta') = 0$ and $m_{i'}(\beta') < m_i(\beta)$.*

Proof. According to Lemma 10.12, for any j with $1 \leq j \leq n$, there exists some $i' > i$ such that $j(i' - 1) = j$. In particular, since $S_i(\beta)$ is nonempty, there exists some $i' > i$ such that $j(i' - 1) \in S_i(\beta)$.

Let us choose the *least* value of $i' > i$ such that $j(i' - 1) \in S_i(\beta)$. Then by $i' - i - 1$ applications of Lemma 10.23, β is contained in $Y_{i'-1}$ and is in fact a local branch of $Y_{i'-1}$ at x. Furthermore, we have $m_{i'-1}(\beta) = m_{i-1}(\beta)$ and $o_{i'-1}(\beta) = 0$. Lemma 10.24 now shows that there is a local branch β' of $Y_{i'}$ at x such that $o_{i'}(\beta') = 0$ and $m_{i'}(\beta') < m_{i'-1}(\beta) = m_i(\beta)$.

10.26. Lemma. *We cannot have $o_i(\beta) = 0$.*

Proof. If $o_i(\beta) = 0$, then 10.25 gives an $i' > i$ and a local branch β' of $Y_{i'}$ such that $o_{i'}(\beta') = 0$ and $m_{i'}(\beta') < m_i(\beta)$. Applying 10.25 again we get an integer $i'' > i'$ and a local branch β'' of $Y_{i''}$ with $o_{i''}(\beta'') = 0$ and $m_{i''}(\beta'') < m_{i'}(\beta')$. Continuing in this way we get a strictly decreasing sequence $m_i(\beta)$, $m_{i'}(\beta')$, $m_{i''}(\beta'')$, ... of positive integers. This cannot be. q.e.d.

10.27.   The proof of Proposition 10.1 is now complete, by virtue of the remarks in 10.19.

11. Proof of the structure theorem in the rank-2 case

This section is devoted to the proof of Theorem 6.2. Throughout the section, Λ will denote a subgroup of $\mathbb{R}$ whose $\mathbb{Q}$-rank is 2, and $\mathcal{T} = (T, \Gamma, \rho)$ will denote a countable Λ-tree with symmetry. We write Λ as a union $\bigcup_{i=1}^{\infty} L_i$, where each L_i is a rank-2 free abelian subgroup of Λ and $L_i \subset L_j$ whenever $i \leq j$.

11.1. Lemma. *There exist a direct system $(\mathscr{S}_i, \mathscr{W}_{ij})$ in the 0-tame category of $\mathcal{T}$-complexes, with $\mathscr{S}_i = (\Sigma_i, \Gamma, \rho_i, v_i, 1)$ and $\mathscr{W}_{ij} = (W_{ij}, 1)$, and compatible base vertices $v_i \in \Sigma_i$, for which the following conditions hold.*

(i) *Each Σ_i is a singular surface and is finite $\mod \Gamma$.*

(ii) *The natural length system l_i for the T-complex Σ_i takes its values in L_i.*

(iii) *The length system l_i is nondegenerate.*

(iv) *$(\mathscr{S}_i, \mathscr{W}_{ij})$ is abundant.*

(v) *For $n = 0$, 1 we have $\varinjlim \pi_n(\Sigma_i, v_i) = 0$.*

Proof. As in 9.15, we can use the L_i to construct a direct system $((X_i, \phi_i); J_{ij})$ of natural T-complexes and 2-tame $\operatorname{Aut}(T)$-equivariant

T-maps. This gives a direct system $(\mathscr{X}_i, \mathscr{J}_{ij})$ in the 2-tame category of $\mathscr{T}$-complexes, where $\mathscr{X}_i = (X_i, \Gamma, \rho_i, \phi_i, 1)$ and $\mathscr{J}_{ij} = (J_{ij}, 1)$. Since the 1-skeleton of X_i is $U_i = U(L_i)$, it follows from 4.7 that $(\mathscr{X}_i, \mathscr{J}_{ij})$ is abundant.

We choose consistent base vertices v_i for the X_i. By Theorem 9.15 we have $\varinjlim \pi_n(X_i, v_i) = 0$ for $n = 0$, 1.

Since T is countable, each X_i consists of countably many simplices. Hence we may write $X_i = \bigcup_{m=1}^\infty Z_{im}$, where each Z_{im} is a nonempty finite subcomplex of X_i, and $Z_{im} \subset Z_{i,m+1}$ for every $m \geq 1$. We choose the Z_{im} so that $v_i \in Z_{i1}$ for all i.

For the purpose of this proof we define a *simplicial path* in X_i to be a path of the form $\gamma = \zeta_1 * \cdots * \zeta_n$, where each ζ_i is an affine isomorphism of $[0, 1]$ onto a closed 1-simplex in X_i. If $\gamma(0) = \gamma(1) = v_i$ we call γ a *simplicial loop based at* v_i.

Since X_i contains only countably many simplices, it contains only countably many simplicial loops based at v_i. We index these as γ_{i1}, $\gamma_{i2}, \ldots$.

We shall recursively define Γ-invariant subcomplexes Σ_i of the X_i which are singular surfaces and are finite $\bmod \Gamma$, and Γ-equivariant maps $W_{ij}: \Sigma_i \to \Sigma_j$ for $i \leq j$. Let $i \geq 1$ be given. Suppose that for every j such that $1 \leq j < i$ we have defined a Γ-invariant singular surface $\Sigma_j \subset X_j$ which is finite $\bmod \Gamma$. Suppose also that for all j, k such that $1 \leq j \leq k < i$ we have defined a Γ-equivariant map $W_{jk}: \sigma_j \to \Sigma_k$. Let us fix a finite subcomplex M_i of X_i with the following properties:

(a) $v_i \in M_i$;

(b) $J_{ji}(Z_{ji}) \subset M_i$ for every $j < i$;

(c) if j and m are positive integers $< i - 1$, such that $|\gamma_{jm}| \subset \Sigma_j$ and such that $W_{j,i-1} \circ \gamma_{jm}$ is null-homotopic in X_{i-1}, then $J_{i-1,i} \circ W_{j,i-1} \circ \gamma_{jm}$ has support contained in M_i and is null-homotopic in M_i.

Now let Y_i be a Γ-invariant subcomplex of X_i which contains M_i and is finite $\bmod \Gamma$. If $i > 1$, we take Y_i to contain $J_{i-1,i}(\Sigma_{i-1})$. By 10.1 there exists a Γ-equivariant T-map $\mathfrak{G}_i: Y_i \to X_i$ such that $\Sigma_i = \mathfrak{G}_i(Y_i)$ is a singular surface. Furthermore, $\mathfrak{G}_i | \operatorname{sk}_1(Y_i)$ is homotopic $\operatorname{rel} \operatorname{sk}_0(Y_i)$ to the inclusion map $\operatorname{sk}_1(Y_i) \to X_i$; and the natural L_i-length system l_i for the T-complex $\Sigma_i \subset X_i$ is nondegenerate.

We define W_{ii} to be the identity map on Σ_i. If $i > 1$ we set $W_{i-1,i} = \mathfrak{G}_i \circ (J_{i-1,i} | \Sigma_{i-1}): \Sigma_{i-1} \to \Sigma_i$, and for $j < i - 1$ we set $W_{ji} = W_{i-1,i} \circ W_{j,i-1}$. This completes the recursive definition of the Σ_i and the W_{ij}.

For each i we set $v_i = \phi_i | \Sigma_i$. Since Σ_i is a Γ-invariant subcomplex of X_i, we have a $\mathcal{T}$-complex $\mathcal{S}_i = (\Sigma_i, \Gamma, \rho_i, v_i, 1)$ for each i. It follows from the construction that the W_{ij} are 0-tame Γ-equivariant T-maps and that we have $W_{jk} \circ W_{ij} = W_{ik}$ for $i \leq j \leq k$. Hence, setting $\mathcal{W}_{ij} = (W_{ij}, 1)$, we have a direct system $(\mathcal{S}_i, \mathcal{W}_{ij})$ in the category of $\mathcal{T}$-complexes. Since the $\mathfrak{G}_i$ fix the vertices Y_i, the maps J_{ij} and W_{ij} agree on the vertices of Σ_i. In particular the v_i are consistent base vertices for the Σ_i.

Note that for any $i > 1$, since $\mathfrak{G}_i | \mathrm{sk}_1(Y_i)$ is homotopic $\mathrm{rel}\,\mathrm{sk}_0(Y_i)$ to the inclusion map $\mathrm{sk}_1(Y_i) \to X_i$, the map $\mathfrak{G}_i : Y_i \to X_i$ induces the same homomorphism from $\pi_1(Y_i, v_i)$ to $\pi_1(X_i, v_i)$ as the inclusion map. Hence $J_{i-1,i}$ and $W_{i-1,i}$ induce the same homomorphism from $\pi_1(\Sigma_{i-1}, v_{i-1})$ to $\pi_1(X_i, v_i)$. It follows that J_{ij} and W_{ij} induce the same homomorphism from $\pi_1(\Sigma_i, v_i)$ to $\pi_1(X_j, v_j)$ whenever $i \leq j$.

By construction, $(\mathcal{S}_i ; \mathcal{W}_{ij})$ satisfies conditions (i)–(iii) of the lemma. We shall complete the proof of the lemma by showing that it satisfies conditions (iv) and (v) as well: that is, that $(\mathcal{S}_i, \mathcal{W}_{ij})$ is abundant and that $\varinjlim \pi_n(\Sigma_i, v_i) = 0$ for $n = 0, 1$.

To show that $(\mathcal{S}_i ; \mathcal{W}_{ij})$ is abundant we must check conditions (i)–(iii) of 4.6. To check 4.6(i) we consider an arbitrary point $x \in T$. Since $(\mathcal{X}_i ; \mathcal{F}_{ij})$ is abundant, there exist an index j and a vertex $w \in X_j$ such that $\phi_j(w) = x$. For some $i \geq j$ we have $w \in Z_{ji}$. Hence $v = J_{ji}(w) \in M_i \subset Y_i$. Since $\mathfrak{G}_i$ fixes the vertices of Y_i we have $v \in \Sigma_i$. Now $v_j(v) = \phi_j(v) = x$, and 4.6(i) is established.

To check 4.6(ii) we consider a positive integer i, a parametrized segment γ in T, and two vertices v_0 and v_1 of X_i such that $\phi_i(v_t) = \gamma(t)$ for $t = 0, 1$. Since $(\mathcal{X}_i ; \mathcal{F}_{ij})$ is abundant, there exist an integer $k \geq i$ and a 0-tame lift $\hat{\gamma}_0$ of γ to X_k such that $\hat{\gamma}_0(t) = J_{ik}(v_t)$ for $t = 0, 1$. For some $j \geq k$ we have $|\hat{\gamma}_0| \subset Z_{kj}$, so that $|J_{kj} \circ \hat{\gamma}_0| \subset Y_j$. Now $\hat{\gamma} = \mathfrak{G}_j \circ J_{kj} \circ \hat{\gamma}_0$ is a 0-tame lift of γ to X_j. Furthermore, since $\mathfrak{G}_k$ fixes the vertices of Y_k we have $\hat{\gamma}_0(t) = J_{ij}(v_t)$ for $t = 0, 1$. This proves 4.6(ii).

Condition 4.6(iii) is trivial in this case since the map $\omega_i : \Gamma_i \to \Gamma$ is the identity for every i. Hence $(\mathcal{S}_i ; \mathcal{W}_{ij})$ is abundant.

To prove that $\varinjlim \pi_0(\Sigma_i, v_i) = 0$ we must show that for any i and any vertex v of Σ_i, there is an integer $j \geq i$ such that $J_{ij}(v) = W_{ik}(v)$ and $J_{ij}(v_i) = W_{ij}(v_i)$ lie in the same connected component of Σ_j. To show this, we note that since $\varinjlim \pi_0(X_i, v_i) = 0$, there is an integer $k \geq i$ such that $J_{ik}(v)$ and $J_{ik}(v_i)$ are in the same component of X_i. Hence there exists some $j \geq k$ such that $J_{ik}(v)$ and $J_{ik}(v_i)$ are in the same component

of Z_{kj}. Therefore $J_{ij}(v)$ and $J_{ij}(v_i)$ are in the same component of Y_j. Since $\mathfrak{G}_j$ fixes the vertices of Y_j it follows that $J_{ij}(v)$ and $J_{ij}(v_i)$ are in the same component of Σ_j, as required.

It remains to prove that $\varinjlim \pi_1(\Sigma_i, v_i) = 0$. We must show that if γ is a simplicial loop in Σ_i, based at v_i, then $W_{ij} \circ \gamma$ is homotopically trivial in Σ_j for some $j \geq i$. We have $\gamma = \gamma_{im}$ for some m. Since $\varinjlim \pi_1(X_i, v_i) = 0$, there exists $j > 1 + \max(i, m)$ such that $J_{i,j-1} \circ \gamma$ is homotopically trivial in X_j. Since J_{ij} and W_{ij} induce the same homomorphism from $\pi_1(\Sigma_i, v_i)$ to $\pi_1(X_j, v_j)$, it follows that $W_{i,j-1} \circ \gamma = W_{i,j-1} \circ \gamma_{im}$ is homotopically trivial in X_{j-1}. By property (c) of the set M_j this implies that $J_{j-1,j} \circ W_{i,j-1} \circ \gamma$ has support contained in M_j, and is homotopically trivial in M_j and therefore in Y_j. Hence $W_{ij} \circ \gamma = \mathfrak{G}_j \circ J_{j-1,j} \circ W_{i,j-1} \circ \gamma$ is homotopically trivial in $\Sigma_j = \mathfrak{G}_j(Y_j)$, as required.

11.2 Lemma. *There exist a direct system* $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ *in the 0-tame category of* $\mathscr{T}$-*complexes, with* $\widetilde{\mathscr{S}}_i = (\widetilde{\Sigma}_i, \Gamma_i, \tilde{\rho}_i, \tilde{v}_i, \omega_i)$ *and* $\widetilde{\mathscr{W}}_{ij} = (\widetilde{W}_{ij}, \omega_{ij})$, *and compatible base points* $v_i \in \widetilde{\Sigma}_i$ *for which the following conditions hold.*

(i) *Each* $\widetilde{\Sigma}_i$ *is a 1-connected singular surface and is finite* $\mod \Gamma_i$.

(ii) *The natural length system* $\tilde{l}_i$ *for the T-complex* $\widetilde{\Sigma}_i$ *takes its values in* L_i.

(iii) *The length system* $\tilde{l}_i$ *is nondegenerate.*

(iv) $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ *is abundant.*

Proof. Let $(\mathscr{S}_i; \mathscr{W}_{ij})$ be the direct system given by Lemma 11.1. Define $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ to be the universal cover (4.8) of $(\mathscr{S}_i; \mathscr{W}_{ij})$. The $\widetilde{\Sigma}_i$ are of course 1-connected. By Proposition 4.8, $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ is abundant. The other asserted properties of $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ follow immediately from the corresponding properties of $(\mathscr{S}_i; \mathscr{W}_{ij})$. q.e.d.

11.3. *Proof of Theorem 6.2.* Let $(\widetilde{\mathscr{S}}_i; \widetilde{\mathscr{W}}_{ij})$ be the direct system given by 11.2. For each i, since $\tilde{l}_i$ is nondegenerate, it defines an L_i-foliation μ_i on the singular surface $\widetilde{\Sigma}_i$. Thus for each i, we have a simply connected, uniform L_i-foliated singular surface with symmetry $(\widetilde{\Sigma}_i, \tilde{\mu}_i, \Gamma_i, \tilde{\rho}_i)$. Let $\mathscr{T}_i = (T_i, \Gamma_i, \overline{\rho}_i)$ denote the preleaf space of $(\widetilde{\Sigma}_i, \tilde{\mu}_i, \Gamma_i, \tilde{\rho}_i)$. By 5.28, $\mathscr{T}_i$ is an L_i-pretree with symmetry; in particular it is a Λ-pretree with symmetry.

According to 7.8, the T-map $\widetilde{W}_{ij}$ is a morphism of Λ-foliated singular surfaces for all $i \leq j$. Hence $\widetilde{\mathscr{W}}_{ij} : (\widetilde{\Sigma}_i, \tilde{\mu}_i, \Gamma_i, \tilde{\rho}_i) \rightarrow (\widetilde{\Sigma}_j, \tilde{\mu}_j, \Gamma_j, \tilde{\rho}_j)$

is a morphism of L_i-foliated singular surfaces with symmetry. It follows from 5.29 that $\widetilde{\mathscr{W}}_{ij}$ induces a morphism $\mathscr{f}_{ij} = (f_{ij}, \omega_{ij}): \mathscr{T}_i \to \mathscr{T}_j$ of trees with symmetry. Thus we have a direct system $(\mathscr{T}_i; \mathscr{f}_{ij})$ in the category of Λ-pretrees with symmetry.

For each positive integer i, it follows from 7.6 that there is a natural morphism $\mathscr{f}_i = (f_i, \omega_i): \mathscr{T}_i \to \mathscr{T}$. Furthermore, the kernel of ω_i is the group of covering transformations of $\widetilde{\Sigma}_i$ and therefore acts freely on $\widetilde{\Sigma}_i$. Thus to complete the proof of Theorem 6.2 we need only show that $(\mathscr{T}_i; \mathscr{f}_{ij})$ converges strongly and that its limit can be identified isomorphically with $\mathscr{T}$ is such a way that the $\mathscr{f}_i$ are the canonical morphisms. To do this we must check that conditions (i)–(iv) of 1.26 are satisfied.

Condition 1.26(i) follows from 7.8. The verification of conditions (ii)–(iv) of 1.26 uses the fact that the direct system $(\widetilde{\mathscr{T}}_i; \widetilde{\mathscr{W}}_{ij})$ is abundant. Indeed, condition (iv) of 1.26 follows immediately from condition (iii) of Definition 4.6. Condition 1.26(ii) follows from 4.6(i) since, by definition, T_i is the image of the $\mathrm{sk}_0(\widetilde{\Sigma}_i)$ under the natural map $\chi_i: \widetilde{\Sigma}_i \to T_i$.

In order to check condition (iii) of 1.26 we must consider an integer $i \geq 1$ and two points $x_0, x_1 \in T_i$. Set $y_t = f_i(x_t)$ for $t = 0$, 1. Let γ be a parametrized segment in T with $\gamma(t) = y_t$ for $t = 0$, 1. Choose vertices v_0, v_1 of $\widetilde{\Sigma}_i$ such that $x_t = \chi_i(v_t)$. Then $\upsilon_i(v_t) = y_t$. By Condition (ii) of 4.6, there exist an index $j \geq i$ a 0-tame lift $\hat{\gamma}_j$ of γ to $\widetilde{\Sigma}_j$ such that $\hat{\gamma}(t) = \widetilde{W}_{ij}(v_t)$ for $t = 0, 1$.

Since $\tilde{\upsilon}_j$ maps each leaf of $(\widetilde{\Sigma}_j, \tilde{\mu}_j)$ to a point of T, and since $\tilde{\upsilon}_j \circ \hat{\gamma} = \gamma$ is a parametrized segment in T, it follows that the arc $|\hat{\gamma}_j|$ meets each leaf of $(\widetilde{\Sigma}_j, \tilde{\mu}_j)$ in at most a single point. In particular, $\hat{\gamma}$ is taut. Hence $\hat{\gamma}_j = \chi_j \circ \hat{\gamma}$ is a parametrized segment in $\mathbb{R}T_j$ joining $f_{ij}(x_0)$ to $f_{ij}(x_0)$. But $f_j \circ \gamma_j = \gamma$ is a parametrized segment in T joining y_0 to y_1. Thus $f_j||\gamma_j|$ is 1-1 and is therefore an isometric embedding by 1.7. Hence $d(f_{ij}(x_0), f_{ij}(x_11)) = d(f_i(x_0), f_i(y_0))$, as required.

References

[1] R. Alperin & H. Bass, *Length functions of group actions on trees*, in Combinatorial Group Theory and Topology (S. M. Gersten and J. R. Stallings, eds.), Ann. of Math. Studies, No. 111, Princeton Univ. Press, Princeton, NJ, 1987, 265–378.

[2] R. Alperin & K. Moss, *Complete trees for groups with a real-valued length function*, J. London Math. Soc. (2) **31** (1985) 55–68.

[3] I. M. Chiswell, *Abstract length functions in groups*, Math. Proc. Cambridge Philos. Soc. **80** (1976) 451–463.

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

[5] A. Fathi, F. Laudenbach, & V. Poénaru, *Travaux de Thurston sur les surfaces*, Séminaire Orsay, Astérisque no. 66–67, Soc. Math. de France, Paris, 1979.

[6] G. H. Hardy & E. M. Wright, *An introduction to the theory of numbers*, 5th ed., Oxford Univ. Press, Oxford, 1979.

[7] J. Hempel, 3-*manifolds*, Ann. of Math. Studies, No. 86, Princeton Univ. Press, Princeton, NJ, 1976.

[8] R. C. Lyndon, *Length functions in groups*, Math. Scand. **12** (1963) 209–234.

[9] J. W. Morgan, *Group actions on trees and the compactifications of the space of classes of* SO(n, 1)-*representations*, Topology **25** (1986) 1–34.

[10] ——, *Ergodic theory and free actions of groups on* $\mathbb{R}$-*trees*, preprint.

[11] J. W. Morgan & J.-P. Otal, *Non-Archimedean measured laminations and degenerations of surfaces*, preprint.

[12] J. W. Morgan & P. B. Shalen, *Degenerations of hyperbolic structures. I*, Ann. of Math. (2) **120** (1984) 401–476.

[13] ——, *Degenerations of hyperbolic structures. III*, Ann. of Math. (2) **127** (1988) 457–519.

[14] ——, *Actions of surface groups on* $\mathbb{R}$-*trees* (to appear).

[15] J.-P. Serre, *Trees*, Springer, New York, 1980.

[16] P. B. Shalen, *Linear representations of certain amalgamated products*, J. Pure Appl. Algebra **15** (1979) 187–197.

[17] H. Zieschang, E. Vogt & H.-D. Coldeway, *Surfaces and planar discontinuous groups*, Lecture Notes in Math., Vol. 835, Springer, Berlin, 1980.

UNIVERSITY OF ILLINOIS, CHICAGO