Dendrology and its applications

Peter B. Shalen[1]

Preface

The organizers of the ICTP workshop on geometric group theory kindly invited me to give a series of lectures about group actions on real trees in the spirit of my earlier expository article [Sh]. In the five hours that were allotted, I was able to cover a number of aspects of the subject that I had not covered in [Sh], and some developments that had taken place since [Sh] was written. In the present article I have extended the scope much further still, and have treated in some depth a number of topics that were barely mentioned in my lectures. I have tried to show what a broad and rich subject this is and how much other mathematics it interacts with.

Section 1 includes Gupta and Sidki's construction of finitely generated infinite p-groups using simplicial trees, and Brown's interpretation of the Bieri-Neumann invariant in terms of real trees. In Section 2 I describe Bestvina and Handel's work on outer automorphisms of free groups and its connection with Culler-Vogtmann and Gersten's outer space and with exotic free actions of free groups on trees; one example is worked out in considerable detail. Section 2 also includes a sketch of Skora's proof that every small action of a surface groups on a real tree is dual to a measured foliation.

In Section 3 I give an account of the Bruhat-Tits tree for $textSL_2$ of a valued field, and its application by Lubotzky, Phillips and Sarnak in their work on Ramanujan graphs. I also discuss in some depth my work with Culler on trees associated to ideal points of curves in the character variety of a group and their applications in 3-manifold theory, including the proof of the Cyclic Surgery Theorem by Culler, Gordon, Luecke and myself. In Section 4 I explain the connection between trees and hyperbolic geometry in much greater depth than in [Sh], and include accounts of both the original approach used in my work with Morgan and based on the Bruhat-Tits tree, and the approach of Bestvina and Paulin based on Gromovian notions of convergence of metric spaces. There is also a brief discussion of Paulin's work on finiteness of outer automorphism groups. The section concludes with some idle speculations.

In Section 5 the emphasis is on the aspects of my work with Gillet on rank-2 trees that were left out of [Sh], especially the notion of strong convergence. This section includes a brand-new, and therefore somewhat tentative, conjecture about how to extend our theory to arbitrary rank. I also discuss some surprising connections between the notion of strong convergence on the one hand, and both Bestvina-Handel theory and the contractibility of outer space on the other.

In order to make this article self-contained, or at least coherent, it has been necessary to allow some slight overlap with [Sh]. I have tried to minimize it by taking a different point of view from that of [Sh] wherever possible. In the few cases where I needed to repeat something that I had said in [Sh], I have tried to be very brief.

I have tried to maintain the same informal style as in [Sh], which is intended to approximate the tone of a lecture rather than a journal article. But in print one is at a disadvantage in that one cannot use one's hands—for example, to point to the board where some important theorem had been written before it was erased. To compensate for this, I have divided the article into lots of subsections and included lots of cross-references, which I hope will keep non-expert readers from getting lost. If you find cross-references irksome, I can only ask you to ignore them.

In writing this article I have had to come to grips with a number of aspects of the subject of which I was only dimly conscious before. I am indebted to Marc Culler for the many hours that he has spent listening to my thoughts on the material and helping me unscramble them.

The tentative conjecture stated in 5.5.6 was formulated with the help of Marc Culler, Henri Gillet and Richard Skora. However, I will take the responsibility if it falls flat.

I am grateful to Mladen Bestvina for working out and lucidly explaining the example given in 2.2.

[1]Partially supported by an NSF grant.

I don't usually enjoy conferences much, and I confess I didn't look forward to stepping off a trans-Atlantic flight and giving five lectures. As it turned out, ICTP is such a delightful place, and the organizers—Alberto Verjovsky, André Haefliger and Etienne Ghys—gave me such a nice welcome, that it was better than being on vacation.

Section 1. Generalities (and digressions)

1.1. Simplicial trees. A *simplicial tree* is a connected, simply connected simplicial 1-complex T. Combinatorially, simple connectivity means that T contains no circuits, or equivalently that every edge separates T into two pieces.

1.1.1. A substantial part of classical combinatorial group theory can be interpreted as the study of (simplicial) actions of groups on simplicial trees. For example, a group is free if and only if it acts on a simplicial tree in such a way that no non-trivial element of the group leaves any simplex invariant. An immediate consequence of this fact is the Nielsen-Schreier theorem, which asserts that any subgroup of a free group is free.

1.1.2. It is usually better to think of a simplicial tree not as a set of simplices, but as a set of vertices with the structure given, say, by the adjacency relation. So when one says that a group acts *freely* on a simplicial tree one means that no non-trivial element fixes any vertex. For example, the tree with two vertices and one edge admits a free $\mathbf{Z}/2\mathbf{Z}$-action. Following Serre [Se], one says that a group acts *without inversions* if no element of the group leaves an edge invariant but interchange its endpoints. We can reformulate the assertion of 1.1.1 by saying that a group is free if and only if it acts freely and without inversions on some simplicial tree.

1.1.3. The assumption that a group acts without inversions is not a serious restriction, because any (simplicial) action of Γ on T induces an action without inversions on the first barycentric subdivision of T.

1.1.4. It is easy to show that every free group acts freely and without inversions on some simplicial tree: one need only take the tree to be the Cayley graph with respect to some free generating set. Topologists often prove the converse by saying that if Γ acts freely and without inversions on T then T/Γ is a graph with universal covering space T and deck transformation group Γ, so that Γ is isomorphic to the fundamental group of a graph and is therefore free. There is a good deal of machinery hidden in this argument, and it is instructive to give a direct combinatorial proof.

This can be done by constructing a fundamental domain for the action, i.e. a subtree K of the first barycentric subdivision of T whose endpoints are midpoints of edges of T, and such that (i) $T = \bigcup_{\gamma\in\Gamma} \gamma \cdot K$ and (ii) $\operatorname{int} K \cap \gamma \cdot K = \emptyset$ for every $\gamma \neq 1$. Now for each endpoint x of K there is a unique endpoint $x' \neq x$ of K that lies in the same Γ-orbit as x, and $\tau\colon x \mapsto x'$ is an involution of the set of endpoints of K. Let S be a complete system of orbit representatives for τ, and for each $s \in S$ let x_s denote the unique element of Γ that maps s to $\tau(s)$. Then it is straightforward to show that $(x_s)_{s\in S}$ is a system of free generators for Γ. For example, to show that a non-trivial reduced word $\prod_{i=1}^{n} x_{s_i}^{\epsilon_i}$ (where $\epsilon = \pm 1$) cannot represent the identity, we consider the sets $K_m = \prod_{i=1}^{m} x_{s_i}^{\epsilon_i}(K)$ for $m = 0, \dots n$. It follows from the definitions that whenever $1 \leq m < n$, the sets K_{m-1} and K_{m+1} meet K_m in distinct endpoints of K_m, and hence lie in distinct components of $T - \operatorname{int} K_m$. It follows easily that $K_0, \dots K_n$ are all distinct. In particular $K_n \neq K_0$, and the assertion follows.

1.1.5. In any event, a free actions without inversions of a group Γ on a simplicial tree T determines a quotient graph T/Γ, and the classification of such actions up to equivariant simplicial isomorphism is equivalent to the classification of connected graphs. This picture was generalized by Bass and Serre in [Se]. They showed that the classification of arbitrary actions without inversions on simplicial trees is equivalent to the classification of what are called graphs of groups.

To define a *graph of groups* one must specify (i) a graph $\mathcal{G}$; (ii) for each cell (vertex or edge) c of $\mathcal{G}$, a group A_c; and (iii) for each oriented edge[2] e of $\mathcal{G}$ with terminal vertex v, an injective homomorphism $J_e\colon A_e \to A_v$, defined modulo inner automorphisms of A_v. An action without inversions of a group Γ on a

[2]Each edge of a graph corresponds to two oriented edges. Each oriented edge has a well-defined terminal vertex. If the edge is a loop, both orientations define the same terminal vertex.

simplicial tree T defines a graph of groups in a natural way. We set $\mathcal{G} = T/\Gamma$. A cell e of $\mathcal{G}$ is an orbit of simplices of T, and we define A_e to be the stabilizer of a simplex in this orbit; thus A_e is a subgroup of Γ defined up to conjugation. Any oriented edge e of $\mathcal{G}$ with terminal vertex v is represented by an oriented edge $\tilde{e}$ in T whose terminal vertex $\tilde{v}$ represents v. We define J_e to be the inclusion homomorphism from the stabilizer of $\tilde{e}$ to the stabilizer of $\tilde{v}$.

Bass and Serre showed that this construction gives a bijective correspondence between equivariant simplicial isometry classes of group actions without inversions on simplicial trees, on the one hand, and isomorphism classes of graphs of groups on the other. Furthermore, the inverse construction can be described explicitly. Thus the inverse correspondence assigns to each graphs of groups a group Γ, called the *fundamental group* of the give graph of groups, and an action of Γ on a tree T, called the *universal covering* tree. (These are not to be confused with the fundamental group and the universal covering tree of the underlying graphs.)

Given the bijectivity, one can of course *define* the fundamental group of a graph of groups to be the unique group which acts without inversions on a simplicial tree so that the given graph of groups appears as the quotient. However, Bass and Serre give an explicit and purely group-theoretical description of the fundamental group.

For example, if the given graph has one edge and two vertices v_1 and v_2, then there are two oriented edges e_1 and e_2, where e_i has terminal vertex v_i. Let us set $C = A_{e_1} = A_{e_2}$, and $A_i = A_{v_i}$ for $i = 1, 2$. Then $J_i = J_{e_i}$ is an injective homomorphism from C to A_i (defined modulo inner automorphisms). In this case, the fundamental group of the graph of groups is obtained from the free product $A_1 * A_2$ by adjoining the relations $J_1(\gamma) = J_2(\gamma)$ for all elements $\gamma \in C$. By definition this is the *free product of A_1 and A_2 with amalgamated subgroup C*, denoted $A_1 *_C A_2$. (The injective homomorphisms J_1 and J_2 are suppressed from the notation.)

The existence of the universal covering tree includes a number of classical facts about an arbitrary free product with amalgamation $\Gamma = A_1 *_C A_2$: in particular, the natural homomorphisms $A_i \to \Gamma$ are injective, so that the A_i can be identified with subgroups of Γ; and we have $A_1 \cap A_2 = C$. (This last statement is just the translation of the fact that the stabilizer of an edge in the universal covering tree is the intersection of the stabilizers of its endpoints.)

Another especially important example is a graph with one edge and one vertex v. In this case we have two oriented edges e_1 and e_2,two groups $A = A_v$ and $C = A_{e_1} = A_{e_2}$, and for $i = 1, 2$ a homomorphism $J_i = J_{e_i} : C \to A$. In this case the fundamental group is obtained from the free product of A with an infinite cyclic group $\langle t \rangle$ by adjoining the relations $tJ_1(\gamma)t^{-1} = J_2(\gamma)$ for all elements $\gamma \in C$. This is called an *HNN* (Higman-Neumann-Neumann) *extension with base group A and associated subgroup C* and is denoted $A*_C$.

In general, the fundamental group of a finite graph of groups can be calculated by successively using the two special constructions, free products with amalgamation and HNN extensions, that I have described above. Furthermore, the fundamental group of an infinite graph of groups can be obtained as the direct limit of the fundamental groups of its finite subgraphs.

If Γ is the fundamental group of a graph of groups, for each vertex (or edge) c the groups A_c has a natural identification—modulo inner automorphisms—with a subgroup of Γ. This again follows from the existence of the universal covering tree. A conjugate of A_c in Γ is called a *vertex group* (or edge group) of Γ.

The proofs given by Bass and Serre are combinatorial, rather like the proof I gave in 1.1.4 that a group acting freely is free. An alternative approach to their theory, using topology, was developed by Scott and Wall [ScW]. I discussed this topological approach in [Sh].

1.2. Metric trees. One useful way to think about a simplicial tree is to regard the set of vertices as forming a metric space T in which the distance function takes integer values: any two vertices x and y are joined by a unique simplicial arc, and the distance $\operatorname{dist}(x, y)$ is the number of 1-simplices that make up this arc. This provides an equivalence between simplicial trees and integer metric spaces satisfying certain conditions. When one writes the conditions down, they are seen to make sense not only for integer metric spaces, but for real metric spaces and—what is curious but important—even for metric spaces in a more general sense. This leads to the following definition, which was first given in [MSh1].

1.2.1. Let Λ be an ordered abelian group, denoted additively. A Λ-*metric space* is a set X equipped with a "distance function" $d \colon X \times X \to \Lambda$ which satisfies the usual formal axioms for a metric space. A *segment* in a Λ-metric space X is a subset which is isometric to a closed interval in Λ.

1.2.2. A Λ-*tree* is a Λ-metric space T with the following properties:

(i) Any two points $x, y \in T$ are the endpoints of a unique segment $[x, y]$;
(ii) For any $x, y, z \in T$ we have $[x, y] \cap [x, z] = [x, w]$ for some $w \in T$; and
(iii) If $[x, y] \cap [x, z] = x$ then $[x, y] \cup [x, z] = [y, z]$.

1.2.3. A **Z**-tree is just the 0-skeleton of a simplicial tree with the natural **Z**-metric that I described above. In this article I will use the terms "simplicial tree" and "**Z**-tree" interchangeably.

1.2.4. On the other hand, **R**-trees are more exotic objects. It was proved in [MSh1] that an **R**-metric space is an **R**-tree if and only if (i) any two points are the endpoints of a unique topological arc (i.e. a subspace homeomorphic to a closed interval) and (ii) every topological arc is isometric to a closed interval. (One could also replace (ii) by the condition that the space is a path space, i.e. that the distance between any two points is the infimum of the lengths of the paths joining them.) I will be mentioning some examples very soon.

1.2.5. When I talk about an action of a group Γ on a Λ- tree, I will always mean an action by isometries. This is the natural generalization of the convention that a group acting on a simplicial tree is understood to act simplicially. There is also a natural generalization of the "no inversions" condition. A group Γ is said to act *without inversions* on a Λ-tree T if whenever an element γ of Γ leaves a segment $[x, y]$ invariant, either (i) γ fixes x and y, or (ii) $\operatorname{dist}(x, y)$ is divisible by 2 in L. (In the latter case, $[x, y]$ has a midpoint z, and γ fixes z.) An action of a group on an **R**-tree is automatically without inversions since every real number is divisible by 2.

1.2.6. There is a natural generalization of the observation (1.1.4) that every free group acts freely and without inversions on its Cayley tree. Any free product Γ of copies of an arbitrary ordered abelian group Λ acts freely and without inversions on a Λ-tree. To see this, we recall that any element γ of Γ can be written uniquely as a reduced word; we define the *length* of γ, a non-negative element of Λ, to be the sum of the absolute values of the letters in the word. Then we can make Γ into a Λ-metric space by defining $\operatorname{dist}(x, y)$ to be the length of $x - y$. It's not hard to show that Γ is a Λ-tree and that the left regular action of Γ on itself is a free action by isometries having no inversions.

1.2.7. These remarks suggest one reason, not to be scoffed at, for introducing the general notion of Λ-tree: it provides a natural formal setting for the study of trees. Many arguments about simplicial trees go through equally well for Λ-trees, and the logic of an argument often becomes clearer if it is phrased in terms of Λ-trees.

Let me give another elementary illustration of this. It is a classical fact that any finite subgroup of a free product with amalagamation $A_1 *_C A_2$ is contained in a conjugate of one of the factors A_i. According to the Bass-Serre theory (1.1.5), this is a special case of the following fact about a **Z**-tree T: *any action without inversions of a finite group G on T has a fixed point.* Now this is not hard to prove, not only for a **Z**-tree, but for a Λ-tree, where Λ is any ordered abelian group.

Indeed, if x is any point of T, then there is a smallest subtree containing the orbit $G \cdot x$. Since $G \cdot x$ is finite, this subtree is finite in the sense that it is a finite union of segments. Any finite Λ-treeT_0 has a well-defined diameter $D = \max_{x,y \in T_0} \operatorname{dist}\,(x, y)$ and a well-defined *barycenter* m. We may define m to be the midpoint of any segment of length D in T_0; it is an exercise in using the tree axioms to show that any two such segments have the same midpoint. The barycenter of the smallest subtree containing $G \cdot x$ is obviously a fixed point for the action of G.

In [Sh] you will find other elementary arguments like the one above, for which Λ-trees seem to provide the natural formal context.

1.2.8. The construction that I described in 1.2.6 gives a free action of the group $\mathbf{R} * \mathbf{R}$ on an **R**-tree. As I said in [Sh], it is good to try to visualize this tree. This example shows how very infinite **R**-trees can be. On the other hand, it is really a completely tame example: it is simply the real analogue of the Cayley graph of $\mathbf{Z} * \mathbf{Z}$.

1.2.9. If $\Gamma = *_i \Gamma_i$ is a free product of subgroups of **R**, we get a free action of Γ on an **R**-tree by restricting the action described in 1.2.6.

It is not obvious whether these are the only groups that act freely on **R**-trees. I shall return to this question in Section 2.

1.2.10. An interesting case of the actions I just described in 1.2.9 is the one in which all the subgroups are infinite cyclic—but with arbitrary, possibly incommensurable, real generators. The actions that one obtains in this case are examples of what I call *polyhedral* actions. By definition, an action on an **R**-tree is polyhedral if it is topologically (but not necessarily metrically) equivalent to a simplicial action.

If you assign lengths to the edges of a graph $\mathcal{G}$, the universal cover $\tilde{\mathcal{G}}$ inherits an **R**-tree metric, and the action of $\pi_1(\mathcal{G})$ on $\tilde{\mathcal{G}}$ is polyhedral. All free polyhedral actions without inversions arise in this way. So there is nothing mysterious about such actions. And yet as I will be explaining in Section 2, the starting point for the fundamental work of Culler and Vogtmann on the outer automorphism group of a free group is to consider such actions. This is another good example of how the generalized notion of tree can be useful in a purely conceptual way.

1.2.11. A second reason for studying Λ-trees is that there are examples of actions of groups on Λ-trees which simplicial experience would not predict. For example, in the simplicial case it is only the free groups that act freely and without inversions. From the analogy in 1.2.6 one might guess that the only groups that act freely on **R**-trees are free products of subgroups of **R**. But this is false, as we shall see in 2.3. Similarly, the obvious free actions of a free group on **R**-trees are the polyhedral actions–and those derived from them in a trivial way[3]. But as we shall see in 2.2, there are free actions of free groups that are not of this type.

So an action of a group on a Λ-tree defines a genuinely new kind of structure on the group. And any new kind of structure provides an opportunity to recover a bit of order from the seemingly chaotic world of infinite groups.

1.2.12. A third reason for studying actions on Λ-trees is that they arise in applications, and not only as a formal device: exotic actions of the kind I mentioned in 1.2.11 come up in applications. In particular, the theory first introduced in [MSh1], and further developed in [MSh2], [M1], [MO], [Be], [Pau2], [Bru], [Bab] and [Ch], provides a connection between hyperbolic geometry and the study of Λ-trees. This allows one to reduce certain important questions about hyperbolic manifolds to questions about group actions on Λ-trees for suitable Λ. In this way it has been possible to re-prove and extend some fundamental results of Thurston's by proving appropriate theorems about Λ-trees; and certain conjectures about Λ-trees that I will be discussing in 2.5 and in Section 5 would allow one to extend these results much further.

I mentioned in 1.2.10 that the theory of Culler and Vogtmann begins by considering polyhedral actions on **R**-trees. But there are now important interactions, which I will discuss in 2.2, between this theory and the study of certain exotic actions. More generally, according to work of Paulin [Pau3] which I will briefly explain in 4.3.5, for a group Γ which is hyperbolic in the sense of the Cannon-Gromov theory, the study of the outer automorphism group $\mathrm{Out}(\Gamma)$ is closely related to the study of actions of Γ on trees.

These applications of dendrology are all aspects of a general theory that I will be discussing in Section 4.

1.3. Foundations and conventions. The remainder of this section was originally planned as a quick trip through some foundational material and conventions that I will be referring to later in the article. However, I will be stopping on the way to point out a couple of three-star views.

1.3.1. The geometric realization of a simplicial tree becomes an **R**-tree if we give each edge the linear metric of length 1. Restricting the **R**-tree metric to the 0-skeleton gives the **Z**-tree metric.

More generally, if $\Lambda_0 \subset \Lambda$ are ordered abelian groups then any Λ_0-tree T_0 has an embedding i in a Λ-tree T. The embedding i is an isometry of Λ_0-metric spaces, and every point of T lies in a segment with endpoints in $i(T_0)$. These properties are expressed by saying that T (or more precisely the pair (T, i)) is a Λ-*completion* of T_0. The completion is unique up to isometry: if (T', i') is another completion then there is a unique isometry $j \colon T \to T'$) such that $j \circ i = i'$. The existence and uniqueness of the Λ-completion were established in [AlpB], after some partial results in [MSh1]. We may write ΛT for the Λ-completion of T.

It follows from the uniqueness of the completion that any action of a group Γ on a Λ_0-tree T has a unique extension to ΛT. This makes it possible to think of actions on Λ_0-trees as being essentially a special case of actions on Λ-trees.

[3]I am referring here to actions for which the induced action on the minimal invariant sub-tree is polyhedral. See 1.5.

1.3.2. There is a natural generalization of the first barycentric subdivision of a simplicial tree. If Λ is any ordered abelian group, the order on Λ extends to an order on $\frac{1}{2}\Lambda = \Lambda \otimes_{\mathbf{Z}} (\frac{1}{2}\mathbf{Z}) \supset \Lambda$. If T is a Λ-tree then $\frac{1}{2}\Lambda T$ is a generalized first barycentric subdivision of T. Generalizing 1.1.3, one can show that the $\frac{1}{2}\Lambda$ completion of any group action on T is an action without inversions on $\frac{1}{2}\Lambda T$.

1.3.3. I am in the habit of saying that an action of a group Γ on a Λ-tree T is *trivial* if it has a fixed point, i.e. if there is a point of T which is fixed by all of Γ. This terminology is motivated by Bass-Serre theory. An action without inversions of a group Γ on a simplicial tree defines a *splitting* of Γ, i.e. an isomorphism between Γ and the fundamental group of a graph of groups. A splitting of a group is said to be *trivial* if some vertex group in the graph of groups corresponds to the entire group Γ under the isomorphism. In this sense, exhibiting Γ as a free product with amalgamation $A_1 \star_C A_2$ gives a splitting which is non-trivial if and only if the A_i are proper subgroups of Γ; this is true if and only if C is a proper subgroup of each A_i. On the other hand, exhibiting a group Γ as an HNN extension $A{*_C}$ always gives a non-trivial splitting of Γ. An action is trivial if and only if it corresponds to a trivial splitting.

Using Bass and Serre's group-theoretical definition of the fundamental group of a graph of groups, one can show that a group Γ admits a non-trivial splitting if and only if either (i) Γ admits a homomorphism onto $\mathbf{Z}$, or (ii) Γ is an amalgamated free product of two proper subgroups.

1.4. Trivial actions and deep mathematics. The use of the term *trivial* to describe an action with a fixed point should not obscure the fact that such actions can be extremely interesting. I would like to illustrate this by describing some beautiful work due to Gupta and Sidki [GuS].

A *torsion group* is a group whose elements are all of finite order. The classical Burnside problem asked whether a finitely generated torsion group must be finite. Golod [Go] gave a strong negative solution by showing that for every prime p there is a finitely generated infinite group which is a p-group in the sense that the order of every element is a power of p. Gupta and Sidki gave a marvelously simple and elementary proof of this for the case where p is odd by producing a two-generator subgroup of the automorphism group of a simplicial tree which is an infinite p-group.

For concreteness let us take $p = 3$. Consider the set $T = \bigcup_{n \geq 0} (\mathbf{Z}/3\mathbf{Z})^n$ of all finite sequences of elements of $\mathbf{Z}/3\mathbf{Z}$. We give T the structure of a simplicial tree by joining two vertices if and only if they are of the form $(r_1, \ldots, r_n)$ and $(r_1, \ldots, r_{n+1})$ for some $n \geq 0$ and some $r_1, \ldots, r_{n+1} \in \mathbf{Z}/3\mathbf{Z}$. (So T is the family tree of descendants of a Martian amoeba that reproduces by three-fold fission. See figure 1.4.1.)

Figure 1.4.1

There is an obvious order-3 automorphism t of T defined by $t(r_1, r_2, \ldots, r_n) = (r_1 + 1, r_2, \ldots, r_n)$. In the above picture, t just cyclically permutes the sub-trees lying below the points (1), (2) and (3). Now we define a second order-3 automorphism t as follows. Consider any vertex $v = (r_i)_{1 \leq i \leq n}$ of T. If $r_i = 0$ for $i = 1, \ldots, n-1$ we set $a(v) = v$. Otherwise let $k \leq n-1$ be the smallest index such that $r_k \neq 0$, and set $a(v) = (r'_i)_{1 \leq i \leq n}$, where $r'_i = r_i$ for $i \neq k+1$, and $r'_{k+1} = r_{k+1} + r_k$.

When we draw the automorphism t in the above picture, something surprising happens. It fixes the top vertex $\emptyset$ and each of the depth-one vertices (1), (2) and (3). Each depth-one vertex (r) is the top point of a sub-tree T_r which looks like the full picture of the tree T. For $r = \pm 1$ the automorphism t acts on T_r exactly as a^r acts on T. But the action of t on T_0 looks exactly like the action of t *itself* on the whole tree T!

This recurrent or self-referential property of t is the key to understanding the properties of the group $\Gamma \subset \mathrm{Aut}(T)$ generated by a and t. In particular it is easy to write down an isomorphism J between an index-3 subgroup of Γ and a subdirect product of three isomorophic copies of Γ. The existence of such an isomorphism J immediately implies that Γ is infinite. A simple but ingenious combinatorial argument, also using the isomorphism J, shows that Γ is a 3-group. (However, the orders of elements of Γ are not bounded: all powers of 3 occur as orders of elements.)

1.5. Foundations and conventions continued: length functions. Let us now turn our attention to non-trivial actions. It was shown in [CuM] (for $\Lambda \leq \mathbf{R}$) and independently in [AlpB] (for the general case) that if Γ acts non-trivially on a Λ-tree T then T has a unique minimal Γ- invariant Λ-subtree. Abstractly speaking, the basic problem in the subject is to classify minimal non-trivial actions up to equivariant isometry.

1.5.1. If a group Γ acts without inversions on an Λ-tree T, then $l(\gamma) = \min_{x \in T} \mathrm{dist}(x, \gamma \cdot x)$ exists for every $\gamma \in \Gamma$. This follows from an elementary argument given in [MSh1] and [AlpB], and discussed in [Sh]. The non-negative-valued function $l \colon \Gamma \to \Lambda$ is called the *length function* associated to the given action. It is obvious that l is constant on each conjugacy class, and that $l(\gamma) = 0$ if and only if γ has a fixed point in T.

1.5.2. The simplest length function on any group Γ is the one that is identically zero; this is realized by the action of Γ on the Λ-tree consisting of a single point. The next simplest kind of length function is of the form $l(\gamma) = |h(\gamma)|$, where h is any homomorphism of a group Γ into Λ. We may realize this as a length function by regarding Λ itself as being a Λ-tree, and letting Γ act on Λ by $\gamma \cdot \lambda = \lambda + h(\gamma)$. Length functions of this form are often called *abelian*.

1.5.3. Let Γ be a finitely generated group. It was proved in [MSh1] and [AlpB], and discussed in [Sh], that the length function associated to a given action of Γ on a Λ-tree T is identically zero if and only if the action is trivial. For a non-trivial action, the length function is determined by the restriction of Γ to its minimal invariant subtree.

1.5.4. Conversely, it was proved in [CuM] (for $\Lambda \leq \mathbf{R}$) and in [AlpB] (for the general case), and discussed in [Sh] that if two non-trivial actions determine the same length function then either (i) their minimal invariant subtrees are equivariantly isometric, or (ii) the length function in question is abelian.

1.6. Abelian length functions and the Bieri-Neumann Strebel invariant. In 1.4 I pointed out that trivial actions–those that define the zero length function–can be extremely interesting. Actions that define abelian length functions can likewise be very interesting. In [Bro], Brown relates such actions to the theory of the Bieri-Neumann-Strebel invariant [BiNS], which is a powerful tool in combinatorial group theory.

Geometrically, to say that a non-trivial action of a group Γ on an $\mathbf{R}$-tree T defines an abelian length function means that it has an invariant end. An *end* of T is by definition an equivalence class of rays in T (i.e. subtrees isometric to $[0, \infty) \subset \mathbf{R}$), where two rays are defined to be equivalent if their intersection is a ray. Thus an action defines an abelian length function if and only if there is a ray $r \subset T$ such that $\gamma \cdot r$ is equivalent to r for every $\gamma \in \Gamma$.

In the obvious case where the minimal invariant subtree is a line, there are clearly two invariant ends for the action. When the minimal invariant subtree is not a line, the invariant end is unique. In the latter case I shall say that the action is an *exceptional abelian action*.

1.6.1. Examples of exceptional abelian actions are provided by *ascending HNN extensions*. Let A be a group and let J be an isomorphism of A onto a proper subgroup C of A. Consider a graph of groups with one vertex and one edge, both labelled with the graph A; the homomorphisms $A \to A$ corresponding to the two orientations of the edge are the identity and J. The fundamental group Γ of this graph of groups is the HNN extension (1.1.5) obtained from a free product $A \star \langle t \rangle$ by adjoining the relations $t\gamma t^{-1} = J(\gamma)$ for all $\gamma \in C$. The Bass-Serre theory gives an action of Γ on a $\mathbf{Z}$-tree, which is an exceptional abelian action as one can check. A simple example is obtained by taking $A = C = \mathbf{Z}$ and $J(n) = 2n$; this gives the group $\langle x, y | yxy^{-1} = x^2 \rangle$.

1.6.2. By definition an abelian length function l on Γ has the form $\gamma \mapsto |h(\gamma)|$ for some homomorphism $h\colon \Gamma \to \mathbf{R}$. I'll write $l = |h|$. Given a non-zero abelian length function l, there are two different homomorphisms with absolute value l. However, an exceptional abelian action of Γ on T gives rise in a canonical way to a homomorphism whose absolute value is the length function defined by the action. In fact, if e is the unique invariant end, then for any $\gamma \in \Gamma$ there exist a number $\epsilon = \pm 1$ and a ray r representing the end e such that $\gamma^{\epsilon}(r) \supset r$. Furthermore, $\epsilon = \epsilon(\gamma)$ is uniquely determined by γ (and the given action of Γ on T); and we have dist $(x, \gamma \cdot x) = l(x)$, where x is the endpoint of r and l is the length function defined by the action. The homomorphism h canonically associated with the action is defined by $h(\gamma) = \epsilon(\gamma) l(\gamma)$.

1.6.3. Let Γ be a finitely generated group. The set Hom $(\Gamma, \mathbf{R})$ of all homomorphisms from Γ to $\mathbf{R}$ is a vector space over $\mathbf{R}$, whose dimension is the number n of infinite cyclic summands in the commutator quotient of Γ. The multiplicative group $\mathbf{R}^+$ of positive reals acts on Hom $(\Gamma, \mathbf{R}) - \{0\}$ by homotheties, and the quotient space Hom $(\Gamma, \mathbf{R}) - \{0\}$ is a sphere S of dimension $n-1$. Let $p\colon$ Hom $(\Gamma, \mathbf{R}) - \{0\} \to S$ denote the projection map. We have a dense set $p($Hom $(\Gamma, \mathbf{Q}) - \{0\}) = p($Hom $(\Gamma, \mathbf{Z}) - \{0\})$ of *rational* points in S.

Let us say that a non-zero homomorphism $h\colon \Gamma \to \mathbf{R}$ is *exceptional* if it is the homomorphism canonically associated to some exceptional abelian action of Γ on an $\mathbf{R}$- tree. The exceptional homomorphisms form a subset of Hom $(\Gamma, \mathbf{R}) - \{0\}$ which is invariant under the action of $\mathbf{R}$. The image of this set under p is a closed subset of S whose complement in S is denoted $\Sigma = \Sigma(\Gamma)$. The open set $\Sigma \subset S$ is the *Bieri-Neumann-Strebel invariant* of the group Γ. If the isomorphism type of Γ is specified then Σ is a subset of the $(n-1)$-sphere $\mathbf{R}^n/\mathbf{R}^*$ which is well-defined modulo the natural action of $\mathrm{GL}_n(\mathbf{Z})$.

1.6.4. To illustrate what can be done with this invariant, let me outline a proof of the following striking and unexpected result, which is proved in [BiNS].

(1.6.4.1) *Let Γ be a finitely presented group whose commutator quotient has at least two infinite cyclic summands. Then either Γ has a finitely generated normal subgroup N with $\Gamma/N \cong \mathbf{Z}$, or Γ has a free subgroup of rank* 2.

The logic of the proof is as follows. First one proves:

(1.6.4.2) Let Γ be a finitely generated group and let $h\colon \Gamma \to \mathbf{ZR}$ be a non-zero homomorphism such that $p(h) \in \Sigma \cap (-\Sigma)$. (Here $-\Sigma$ of course denotes the image of Σ under the antipodal map.) Then $\ker h$ is finitely generated.

Second, one proves:

(1.6.4.3) If Γ is a finitely presented group which contains no free subgroup of rank 2, then $\Sigma(\Gamma) \cup (-\Sigma(\Gamma)) = S$.

Given (1.6.4.2) and (1.6.4.3), one can quickly prove (1.6.4.1) as follows. If the commutator quotient of the finitely presented group Γ has at $n \geq 2$ infinite cyclic summands, then $\Sigma = \Sigma(\Gamma)$ is a subset of a sphere S of dimension $n-1 \geq 1$. In particular S is connected. If Γ does not contain a free subgroup of rank 2, then by (1.6.4.3), S is the union of the two open sets Σ and $-\Sigma$. By connectedness we have $\Sigma \cap (-\Sigma) \neq \emptyset$. Since $\Sigma \cap (-\Sigma)$ is open and the set of rational points $p(($Hom $(\Gamma, \mathbf{Z}) - \{0\})$ is dense in S, we have $p(h) \in \Sigma \cap (-\Sigma)$ for some non-zero homomorphism $h\colon \Gamma \to \mathbf{Z}$. By (1.6.4.2), $N = \ker h$ is finitely generated, and of course we have $\Gamma/N \cong \mathbf{Z}$.

In [Bro], Brown proves (1.6.4.2) from the dendrological point of view. We can assume that h is surjective. Let's fix an element $t \in \Gamma$ with $h(t) = 1$. Since h belongs to Σ, there is in particular no exceptional abelian action of Γ on a $\mathbf{Z}$-treewhose associated homomorphism is h. Translating this statement via Bass-Serre theory one concludes that Γ cannot be expressed as a properly ascending HNN extension with associated homomorphism h; that is, there is no proper subgroup B of $N = \ker h$ such that $tBt^{-1} \supset B$ and $N = \bigcup_{n \geq 0} t^n B t^{-n}$. Using this fact–and the hypothesis that Γ is finitely generated–it is elementary to conclude that N has a finitely generated subgroup C such that $N = \bigcup_{n \geq 0} t^{-n} C t^n$. This exhibits Γ as an ascending HNN extension with associated homomorphism $-h$; since $-h \in \Sigma$, this ascending HNN cannot be proper. Hence $C = N$, so that N is finitely generated.

A proof of (1.6.4.3) in the language of trees does not appear to have been written down, but one can translate the argument given by Bieri, Neumann and Strebel in roughly the following way. Since Γ is finitely presented, it can be realized as the fundamental group of a compact manifold M. Given any $h \in$ Hom $(\Gamma, \mathbf{R}) - \{0\}$, one can construct a piecewise-linear map $\tilde{f}$ from the universal cover $\tilde{M}$ of M to $\mathbf{R}$ such that $\tilde{f}(\gamma \cdot x) = \tilde{f}(x) + h(\gamma)$ for every $\gamma \in \Gamma$. Let N denote the universal abelian cover of M, i.e. the covering space corresponding to the commutator subgroup Γ' of $\Gamma = \pi_1(M)$. Then f induces a map $g\colon N \to \mathbf{R}$. Let $z \in \mathbf{R}$ be a regular value of g, so that $X_+ = g^{-1}[z,\infty)$ and $X_- = g^{-1}(-\infty, z]$ are manifolds-with-boundary of the same dimension as N, and $Y = g^{-1}(\{z\}) = \partial X_+ = \partial X_-$. For the moment let us assume that Y is connected. Then X_+ and X_- are connected, and $\Gamma' = \pi_1(N)$ is the free product of the groups $A_+ = \mathrm{im}(\pi_1(X_+) \to \pi_1(N)$ and $A_- = \mathrm{im}(\pi_1(X_-) \to \pi_1(N)$ amalgamated over the subgroup $C = \mathrm{im}(\pi_1(Y) \to \pi_1(N)$.

If h is exceptional then A_+ is a proper subgroup of Γ'. To prove this we consider an exceptional abelian action of Γ on an $\mathbf{R}$-tree T. We can take this action to be minimal. It is not hard to factor the map $\tilde{f}$ as $\alpha \circ \tilde{f}'$, where $\tilde{f}'\colon \tilde{M} \to T$ is a Γ-equivariant map and $\alpha\colon T' \to T$ satisfies $\alpha(\gamma \cdot y) = \alpha(y) + h(\gamma)$ for every $y \in T'$. Using that the action of Γ on T is exceptional, one can show that $P = \alpha^{-1}([z,\infty))$ is disconnected; this implies, using minimality, that $(\tilde{f}')^{-1}(P) = \tilde{f}^{-1}(Z)$ is disconnected. Thus X_+ has disconnected pre-image in the universal cover $\tilde{M}$ of N and therefore does not carry the fundamental group of N. This means that $A_+ \neq \Gamma'$, as asserted.

The same argument shows that if $-h$ is exceptional then A_- is a proper subgroup of Γ'. So if h and $-h$ are both exceptional, that is, if $p(h) \notin \Sigma \cup (-\Sigma)$, then Γ' is an amalgamated free product of two proper subgroups. By pushing this argument a bit further one can arrange that the amalgamated subgroup has infinite index in both factors. This implies that Γ' has a free subgroup of rank 2, contradicting the hypothesis. So we must have $h \in \Sigma \cup (-\Sigma)$.

I have sketched a proof that $h \in \Sigma \cup (-\Sigma)$ under the assumption that the set C, and hence the sets A_+ and A_-, are connected. In general this need not be the case. However, using the compactness of M (which reflects the hypothesis that Γ is finitely presented) one can show that there is always a unique component of A_+ whose image under g is an unbounded subset of $[0,\infty)$. The analogous assertion holds for A_-. Using this it is not hard to modify the argument so that it applies to an arbitrary $h \in \mathrm{Hom}\,(\Gamma, \mathbf{R}) - \{0\}$. This shows that $h \in \Sigma \cup (-\Sigma) = S$, as asserted in 1.6.4.3.

1.7. Foundations and conventions, concluded: coordinates for actions, and morphisms between trees. Apart from the fascinating ambiguity arising from abelian length functions, 1.5.4 says that minimal non-trivial actions are determined by their length functions. One can think of the length function associated with an action as defining coordinates for the action. This is a particularly useful point of view for the case $\Lambda = \mathbf{R}$.

1.7.1. To formalize it, we let $\mathcal{C}(\Gamma)$ denote the set of conjugacy classes in Γ, and we consider the Cartesian power $[0,\infty)^{\mathcal{C}(\Gamma)}$, where $[0,\infty)$ of course denotes the non-negative reals. We give $[0,\infty)^{\mathcal{C}(\Gamma)}$ the product topology and the subset $[0,\infty)^{\mathcal{C}(\Gamma)} - \{0\}$ the subspace topology. Any length function defined by a non-trivial action of Γ on an $\mathbf{R}$-tree may be regarded as a function on $\mathcal{C}(\Gamma)$ and may therefore be identified with a point of $[0,\infty)^{\mathcal{C}(\Gamma)} - \{0\}$. The set $\mathcal{L}(\Gamma)$ of all length functions for non-trivial actions is therefore identified with a subset of $[0,\infty)^{\mathcal{C}(\Gamma)} - \{0\}$. It turns out that this is a compact set; as I explained in [Sh], this was first proved in [CuM], and a more direct proof was given in [Par]. (I will mention yet another proof in 4.3.2.)

Life becomes even more pleasant if we work in the "projective space" $\mathcal{P}^\Gamma$ defined as the quotient of $[0,\infty)^\Gamma - \{0\}$ by the homothetic action

$$r \cdot (x_c)_{c \in \mathcal{C}(\Gamma)} = (r x_c)_{c \in \mathcal{C}(\Gamma)}$$

of the multiplicative group of positive reals. The space $\mathcal{P}^\Gamma$ is given the quotient topology, and the image $\mathcal{PL}(\Gamma)$ of $\mathrm{L}(\Gamma)$ in $\mathcal{P}^\Gamma$ is then a compact set. The points of $\mathcal{PL}(\Gamma)$ are called *projectivized length functions.*

1.7.2. Completing an action does not change its length function. More precisely, suppose that $\Lambda_0 \subset \Lambda$ are ordered abelian groups. If a group Γ acts without inversions on a Λ_0-tree T_0, then the completed action of Γ on $T = \Lambda T_0$ has no inversions, and it defines the same length function as the action of Γ on T. In particular the length function defined by the completed action takes values in L_0. Conversely, if Γ acts on a Λ-tree T and the length function defined by the action takes values in Λ_0, then the action is the completion of some action of Γ on a Λ_0-tree. These facts are proved in [AlpB].

1.7.3. It is often useful to think of Λ-trees, for any fixed Λ, as forming a category. Let T and T' be Λ-trees. A map of sets $f\colon T \to T'$ is called a *morphism* if each segment in T can be written as a finite union of subsegments, each of which is mapped isometrically into T' by f. (Thus a morphism crumples up any segment in at most a finite way.) An injective morphism is an isometry onto its image. If a morphism $f\colon T \to T'$ fails to be injective, it has a *fold*: that is, there are two segments in T with a common endpoint which are mapped isometrically by f onto the same segment in T'.

Section 2. Outer space, limits and exotic free actions

2.1. Outer space and its boundary; small actions. As an example of how the coordinatization of actions described in 1.7 can be used, consider the free group F_n of rank $n < \infty$, and let Y_n denote the subset of $\mathcal{PL}(F_n)$ consisting of all projectivized length functions defined by free actions of F_n on **R**-trees which are polyhedral in the sense of 1.2.10.

2.1.1. There is a natural action of the outer automorphism group $\mathrm{Out}(F_n)$ on Y_n given by $\alpha \cdot [l] = [l \circ \alpha^{-1}]$. Culler and Vogtmann [CuV1,2] have used this action to study the group $\mathrm{Out}(F_n)$ through properties of the space Y_n. For example, they have shown that Y_n is a contractible triangulable space and have computed its dimension. (The contractibility of Y_n was also proved independently by S. Gersten. I will be discussing a proof of the contractibility later, in 5.4.) The group $\mathrm{Out}(F_n)$ acts properly discontinuously on this space; in particular, the stabilizer of any point in Y_n is a finite subgroup of $\mathrm{Out}(F_n)$. (This is because if a point of Y_n corresponds to an action of F_n on a polyhedral tree T, the stabilizer of the point is isomorphic to a group of automorphisms of the quotient graph T/F_n, and the automorphism group of a finite graph is finite.)

Using the action of $\mathrm{Out}(F_n)$ on Y_n, Culler and Vogtmann have succeeded in calculating the cohomological dimension of $\mathrm{Out}(F_n)$. The same formalism can be used to calculate, for example, the virtual Euler characteristic of $\mathrm{Out}(F_n)$. This was done by Smillie and Vogtmann in [SmV].

I am pleased to note that my own term *outer space* for the space Y_n, reflecting its close connection with the group $Out(F_n)$, has been gaining currency.

The compactness result stated in 1.7.1 implies that the closure $\hat{Y}_n$ of Y_n in $\mathcal{PL}(F_n)$ is a compact space which is again invariant under $\mathrm{Out}(F_n)$. There is considerable evidence by now that deeper properties of outer automorphisms of F_n can be understood by studying their extended action on $\hat{Y}_n$.

2.1.2. The space $\hat{Y}_n$ does not consist entirely of length functions defined by free actions, because the length functions defined by free actions do not form a closed subset of $\mathcal{PL}(F_n)$. However, there is a somewhat weaker property than being free which is closed. Let us say that a group is *small* if it has no rank-2 free subgroup. An action of a group Γ on a Λ-tree is termed *small* if the stabilizer of every non-degenerate segment is a small subgroup of Γ. A (projectivized) length function is *small* if it is defined by some small action. In [CuM], Culler and Morgan proved that for any finitely generated group Γ which is not small, the small projectivized length functions on Γ form a closed (and hence compact) subset $\mathcal{SPL}(\Gamma)$ of $\mathcal{PL}(\Gamma)$. In particular we have $\hat{Y}_n \subset \mathcal{SPL}(\Gamma)$.

The only small subgroups of a free group are cyclic subgroups. So $\hat{Y}_n$ consists of length functions defined by actions with cyclic segment-stabilizers.

2.2. Exotic actions of free groups. The space $\hat{Y}_2$ has been completely described by Culler and Vogtmann in their paper [CuV2]. In particular, the small actions that define the points of $\hat{Y}_2 - Y_2$ are all well understood. It turns out that none of these actions is free. In fact, a result due to Harrison [H] and re-proved geometrically by Morgan [M2] implies that the only free actions of F_2 on **R**-trees are the free polyhedral actions, i.e. the actions whose length functions lie in Y_2.

By contrast, Bestvina and Handel have shown that for every $n \geq 3$ the set $\hat{Y}_n - Y_n$ does contain free actions. This is very striking because it shows that there exist free actions of F_3 on **R**-trees which are exotic in the sense that they are not polyhedral. (I understand that G. Levitt has also constructed exotic free actions of F_n for $n \geq 3$ from a different point of view.) This is one way in which the theory of actions on **R**-trees is genuinely different from the theory of actions on **Z**-trees.

2.2.1. The Bestvina-Handel construction is very natural because it uses the action of $\mathrm{Out}(F_n)$. For any $[l] \in Y_n$ and any $\alpha \in \mathrm{Out}(F_n)$, we can consider the sequence $(\alpha \cdot [l])$. For many choices of $\alpha \neq 1$, Bestvina and Handel show that—for an arbitrary point $[l] \in Y_n$—this sequence converges and its limit is a point of $[l_{\infty]} \in \hat{Y}_n$ which is defined by a free action. The point l_∞ is clearly fixed by α. It follows that l_∞ cannot lie in Y_n. Indeed, if we had $l_\infty \in Y_n$ then l_∞ would have a finite stabilizer by 2.1.1, so that α would have finite order. But in this case the sequence $(\alpha \cdot [l])$ would not converge.

2.2.2. Let me illustrate this by a concrete example, which was kindly provided to me by Mladen Bestvina. Consider the free group $F = F_3$ on the generators x, y and z. The Cayley graph T_0 of F with respect to the generators x, y and z is a simplicial tree, and its length function l_0 assigns to each element $\gamma \in F$ the length of a cyclically reduced word W in the conjugacy class of γ.

Let α be the automorphism of F defined by $\alpha(x) = xy$, $\alpha(y) = yz$ and $\alpha(z) = zxy$. For any cyclically reduced word W in F, we can of course calculate $\alpha(W)$ as a cyclically reduced word by replacing each generator in W by its image under α and making cyclic cancellations. (By a cyclic cancellation I mean either an ordinary cancellation, i.e. removing a subword of the form $u\bar{u}$ or $\bar{u}u$ where u is a generator and u is its inverse, or removing u from one end of the word and $\bar{u}$ from the other.) I'll say that a word W is *legal* if it is reduced and, in addition, it does not contain either $x\bar{z}$ or $z\bar{x}$ as a subword. If W is cyclically reduced and legal, and does not begin with $\bar{x}$ and ends with z, or begin with $\bar{z}$ and end with x, I'll say that W is *cyclically legal*. It is straightforward to check that if W is cyclically legal then no cyclic cancellations occur in calculating $\alpha(W)$, and what is more, that $\alpha(W)$ is itself a cyclically legal word.

For any cyclically legal word W, this makes it easy to study the behavior of $l_0(\alpha^n(W)$ as n increases. To do so let us associate to W the vector $v_W = (\xi, \eta, \zeta) \in \mathbf{R}^3$, where ξ denote the sum of the absolute values of the exponents of the generator x in W, and η and ζ are defined similarly in terms of the generators y and z. Then $l_0(W)$ is the sum of the coordinates of v_W. Since there are no cyclic cancellations involved in computing $\alpha(W)$, we have $v_{\alpha(W)} = A \cdot v_W$, where

$$A = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

Since, moreover, $\alpha(W)$ is cyclically legal, we can iterate this observation and conclude that $v_{\alpha^n(W)} = A^n \cdot v_W$ for every $n \geq 0$. So when W is cyclically legal, $l_0(\alpha^n(W)$ is simply the sum of the coordinates of $A^n \cdot v_W$.

Since A has non-negative entries, and A^2 has strictly positive entries, the Perron-Frobenius theory of positive matrices guarantees that A has a positive eigenvalue, and that if λ denotes its largest positive eigenvalue then, for any vector v in the open positive octant $(0, \infty)^3$, the sequence $(\lambda^{-n} A^n(v))_{n \leq 0}$ has a limit in $(0, \infty)^3$. It follows that $\lim_{n\to\infty} \lambda^{-n} l_0(\alpha^n(W)$ exists and is strictly positive for every non-trivial cyclically legal word W.

2.2.3. Remarkably enough, for every non-trivial element γ of F, there is a positive power α^k of α such that $\alpha^k(W)$ is conjugate to a cyclically legal word. This implies that $\lim_{n\to\infty} \lambda^{-n} l_0(\alpha^n(W)$ exists and is strictly positive for every non-trivial $\gamma \in F$. So the sequence $([l_0 \circ \alpha^n])_{\alpha \geq 0}$ has a limit $l_\infty \in \mathcal{SPL}(F)$, and furthermore we have $l_\infty(\gamma) \neq 0$ for every $\gamma \neq 1$. This means that l_∞ is the length function for a free action of F on an $\mathbf{R}$-tree. As I pointed out in 2.2.1, we have $l_\infty \notin Y_n$; that is, the free action defining l_∞ cannot be polyhedral.

2.2.4. The assertion of 2.2.3, that every $\gamma \neq 1$ is mapped to a conjugate of a cyclically legal word by some positive power of α, is proved by an elementary combinatorial argument. The key step is to show that if U and V are legal words then the element $\alpha^2(UV)$ is always represented by a legal word. In proving this, one can assume, after possibly replacing U and V by subwords, that the word UV is reduced. If UV is legal, then of course so are $\alpha(UV)$ and $\alpha^2(UV)$. If UV is reduced but illegal, we may assume that U ends with x and that V begins with $\bar{z}$. Let's write $U = U'x$ and $V = \bar{z}V'$. Then $\alpha(UV) = \alpha(U')\bar{z}\alpha(V')$. Since U' and V' are legal, $\alpha(U')$ and $\alpha(V')$ are legal, and moreover $\alpha(U')$ does not end with x. Hence if $\alpha(U')\bar{z}\alpha(V')$ is reduced then it is legal, so that the element $\alpha(UV)$ is represented by a legal word, and therefore so is the element $\alpha^2(UV)$.

So we may assume that $\alpha(U')\bar{z}\alpha(V')$ is not reduced. Now since V is reduced, V' does not begin with z, and hence $\alpha(V')$ does not begin with z either. So $\alpha(U')\bar{z}\alpha(V')$ can fail to be reduced only if the word $\alpha(U')$ ends with z; this means that U' must end with y, i.e. U must end with yx. By continuing this sort of analysis we see that U must in fact end with y^2x, and that V must begin with either $\bar{z}\bar{x}$ or $\bar{z}^2$. In the first case we can write $U = U''y^2x$ and $V = \bar{z}\bar{x}V''$. This gives $\alpha^2(UV) = \alpha^2(U'')yz^2\alpha^2(V'')$. Here, since $\alpha^2(U)$ is legal and in particular reduced, $\alpha^2(U'')$ is legal and does not end with $\bar{y}$. On the other hand, since $\alpha^2(V'')$ is the image of legal path under α, it does not begin with $\bar{z}^2$. Hence the word $\alpha^2(U'')yz^2\alpha^2(V'')$ representing the element $\alpha^2(UV)$ is legal, as required. The case in which V begins with $\bar{z}^2$ is handled similarly.

2.2.5. To prove the assertion of 2.2.3 we consider an element $\gamma \neq 1$ of F. Let n denote the smallest integer such that W is a product of n legal words. If $n > 1$, i.e. if γ is not represented by a legal word, it follows from 2.2.4 that α^2 is a product of fewer than n legal words. By induction it follows that $\alpha^{2k}(\gamma)$ is represented by a legal word. It is then easy to adapt the argument of 2.2.4 to conclude that the conjugacy class of $\alpha^{2k+2}(\gamma)$ is represented by a cyclically legal word.

2.2.6. Bestvina and Handel's general theory applies to any outer automorphism α of F_n which is *irreducible* in the sense that (say) no power of α leaves the conjugacy class of any free factor of F_n invariant. (Their definition of irreducibility is actually somewhat weaker than this.) For an irreducible α they prove that there exist a finite graph $\mathcal{G}$, an isomorphic identification of $\pi_1(\mathcal{G})$ with F_n and a map $f\colon \mathcal{G} \to \mathcal{G}$ inducing α, such that for any edge e of $\mathcal{G}$ and any $n \geq 0$, the map $\alpha^n|e\colon e \to \mathcal{G}$ is a reduced path, i.e. a locally 1-1 map of e into $\mathcal{G}$. Such a graph $\mathcal{G}$ is said to define a *train-track structure* for α. For example, for the automorphism α defined in 2.2.2 we can take $\mathcal{G}$ to be a three-leaf clover graph with one vertex and three edges. More generally, if an automorphism α is *positive* in the sense that it maps each generator to a positive word, then we can take $\mathcal{G}$ to be an n-leaf clover graph. Using the train-track structure and an analysis which is qualitatively similar to the one that I sketched in 2.2.2–2.2.5, Bestvina and Handel show that if for any irreducible outer automorphism α of F_n and any point $l \in Y_n$, the sequence $(\alpha \cdot [l])$ converges to a point $l_\infty \in \hat{Y}_n - Y_n$. Furthermore, they show that if no positive power of α fixes any conjugacy class in F_n then l_∞ is defined by a free action of F_n on an $\mathbf{R}$-tree.

The existence of a train-track structure for an irreducible automorphism has a number of other important consequences. In particular, Bestvina and Handel used it to prove a conjecture of G. P. Scott's, that if α is an *arbitrary* automorphism of F_n then the fixed elements of α form a subgroup of rank at most n.

2.2.7. The non-polyhedral free actions that I have described are by construction *limits* of free polyhedral actions, in the sense that the points of $\mathcal{SPL}(\Gamma)$ that they define are limits of sequences of points defined by free polyhedral actions. But it is easy to see that every free polyhedral action is a limit of simplicial actions: one need only approximate the lengths of the edges of the quotient graph by rational numbers. (This gives a $\frac{1}{m}\mathbf{Z}$-valued length function approximating the given polyhedral length function. But any $\frac{1}{m}\mathbf{Z}$-valued length function defines the same point of $\mathcal{SPL}(\Gamma)$ as some $\mathbf{Z}$-valued length function.) So the actions described in 2.2.1–2.2.6 are limits of simplicial actions.

2.3. Free actions of surface groups. I mentioned in 1.2.9 that the only groups that obviously admit free actions on $\mathbf{R}$-trees are free products of subgroups of $\mathbf{R}$. In 2.2 we saw that free groups, which of course belong to this class, admit surprising free actions. There are also examples of groups that are not free products of subgroups of $\mathbf{R}$ and which admit free actions. The first such examples, given by Alperin and Moss in [AlpM], are infinitely generated. In [MSh3], Morgan and I showed, by applying fundamental work of Thurston's, that the fundamental groups of most closed surfaces admit free actions on $\mathbf{R}$-trees.

2.3.1. As I pointed out in 2.2.7, the surprising free actions that I described for a free group F_n can be obtained as limits (in $\mathcal{SPL}(F_n)$) of free simplicial actions. General principles guaranteed that these limit actions would be small, and in fact they were often free. In the case of a surface group Γ, one can obtain surprising actions, including free actions, as limits of small simplicial actions. In this case the simplicial actions cannot be free, since the group is not free. Nevertheless, the limit actions are often free.

2.3.2. So I will begin by talking about some examples of small actions on $\mathbf{Z}$-trees. (Since I also described these examples in [Sh], I will be brief.) Let Σ be a closed surface and let $C \subset \Sigma$ be a disjoint union of two-sided simple closed curves. Then the universal cover $\tilde{\Sigma}$ of Σ is a simply-connected surface, and the pre-image $\tilde{C}$ of C in $\tilde{\Sigma}$ under the covering map $p\colon \tilde{\Sigma} \to \Sigma$ is a properly embedded 1-manifold. There is a dual graph of $\tilde{C}$ in $\tilde{\Sigma}$ whose vertices and edges correspond to components of $\tilde{\Sigma} - \tilde{C}$ and $\tilde{C}$ respectively; an edge e is incident to a vertex v if and only if the component of $\tilde{C}$ corresponding to e is contained in the closure of the component of $\tilde{\Sigma} - \tilde{C}$ corresponding to v. Since $\tilde{\Sigma}$ is 1-connected, it is not hard to show that the dual graph is 1-connected as well. The action of $\Gamma = \pi_1(\Sigma)$ on $\tilde{\Sigma}$ induces an action on this dual graph, and hence on a $\mathbf{Z}$-tree. We call this action the *dual action* to the curve system C.

It is easy to show that the dual action is minimal if and only if C is reduced in the sense that it is non-empty and has no homotopically trivial components. Under the dual action, the stabilizer of a vertex is the stabilizer of the corresponding component of $\tilde{\Sigma} - \tilde{C}$, which up to conjugacy is the image of the fundamental group of a component of $\Sigma - C$ in $\pi_1(\Sigma)$. Likewise, the stabilizer of an edge is (up to conjugacy) the image of the fundamental group of a component of C in $\pi_1(\Sigma)$. In particular, the stabilizer of every edge is an infinite cyclic group. It follows that the action dual to a reduced curve system is small.

2.3.3. An elementary argument, given in [MSh1] and essentially due to Stallings, shows that every small minimal action of a surface group on a **Z**-treecan be obtained by the above construction.

2.3.4. As I said in 2.3.1, one can obtain interesting small actions of surface groups on **R**-trees, including free actions, as limits of small actions on **Z**-trees. This is the way these actions were described in [Sh]. However, in order to get a deeper understanding of these actions it is necessary to give a more direct description of them. They can in fact be described by generalizing the construction given in 2.3.2. This was done in [MSh3], and from a somewhat different point of view in Section 5 of [GiS1]. Here I will try to give a self-contained account of the approach of [GiS1], which involves Thurston's theory of measured foliations (see [FaLP]). First it will be helpful to look at reduced curve systems a bit more closely.

2.3.5. A reduced curve system C can be described by combinatorial data in terms of a triangulation of Σ. In fact, given C one can construct a triangulation of Σ which is adapted to C in the sense that C intersects each 2-simplex σ in the way shown in Figure 2.3.5.1.

Figure 2.3.5.1

In somewhat fewer than a thousand words, this means that C meets each edge of σ, and that we can label the edges of σ as τ , τ' and τ'' in such a way that every component of $C \cap \sigma$ is an arc joining a point of tau to a point of τ' or τ''. In particular, C then intersects each 1-simplex τ in a set of finite cardinality $x_{tau} > 0$. The positive integers x_τ clearly have the property that

(i) for any 2-simplex σ of Σ we may label the edges of σ as τ , τ' and τ'' in such a way that $x_\tau = x_{\tau'} + x_{\tau''}$.

A family (x_τ) of positive integers indexed by the 1-simplices of a given triangulation will be called an integer *length system* for Σ if it satisfies condition (i). We call τ the *long edge* of σ relative to the given length system. For each vertex v of Σ, the *order* o_v of v with respect to the length system is defined to be the number of 2-simplices incident to v whose long edge is not incident to v. I'll say that a length system is *non-degenerate* if

(ii) $o_v > 2$ for every vertex v of Σ.

2.3.6. Given a reduced curve system, it is possible to choose a triangulation which is adapted to C and such that the corresponding length system is non-degenerate. Conversely, it is not hard to show that every non-degenerate integer length system is defined by some reduced curve system. So one can think of the small actions on **Z**-trees described in 2.3.2 as being associated with **Z**-valued length systems. This suggests defining small actions on **R**-trees by means of *real-valued length systems.*

The definition of a real-valued length system is the same as that of an integer-valued length function except that the x_τ are positive *real* numbers. The definition of the long edge of a 2-simplex and the order of a vertex go through without change, and so does the definition of non-degeneracy. Now let me explain how a non-degenerate real-valued length system gives rise to an action of $\pi_1(\Sigma)$ on an **R**-tree.

2.3.7. To construct the tree we note that for each 2-simplex σ there is an affine map $f_\sigma\colon \sigma \to \mathbf{R}$ which maps each edge τ of σ onto an interval of length x_τ. Such a map f_σ is unique up to composition with isometries of $\mathbf{R}$. We define the *length* $\lambda(\alpha)$ of an affine path $\alpha\colon [0,1] \to \sigma$ to be the length of the interval $f_\sigma(\alpha([0,1]))$. A piecewise-linear path α in Σ can be written as a composition $\alpha_1 * \cdots * \alpha_n$ of affine paths in 2-simplices, and it has a well-defined length $\lambda(\alpha) = \lambda(\alpha) + \cdots + \lambda(\alpha)$.

The length of a path $\tilde{\alpha}$ in the universal cover $\tilde{\Sigma}$ is defined to be the length of $p \circ \tilde{\alpha}$.

It can be shown that any two points x and y of $\tilde{\Sigma}$ are joined by a path of minimal length. The existence of such a minimal path is not at all obvious. It depends strongly on condition (ii) of 2.3.5; it also depends on the compactness of the surface Σ, which guarantees the existence of a group of simplicial homeomorphisms of $\tilde{\Sigma}$ which has compact quotient and preserves the induced length system of $\tilde{\Sigma}$.

If we denote the minimal length of a path from x to y by dist (x, y), it follows that *dist* is a pseudo-metric on $\tilde{\Sigma}$. This pseudo-metric gives rise to a metric space T by the standard construction: we say that two points $x, y \in \tilde{\Sigma}$ are *equivalent* if $\mathrm{dist}(x, y) = 0$; the points of T are equivalence classes and the distance function on T is induced by dist.

It may be shown that T is an $\mathbf{R}$-tree. The action of $\pi_1(\Sigma)$ on Σ by deck-transformations induces an action (by isometries) on T. The stabilizer of every non-degenerate segment is cyclic. When the given length system is integer-valued, this action is the completion of the action described in 2.3.2 in terms of the corresponding reduced curve system.

2.3.8. The above description of the tree associated to a real-valued length system is a paraphrase of the description given in [GiS1] in terms of Thurston's theory. His theory is needed to prove the assertions that I made in 2.3.7, including the existence of a minimal path between two points in $\tilde{\Sigma}$. Geometrically, the points of T are *leaves* of a *foliation with singularities* of $\tilde{\Sigma}$. Each leaf is locally Euclidean of dimension 1 except at the vertices. The foliation of $\tilde{\Sigma}$ induces a foliation of Σ.

A neighborhood of a vertex v in a leaf in Σ is a cone on o_v points. The assignment of a length to each piecewise-linear path may be interpreted as a *transverse measure* on the foliation. See Figure 2.3.8.1 for the local picture.

Figure 2.3.8.1

The vertices v such that $o_v \neq 2$ are *singularities* of the foliation. Condition (ii) of 2.3.5 says that the singularities are all *non-degenerate*, i.e. that the situations shown in Figure 2.3.8.2 never occur.

The action of $\pi_1(\Sigma)$ on T is often said to be *dual* to the measured foliation determined by the given length system.

2.3.9. Let me use the geometric picture to illustrate the comment that I made in 2.3.7, that the restrictions on the induced length system of $\tilde{\Sigma}$, including invariance under a group of simplicial homeomorphisms of $\tilde{\Sigma}$ which has compact quotient, are needed to prove the existence of a minimal path between two points of $\tilde{\Sigma}$. Let $f\colon (-1,1) \to \mathbf{R}$ be a piecewise-linear continuous even function which tends to $+\infty$ as $x \to \pm 1$. There is a foliation of $\mathbf{R}^2$ with no singularities, whose leaves are the lines $x = c$ for $|c| \geq 1$ and the curves $y = f(x) + b$ for $b \in \mathbf{R}$.

Figure 2.3.8.2

This foliation has a unique transverse measure for which the length of every vertical line segment in the y-axis, and of every horizontal line segment in the complement of $(-1, 1) \times \mathbf{R}$, is equal to its Euclidean length. It is possible to realize the resulting measured foliation by a length function on a triangulation.

But now look at the points $P = (0, -1)$ and $Q = (0, -1)$. For any ϵ with $0 < \epsilon < 1$ we can join P and Q by an arc consisting of the line segments $[-1, -1 + \epsilon] \times \{0\}$ and $[1 - \epsilon, 1] \times \{0\}$ and an arc in the leaf. As $\epsilon \to 0$ the length of this arc (with respect to the measure) tends to 0. On the other hand, P and Q lie in different leaves, so there is no path of length 0, and hence no minimal path, joining them. What's more, there is no reasonable way of associating a tree with this measured foliation.

2.3.10. If Σ is a compact orientable surface of positive genus, it may be shown that there always exsits a non-degenerate length system for some triangulation of Σ, and that for a generic choice of such a length system, the corresponding action of $\pi_1(\Sigma)$ on an $\mathbf{R}$-tree is free. Furthermore, this construction, applied with a bit more care, gives free actions of the fundamental groups of all non-orientable surfaces except those of Euler characteristic 1, 0 and -1. These facts were established in [MSh3], although from a different point of view.

2.3.11. Let (x_τ) be a non-degenerate length system on Σ, and let $[l] \in \mathcal{SPL}(\pi_1(\Sigma))$ be the projectivized length function defined by the action of $\pi_1(\Sigma)$ on the associated $\mathbf{R}$-tree. It is a matter of linear algebra to show that we can approximate the numbers x_τ arbitrarily well by positive rational numbers x'_τ which define a length system in the same triangulation. After multiplying the x'_τ by a suitable integer we obtain an integer length system which defines an action on a $\mathbf{Z}$-treeand a corresponding projectivized integer-valued length function $[l'] \in \mathcal{SPL}(\pi_1(\Sigma))$. If we choose the approximation x'_τ with a bit of care, it may be shown[4] that $[l']$ approaches $[l]$ as x'_τ approaches x_τ. This shows that the small action defined by the length system (x_τ) is, as I promised, a limit of actions of the type defined in 2.3.2.

[4]This is not at all obvious. It is equivalent to a result due to Thurston. You may find a proof of a similar result in a somewhat more general context in [GiSSk].)

2.3.12. The actions of surface groups on **R**-trees that are defined by non-degenerate length systems—or, if you prefer, are dual to measured foliations—are of course generally not polyhedral; nevertheless, these actions have a striking finiteness property. Let us define a *partition* of a topological arc S to be a finite set $\mathcal{S}$ of sub-arcs of S such that S is the union of the segments in $\mathcal{S}$, and any two segments in $\mathcal{S}$ meet in at most a single point. If an action of a surface group Γ on a tree T is dual to a measured foliation, then for any segment $S \subset T$ there exist a partition $\mathcal{S}$ of S and an indexed family $(\gamma^s)_{s\in\mathcal{S}}$ such that the set $\{\gamma^s(s) \colon s \in \mathcal{S}\}$ is again a partition of S.

Note that each γ^s maps $s \subset S$ isometrically onto $\gamma^s \subset S$. So if we identify S isometrically with an interval in $\mathbf{R}$, each of the maps $\gamma^s|s$ is the restriction to s of a translation or a reflection in $\mathbf{R}$. If S' denotes the complement in S of the set of all endpoints of intervals in $\mathcal{S}$, we can define a measure-preserving map $\phi \colon S' \to S$ by setting $\phi|s = \gamma^s|s$ for each $s \in \mathcal{S}$. This map ϕ, is made up of a finite number of isometries. More precisely, its domain is the complement of a finite set of points in S, each component of the domain is mapped isometrically onto an interval in S, and the closure of these image intervals form a partition of S. Such a map is called an *interval exchange* on the topological arc S.

2.3.13. The property of a dual action described in 2.3.12 is a consequence of the geometric theory of measured foliations. If the given segment S has a partition into subsegments for which the assertion is true, then it is true for S as well. This observation allows one to assume that there is an arc $\tilde{A} \subset \tilde{\Sigma}$ whose interior avoids the singularities of the foliation and is transverse to the leaves, and such that the quotient map $\tilde{\sigma} \to T$ maps $\tilde{A}$ homeomorphically onto S. Likewise one can assume that the covering projection maps $\tilde{A}$ homeomorphically onto an arc A in Σ. So we have a natural homeomorphism between A and S.

Let us fix a transverse orientation for A in Σ. For a generic point $x \in A$, there is a unique arc α_x such that (i) α_x is contained in a leaf of the foliation and contains no singular points, (ii) $x \in \partial\alpha_x = \alpha_x \cap S$, and (iii) a small neighborhood of x in α_x lies on the positive side of A (with respect to the transverse orientation). This follows from the compactness of Σ and the existence of the transverse measure. Let $\phi(x)$ denote the endpoint $\neq x$ of α_x. It is not hard to show that the *Poincaré first return map* ϕ is an interval exchange map. By identifying S with A via the natural homeomorphism we can regard ϕ as an interval exchange on S. Furthermore, for each point $x \in \text{dom } \phi$, the arc α_x may be thought of a as a path from x to $\phi(x)$, and therefore determines an element of $\Gamma = \pi_1(\Sigma, S)$. This element depends only on the component of dom ϕ containing x, so we can denote it γ_s. With these definitions of ϕ and γ_s it is straightforward to verify the assertion of 2.3.12.

2.4. Classification of small actions of surface groups. Richard Skora has recently proved that if Σ is a closed surface and if $\pi_1(\Sigma)$ acts minimally on an **R**-tree in such a way that the stabilizer of every non-degenerate segment is cyclic, then the given action is associated to a non-degenerate length system by the construction that I described in 2.3.7. In more invariant terms this says that every action of $pi_1(\Sigma)$ on an **R**-treewith cyclic segment stabilizers is dual to a measured foliation on Σ. This result, which is the natural generalization of the elementary fact mentioned in 2.3.3 above, gives an affirmative answer[5] to [Sh, Question C].

I would like to give a sketch of Skora's proof, but using a point of view slightly different from Skora's and suggested by Culler and Vogtmann's paper [CuV2].

2.4.1. The starting point for Skora's proof is an idea that was first introduced by Stallings for **Z**-trees, was applied for general **R**-trees in [MSh2], and was further developed in [MO]. Suppose that we are given an arbitrary action of the fundamental group of a surface Σ on an **R**-tree$T$. Let us fix any triangulation of Σ, and let the universal cover $\tilde{\Sigma}$ be given the triangulation induced by the first barycentric subdivision. Then $\Gamma = \pi_1(\Sigma)$ acts on $\tilde{\Sigma}$ by deck transformations, and it is not hard to construct a Γ-equivariant map $\tilde{f} \colon \tilde{\Sigma} \to T$. Furthermore, $\tilde{f}$ can be constructed so as to map each 1-simplex of Σ onto a segment in T, and so that each 2-simplex σ of Σ has an edge τ such that $\tilde{f}(\tau) = \tilde{f}(\sigma)$. The other two edges τ' and τ'' of σ are mapped onto

[5]In [Sh], I discussed the natural small actions of surface groups in terms of measured *laminations.* Here I am discussing them in terms of measured *foliations.* These are two alternative ways of thinking about a class of objects introduced by Thurston. The equivalence between the two points of view is not obvious, however. Measured laminations were used in [MSh3], and measured foliations in [GiS1].

subsegments whose union is $\tilde{f}(\tau)$. In particular we have

$$\text{length}(\tilde{f}(\tau)) = \text{length}(\tilde{f}(\tau')) + \text{length}(\tilde{f}(\tau'')).$$

Hence we can define a measured foliation with (possibly degenerate) singularities by assigning to each 1-simplex τ of Σ the number $x_{tau} = \text{length}(\tilde{f}(\tau))$. The leaves of this foliation are the connected components of the pre-images under $\tilde{f}$ of points of T. The length of a path α in $\tilde{\Sigma}$ is the length of the path $\tilde{f} \circ \alpha$ in T. This measured foliation is $\pi_(\Sigma)$-invariant on account of the equivariance of f, and hence induces a measured foliation on Σ.

One of the results proved by Morgan and Otal in [MO] (although stated in different language) is that if the action of $\pi_1(M)$ on T is non-trivial, then by modifying the map $\tilde{f}$ one can arrange that the foliation of Σ that it defines has non-degenerate singularities. When this has been done, the measured foliation has a dual tree T_0, and $\tilde{f}$ induces a Γ-invariant morphism $\phi\colon T_0 \to T$. (Morphisms of trees were defined in 1.7.3.)

2.4.2. None of this involves the hypothesis that the action of Γ on T is minimal, or—more important—that it is small. When these hypotheses hold, Skora shows that ϕ is an isomorphism and hence that the given action of Γ on T is dual to a measured foliation. Surjectivity follows from the minimality of the action. So one needs only to prove that if the action of Γ on T is small then ϕ is injective, which by 1.7.3 amounts to proving that it doesn't fold. The proof is slightly simpler under the stronger hypothesis that Γ acts freely on T, and it is this case that I will discuss.

Suppose that ϕ does fold, so that there are two segments $S_1, S_2 \subset T_0$ which meet in a common endpoint and are mapped isometrically by ϕ onto the same segment $S \subset T$. According to 2.3.12 there exist for $i = 1, 2$ a partition $\mathcal{S}_i$ of S_i and an indexed family $(\gamma^s)_{s \in \mathcal{S}_i}$ such that the set $\{\gamma^s(s)\colon s \in \mathcal{S}_i\}$ is again a partition of S_i. As in 2.3.12 we get an interval exchange map E_i on each S_i. Conjugating E_1 and E_2 by the isometry $\phi_i = \phi|S_i\colon S_i \to S$, we get two interval exchanges E'_1 and E_2 on S. Now let x be a generic point of S. For each $n \geq 0$ and each $I = (i_1, \ldots, i_n) \in \{0,1\}^n$, we can define $x_I = E'_{i_n} \circ \cdots \circ E'_{i_1}(x)$. From the definition of the E'_i and the equivariance of ϕ you can see that for each $I \in \{0,1\}^n$ there is an element γ_I of Γ such that $x_I = \Gamma_I(x)$. Furthermore, γ_I is a word of length at most n in the elements $\gamma^s, s \in \mathcal{S}_1 \cup \mathcal{S}_2$.

Using the geometry of measured foliations one can arrange, after possibly replacing the group Γ by a finite-index subgroup, and if necessary replacing S by a shorter segment, that the function $I \to \gamma_I$ from $\{0,1\}^n$ to Γ is one-to-one. In this case, the cardinality of the set $\mathcal{G}_n = \{\gamma_I\colon I \in \{0,1\}^n$ is an exponentially growing function of n.

On the other hand, since an interval exchange is made up of finitely many translations and reflections, the cardinality of the set $X_n = \{x_I\colon I \in \{0,1\}^n$ has (at most) polynomial growth as a function of n. Hence when n is large enough, there are two elements $I, I' \in \{0,1\}^n$ such that $\gamma_I \neq \gamma_{I'}$ but $x_I = x_{I'}$, i.e. $\gamma_I(x) = \gamma_{I'}(x)$. This contradicts the hypothesis that the action is free.

2.5. Conjectures and Questions. We saw in 1.2.9 that any free product of subgroups of $\mathbf{R}$ acts freely on an $\mathbf{R}$-tree. A similar construction shows that the class of groups that act freely on $\mathbf{R}$-trees is closed under the formation of free products. In particular a free product of surface groups and free abelian groups admits a free action on an $\mathbf{R}$-tree provided that no factor has the form $\pi_1(\Sigma)$ where Σ is non-orientable and has Euler characteristic > -1.

2.5.1. In [Sh, Question B] I asked whether the following statement was true.

Every finitely generated group which acts freely on an $\mathbf{R}$*-tree is a free product of free abelian groups and surface groups.*

As there is now a great deal of evidence for the truth of this statement, I think it deserves to be called a conjecture. I will discuss some of the evidence in Section 5.

2.5.2. Since it is relatively elementary to show that an action of a surface group $\pi_1(\Sigma)$ is defined by a measured foliation on Σ if and only if it is a limit of small simplicial actions, we can interpret Skora's theorem (2.4) as saying that $\mathcal{SPL}(\pi_1(Sigma))$ has a dense subset consisting of projectivized integer-valued length functions. More generally, for any finitely generated group Γ, we can ask:

Does $\mathcal{SPL}(\Gamma)$ *have a dense subset consisting of projectivized integer-valued length functions?*

2.5.3. A weaker form of Question 2.5.2 was asked as Question D of [Sh]:

If the finitely generated group Γ admits a non-trivial small action on an **R***-tree, does it admit a non-trivial small action on a* **Z***-tree?* By the Bass-Serre theory this is equivalent to asking whether Γ can be exhibited as either a free product of two proper subgroups with a small amalgamated subgroup, or an HNN extension with a small associated subgroup.

2.5.4. As I explained in [Sh] and will explain in more detail in Section 4 below, Question 2.5.3 is especially important for applications. It is also closely related to the conjecture about groups that act freely. In fact, Morgan and Skora [MSk] have shown that if Γ is a finitely presented group which is indecomposable (i.e. is not a non-trivial free product), and if Γ admits both a small action on an **Z**-treeand a free action on an **R**-tree, then Γ is either a free abelian group or a surface group. It follows easily from this result that if Question 2.5.3 has an affirmative answer for finitely presented groups, then Conjecture 2.5.1 is true with the additional hypothesis of finite presentation.

2.5.5. There is considerable evidence that Question 2.5.2 has an affirmative answer under a slightly stronger hypothesis. I will state this here as a conjecture:

Let Γ be a finitely presented group whose small subgroups are all finitely generated. Then $SPL(\Gamma)$ has a dense subset consisting of projectivized integer-valued length functions.

2.5.6. I will be discussing the evidence for Conjectures 2.5.1 and 2.5.5 in Section 5.

Actually, the evidence points to a stronger conjecture which implies both the above conjectures, and which I find very natural and appealing in its own right. As the statement requires a bit of background I will defer it to Section 5.

2.5.7. Question 2.5.2 has an appealing analogue for arbitrary (not necessarily small) non-trivial actions:

Let Γ be a finitely generated group. Does $\mathcal{PL}(\Gamma)$ have a dense subset consisting of projectivized integer-valued length functions?

2.5.8. An affirmative answer would imply an affirmative answer to the following question, which was asked in [Sh] as Question A: *If Γ admits a non-trivial action on an* **R***-tree, does it admit a non-trivial action on a* **Z***-tree?* By the Bass-Serre theory this is equivalent to asking whether Γ admits a non-trivial splitting (1.3.3).

The evidence for an affirmative answer to Questions 2.5.7 and 2.5.8 is rather weak so far. I shall discuss it in Section 5.

2.6. Non-Archimedean trees. The questions that I have been discussing in this section for **R**-trees, and to which I will return in Section 5, take on a rather different flavor if we consider actions on Λ-trees where Λ is not a subgroup of **R**. Such actions, besides being a natural object of study in their own right, arise in connection with Tits buildings for general valued fields (Section 3) and degeneration of hyperbolic structures (see Section 4, especially 4.1.7, 4.2.3 and 4.4.1).

2.6.1. An arbitrary ordered abelian group Λ can be analyzed in terms of its convex subgroups. A subgroup Λ_0 of Λ is said to be *convex*, or *isolated*, if for any $x, y, z \in \Lambda$ such that $x \leq y \leq z$ and $x, z \in \Lambda_0$, we have $y \in \Lambda_0$. The group Λ is order-isomorphic to a subgroup of **R** if and only if it has no non-trivial proper convex subgroups. The simplest example of a group Λ which does have a non-trivial proper convex subgroup is the group $\Lambda = \mathbf{Z} \times \mathbf{Z}$ with the lexicographical order; here the unique non-trivial proper convex subgroup is $\{0\} \times \mathbf{Z}$.

In the general case, the convex subgroups of Λ form a linearly ordered set under inclusion. In most casea that arise in applications, there are only finitely many convex subgroups. The number of non-trivial convex subgroups (including Λ itself) is the *order-rank* of Λ. If Λ has order-rank n and its convex subgroups are $\{0\} = \Lambda_0 \leq \Lambda_1 \leq \cdots \leq \Lambda_n = \Lambda$, the quotient groups Λ_i/Λ_{i-1} have induced orders, and are order-isomorphic to subgroups of **R**.

2.6.2. Bass has shown in [Bas3] that it is largely possible to reduce the study of group actions on Λ-trees to the study of actions on Λ_i/Λ_{i-1}-trees. In particular, when the groups Λ_i/Λ_{i-1} are all cyclic, as in the case $\Lambda = \mathbf{Z}^n$ with the lexicographical order, these actions can be analyzed via the Bass-Serre theory. The results are quite different from the corresponding results for **R**-trees. For example, any free product with amalgamation of the form $F_1 *_C F_2$, where the F_i are free and C is an infinite cyclic group which is identified inder the amalgamation with a maximal cyclic subgroup of each F_i, admits a free action without inversions on a $\mathbf{Z} \times \mathbf{Z}$-tree, where $\mathbf{Z} \times \mathbf{Z}$ has the lexicographical order. By contrast, it was shown by Morgan in [M2], and is included in the results of Morgan and Skora discussed in 2.5.4, that a group of this type cannot act freely and without inversions on a $\mathbf{Z}$-tree unless it is a free product of cyclic groups and surface groups, as predicted by Conjecture 2.5.1.

Section 3. The Bruhat-Tits tree and its applications

3.1. Valuations and Serre's construction. Many of the applications of dendrology involve the Bruhat-Tits building for PSL_2 of a valued field. In this section I am going to describe the Bruhat-Tits construction from a point of view that is essentially due to Serre, who presented it in his book [Se] for the case $\Lambda = \mathbf{Z}$. The general case was presented in my paper [MSh1] with John Morgan. In this section and in 4.1–4.2, I will talk about some applications of this construction to the study of matrix groups, group representations, 3-manifold topology and hyperbolic geometry.

3.1.1. Let me start with a standard definition from commutative algebra. Let K be a field. A *Krull valuation* of K is a homomorphism v of the multiplicative group $K^* = K - \{0\}$ onto an ordered abelian group *Lambda*, called the *value group*, such that for any $x, y \in K^*$ with $x + y \neq 0$ we have

$$v(x + y) \geq \min(v(x), v(y)).$$

To understand this definition, one should bear two basic examples in mind. In both examples we have $\Lambda = \mathbf{Z}$.

First, let $K = \mathbf{Q}$. Any prime p defines a Krull valuation of $\mathbf{Q}$. For $0 \neq a \in \mathbf{Z}$ we define $v(a)$ to be the exponent of$p$ in the prime factorization of a. Then v is a map from $\mathbf{Z} - \{0\}$ to the non-negative integers, and it has a well-defined extension to a map $v\colon \mathbf{Q}^* \to \mathbf{Z}$ given by $v(\frac{a}{b}) = v(a) * v(b)$ for $a, b \in \mathbf{Z}$ and $b \neq 0$.

For the second example we take K to be the function field ofa complex algebraic curve C. Any smooth point z of C defines a Krull valuation $v = v_z$ of K as follows. Let $f \in K$ be given. If f has a zero of order $n > 0$ at z we set $v(f) = n$; if$f$ has a pole of order $n > 0$ at z we set $v(f) = -n$; and otherwise we set $v(f) = 0$.

If v is a Krull valuation of K then

$$\mathcal{O}_v = \{0\} \cup \{x \in K | v(x) > 0\}$$

is a sub-ring of K, the *valuation ring* defined by v. This ring has a unique maximal ideal, namely

$$\mathcal{M}_v = \{0\} \cup \{x \in K | v(x) > 0\}.$$

3.1.2. A Krull valuation v of K with value group Λ determines an action of the group $\mathrm{GL}_2(K)$ on a Λ-tree, the *Bruhat-Tits building* for $\mathrm{GL}_2(K)$. To define this tree we consider the set of all $(\mathcal{O}_v)$-*lattices* in the 2-dimensional vector space $V = K^2$. A lattice is by definition a finitely generated $\mathcal{O}_v$-submodule of V which spans V as a K-vector space. Any lattice is isomorphic to $\mathcal{O}_v$. We shall call two lattices L and L' *equivalent* if $L' = \theta\lambda L$ for some $\theta \in K^*$. We let T denote the set of all equivalence classes oflattices, and write $[L]$ for the equivalence class of a lattice L.

Given any two lattices L_1 and L_2 one can show that L_2 is equivalent to a lattice L_2' such that $L_2' \subset L_1$, and such that L_1/L_2' is isomorphic (as an $\mathcal{O}_v$-module) to $\mathcal{O}_v/\alpha\mathcal{O}_v$ for some non-zero element α of $\mathcal{O}_v$. One can also show that $v(\alpha)$ is uniquely determined by the equivalence classes of L_1 and L_2; we write$v(\alpha) = \mathrm{dist}([L_1], [L_2])$. Finally, one can show that with this definition of dist: $T \times T \to \Lambda$, the set T becomes a Λ-tree. (You will find proofs of these facts in [Se] for the case $\Lambda = \mathbf{Z}$ and in [MSh1] for the general case.)

3.1.3.

In the case $\Lambda = \mathbf{Z}$, there is a simple description of the link of a vertex s of T, i.e. the metric sphere of radius 1 about s. Let us write $s = [L_0]$, where L_0 is a lattice. Let k denote the *residue field* $\mathcal{O}/\mathcal{M} = \mathcal{O}_v/\mathcal{M}_v$. Since $\Lambda = \mathbf{Z}$, any element π of $\mathcal{O} = \mathcal{O}_v$ with $v(\pi) = 1$ generates $\mathcal{M}_v$. Hence the link of v consists of all classes represented by lattices $L \subset L_0$ such that L_O/L is isomorphic as an $\mathcal{O}$-module to $\mathcal{O}/\Updownarrow$. Such lattices L are in bijective correspondence, via the natural map $\mathcal{O}^2 \to k^2$, with 1-dimensional subspaces of the k-vector space k^2. Thus the link of any vertex has the structure of a projective line over the field k. In particular, if k is a finite field of order q then T is a $(q + 1)$-regular tree. (A graph is *k-regular*, where k is a positive integer, if every vertex has valence k.)

3.1.4. The natural (linear) action of $\mathrm{GL}_2(K)$ on V induces an action (by isometries) of $\mathrm{PGL}_2(K)$ on T. The restriction of this action to $PSL_2(K)$ an action without inversions.

It is easy to describe the stabilizers of points of T under the action of $\mathrm{PSL}_2(K)$. First one observes that the action of $\mathrm{PGL}_2(K)$ on lattices is transitive; hence the stabilizer in $\mathrm{PSL}_2(K)$ of any point of T is conjugate in $\mathrm{PGL}_2(K)$ to the stabilizer of the equivalence class of the standard lattice $\L_0 = \mathcal{O}_v^2 \subset K^2 = V$. Next one checks that a unimodular matrix can fix $[L_0]$ only if it fixes L_0. But the stabilizer of L_0 is by definition the subgroup $\mathrm{PSL}_2(\mathcal{O}_v)$ of $\mathrm{PSL}_2(K)$. So the point-stabilizers in $\mathrm{PSL}_2(K)$ are just the conjugates of $\mathrm{PSL}_2(\mathcal{O}_v)$ by elements of $\mathrm{PGL}_2(K)$.

3.2. Ihara's theorem. In [Se], Serre gives a number of applications of this construction. One celebrated application is a proof of a theorem due to Ihara about groups of matrices over p-adic numbers.

One nice way to define p-adic numbers is to use the concept of a Krull valuation. If p is any prime one can make $\mathbf{Q}$ into a metric space by using the valuation $v = v_p$ of $\mathbf{Q}$ that I defined above: the standard way is to set $\mathrm{dist}(x, y) = p^{v(x-y)}$.

The set $\mathbf{Q}_p$ of p-adic numbers is by definition the completion of this metric space. The field operations of $\mathbf{Q}$ have unique continuous extensions to $\mathbf{Q}_p$, making $\mathbf{Q}_p$ a field; and $v = v_p$ extends uniquely to a valuation of $\mathbf{Q}_p$ with value group $\mathbf{Z}$. I shall denote this extension by v_p as well. It is easy to show that $\mathbf{Q}_p$ is locally compact in the topology defined by its complete metric. Furthermore, the valuation ring $\mathbf{Z}_p$ is compact.

It follows that $\mathrm{PSL}_2(\mathbf{Q}_p)$ has in a natural way the structure of a locally compact topological group, and that $\mathrm{PSL}_2(\mathbf{Z}_p)$ is a compact subgroup. Ihara's theorem asserts that *every torsion-free, discrete subgroup of $PSL_2(\mathbf{Q}_p)$ is free.*

Here is Serre's proof of Ihara's theorem. By the construction that I described in 3.1, the valuation v_p determines a natural action without inversions of $\mathrm{PSL}_2(\mathbf{Q}_p)$ on a $\mathbf{Z}$-tree $T_p = T_{v_p}$. The stabilizer in $\mathrm{PSL}_2(\mathbf{Q}_p)$ of any point (vertex) of T_p is a conjugate of $\mathrm{PSL}_2(\mathbf{Z}_p)$ by an element of $\mathrm{PGL}_2(\mathbf{Q}_p)$, and is therefore compact. Hence if $\Gamma \leq \mathrm{PSL}_2(\mathbf{Q}_p)$ is discrete then the stabilizer in Γ of any point of T_p is both compact and discrete, and is therefore finite. If Γ is torsion-free it follows that the point-stabilizers in Γ are trivial; that is, the action of Γ is free. But as we saw in 1.1.2 and 1.1.4, a group which acts freely and without inversions on a $\mathbf{Z}$-tree is free.

3.3. Ramanujan graphs. Lubotzky, Phillips and Sarnak [LPS] have discovered an amazing and beautiful application of the Bruhat-Tits tree to graph theory. We have seen that the group $\mathrm{PGL}_2(\mathbf{Q}_p)$ has a natural action on a $\mathbf{Z}$-tree T_p. For any subgroup Γ of $\mathrm{PSL}_2(\mathbf{Q}_p)$ we can restrict the action to Γ and form the quotient T_p/Γ, which is a graph. This construction is especially interesting when Γ is a *lattice*, i.e. when Γ is a discrete subgroup of the locally compact group $\mathrm{PGL}_2(\mathbf{Q}_p)$ and the volume of the quotient space $\mathrm{PGL}_2(\mathbf{Q}_p)1/\Gamma$–defined in terms of the Haar measure on $\mathrm{PGL}_2(\mathbf{Q}_p)$–is finite. In this case T_p/Γ is a finite graph. As in 3.2, if Γ happens to be torsion-free then it acts freely and without inversions on T_p. In this case the graph T_p/Γ is $(p+1)$-regular, since T_p is a $(p+1)$-regular tree by 3.1.3.

This construction turns out to give examples of finite graphs that are extremely interesting from the perspective of combinatorics and computer science.

3.3.1. The most natural lattices in $\mathrm{PSL}_2(\mathbf{Q}_p)$ are the arithmetic ones, which are constructed from quaternion algebras. For any two positive integers u and v, we may construct a 4-dimensional associative $\mathbf{Q}$-algebra $D = D_{uv}$ with a linear basis $\{1, i, j, k\}$ satisfying the relations

$$i^2 = -u, j^2 = -v, ij = -ji = k.$$

If u and v are fixed, then for all but finitely many primes p, the $\mathbf{Q}_p$-algebra $D \otimes \mathbf{Q}_p$ is isomorphic to the matrix algebra $\mathcal{M}_2(\mathbf{Q}_p)$. In this case $D = D_{uv}$ is said to be *unramified* at p. Note that when D is unramified, its group of units $D^* = D^*_{uv}$ is isomorphic to $\mathrm{GL}_2(\mathbf{Q}_p)$.

For any prime p there is a sub-algebra $D(\mathbf{Z}[\frac{1}{p}])$ of D consisting of all elements whose coefficients in the basis $\{1, i, j, k\}$ belong to the ring $\mathbf{Z}[\frac{1}{p}]$. The group of units $D^*(\mathbf{Z}[\frac{1}{p}])$ is a subgroup of $D^*(u, v)$. Hence when $D(u, v)$ is unramified at p we can identify $D^*(\mathbf{Z}[\frac{1}{p}])$ with a subgroup of $\mathrm{GL}_2(\mathbf{Q}_p)$. The image Γ of $D^*(\mathbf{Z}[\frac{1}{p}])$ in $\mathrm{PGL}_2(\mathbf{Q}_p)$ is a lattice.

3.3.2. This lattice Γ does not define an interesting graph, because it acts transitively on T_p. However, for every positive integer N prime to p there is a *congruence subgroup* $\Gamma(N)$ of Γ modulo N. To define $\Gamma(N)$ we note that the unique ring homomorphism $\mathbf{Z}[\frac{1}{p}] \to \mathbf{Z}/n\mathbf{Z}$ induces a homomorphism h from $D(\mathbf{Z}[\frac{1}{p}])$ to $D(\mathbf{Z}/n\mathbf{Z})$, where $D(\mathbf{Z}/n\mathbf{Z}) = D_{uv}(\mathbf{Z}/n\mathbf{Z})$ is the (finite) 4-dimensional algebra defined over $\mathbf{Z}/n\mathbf{Z})$ in the same way that D_{uv} is defined over Q. The kernel of $h|D^*(\mathbf{Z}[\frac{1}{p}])$ is a finite-index subgroup of $D^*(\mathbf{Z}[\frac{1}{p}])$. We define $\Gamma(N)$ to be the image of the latter subgroup in Γ. Since $\Gamma(N)$ has finite index in Γ, it is again a lattice. Furthermore, if p is not congruent to a square modulo p (as is the case for 50% of all integers N) then $\Gamma(N)$ is torsion-free, and so $T_p/\Gamma(N)$ is a finite $(p+1)$-regular graph.

3.3.3. These regular graphs have many remarkable properties. The most important can be expressed in terms of adjacency matrices. If $\mathcal{G}$ is a finite graph with vertex set V, its *adjacency matrix* is defined to be $A = (a_{vw})_{v,w\in V}$, where $a_{v,w}$ is equal to 1 if v and w are joined by an edge, and is equal to 0 otherwise. One can encode subtle information about a finite graph $\mathcal{G}$ in the eigenvalues of its adjacency matrix A, or equivalently in those of the matrix $\Delta = (\#V)I - A$, which can be interpreted as the matrix of a combinatorial Laplacian operator for $\mathcal{G}$. In particular, let λ_1 denote the smallest positive eigenvalue of Δ, and for any set $X \subset V$ define the *coboundary* δX of X to be the set of all edges having exactly one endpoint in X. It was proved by Tanner [T] and Alon and Milman [AloM] that for any subset X of V with $\#X \le \frac{1}{2}\#V$ we have

$$\#\delta X \ge \frac{\lambda_1}{2}\#X.$$

This is expressed by saying that the *Cheeger constant* of $\mathcal{G}$ is at least $\frac{\lambda_1}{2}$. Roughly speaking, if the Cheeger constant is large then there are a lot of edges joining any subgraph to its complement. As you may imagine, this makes graphs with large Cheeger numbers, known as *expanders*[6] ,important as components in communication networks. So it is of practical importance to have graphs for which λ_1 is large.

For simplicity, let's fix an integer $k > 1$ and look at k-regular graphs. It is a theorem due to Alon and Boppana [AloB] that for any $\epsilon > 0$ there are only finitely many k-regular graphs with $\lambda_1 > k - \sqrt{2k-1} + \epsilon$. So in an asymptotic sense, the best estimate one can hope for is $\lambda_1 \ge k - sqrt 2k - 1$.

3.3.4. A k-regular graph $\mathcal{G}$ is called a *Ramanujan graph* if every eigenvalue of its adjacency matrix is either $= \pm k$ or $\le \sqrt{2k-1}$. For a Ramanujan graph we automatically have $\lambda_1 \ge k - sqrt 2k - 1$. Lubotzky, Phillips and Sarnak proved that the $(p+1)$-regular graphs $T_p/\Gamma(N)$ described in 3.3.2 are all Ramanujan graphs.

In particular this implies that for $k = p+1$, where p is a prime, there are infinitely many k-regular Ramanujan graphs. Morgenstern [Mst] has generalized this to the case $k = q+1$ where q is a prime power. Apparently nothing is known for other values of k.

For q a prime power, these methods allow one to describe infinite families of $(q+1)$-regular Ramanujan graphs quite explicitly, so that in principle they can be used for constructing networks.

3.3.5. The proof that the graphs $T_p/\Gamma(N)$ are Ramanujan graphs uses deep results in number theory and representation theory, including the Ramanujan Conjecture[7] (proved by Deligne using his proof of the Riemann-Weil conjecture) and the Jacquet-Langlands correspondence. Well, you get the idea. I had better refer you to Lubotzky's book [L] for details. The whole thing is an impressive demonstration of the power of pure mathematics.

3.4. Character varieties, ideal points and trees. I want to describe another way of applying the Bruhat-Tits construction. It involves some additional machinery that Marc Culler and I developed in our paper [CuS1]. Let me begin with some elementary background.

3.4.1. The group $\mathrm{SL}_2(\mathbf{C})$ has in a natural way the structure of a complex affine algebraic variety: it is the solution set of the equation $xw - yz = 0$ in the 4-dimensional affine space $\mathcal{M}$ of complex 2×2- matrices $\begin{pmatrix} x & y \\ z & w \end{pmatrix}$.

The group $\mathrm{PSL}_2(\mathbf{C})$ has likewise the structure of a complex affine variety. This can be seen, for example, by mapping the variety $\mathrm{SL}_2(\mathbf{C})$ into the affine space of linear transformations of the vector space $\mathcal{M}$ via

[6]For a more precise definition of an expander, see [L].

[7]Yes, that's where they got the term.

the map $A \mapsto C_A$, where $C_A\colon \mathcal{M} \to \mathcal{M}$ is defined by $C_A(M) = AMA^{-1}$. For any $A, B \in \mathrm{SL}_2(\mathbf{C})$ we have $C_A = C_B$ if and only if $A = \pm B$. Thus the image of the map $A \mapsto C_A$ is identified with $\mathrm{PSL}_2(\mathbf{C})$. But the map $A \mapsto C_A$ is proper and is defined by polynomials in the affine coordinates; this implies that its image is again an affine variety.

3.4.2. For any finitely generated group Γ , we may consider the set $\mathrm{Hom}(\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ of all representations of Γ in $\mathrm{PSL}_2(\mathbf{C})$. If we fix a set S of generators of Γ then a representation $\rho \in \mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ is determined by the images of the elements of S under ρ. So we can identify $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ with a subset of the algebraic variety $\mathrm{PSL}_2(\mathbf{C})^S$. It is easy to see that $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ is in fact an algebraic subset of $\mathrm{PSL}_2(\mathbf{C})^S$: to each defining relation in Γ there corresponds a matrix equation, which translates into a system of algebraic equations in the affine coordinates. Up to isomorphism the structure of an algebraic set on $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ is independent of the choice of generating set S.

3.4.3. Now let C be a curve in $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$, i.e. an irreducible affine algebraic subset of complex dimension 1. Let K be the function field of C: concretely we can think of K as consisting of rational maps from C to $\mathbf{C}$, i.e. functions $C \to \mathbf{C}$ of the form $\frac{f}{g}$, where f and g are restrictions of polynomial functions on the ambient affine space, and g does not vanish identically on C. Next consider the group $\mathrm{PSL}_2(K)$. We may interpret an element of this group as a rational mapping from C to $\mathrm{PSL}_2(\mathbf{C})$. With each element $\gamma \in \Gamma$ we can associate the rational mapping $P(\gamma)\colon C \to \mathrm{PSL}_2(\mathbf{C})$ defined as follows: an arbitrary point of C is a representation $\rho\colon \Gamma \to \mathrm{PSL}_2(\mathbf{C})$. We define $P(\gamma)$ to map each $\rho \in C$ to $\rho(\gamma)$. If for each $\gamma \in \Gamma$ we now identify $P(\gamma)$ with an element of $\mathrm{PSL}_2(K)$, we obtain the *tautological representation* $P\colon \Gamma \to \mathrm{PSL}_2(K)$.

Any affine curve C has (up to isomorphism) a unique projective completion $\hat{C}$ such that all the (finitely many) points of $\hat{C} - C$ are smooth. I'll call these points *ideal points* of C. One of the basic ideas of [CuS1] is that if Γ is a finitely generated group and $C$is a curve in $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$, ideal points of C tend to determine interesting actions of Γ on $\mathbf{Z}$-trees. It is not hard to see why this should be so. The inclusion map from C to $\hat{C}$ induces an isomorphism of function fields; that is, any rational function on C has a unique extension to $\hat{C}$ which is meromorphic at z. So we can identify the function field of $\hat{C}$ with K. Any point z of $\hat{C} - C$ then defines a valuation $v = v_z$ of K. This valuation in turn defines a $\mathbf{Z}$-tree $T = T_z$, and by pulling back the natural action of $\mathrm{PSL}_2(K)$ by the homomorphism P we obtain an action of Γ on T_z.

3.4.4. This action contains important information. For example, for any element γ of Γ we have a function $I_\gamma \in K$ defined by $I_\gamma = (\text{trace } \rho(\gamma))^2$ for any $\rho \in \mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$. (For an element of $\mathrm{PSL}_2(\mathbf{C})$ the square of the trace is well-defined although the trace itself is not.) The restriction of I_γ to C extends uniquely to a function $\hat{I}_\gamma\colon \hat{C} \to \mathbf{C}$ This function $\hat{I}_\gamma$ is finite-valued at a given ideal point z if and only if γ fixes some point of T_z. This is not hard to prove: if you chase through the definitions you will find that $\hat{I}_\gamma$ is finite-valued at z if and only if $(\text{trace} P(\gamma))^2 \in \mathcal{O}_v$, which by an elementary argument is equivalent to saying that $P(\gamma)$ is in a conjugate (within $\mathrm{GL}_2(K)$) of $\mathrm{PSL}_2(\mathcal{O}_v)$. But the conjugates of $\mathrm{PSL}_2(\mathcal{O}_v)$ are just the stabilizers of points of T.

A similar analysis shows that if $\hat{I}_\gamma$ has a pole at z then the order of the pole is $-2l(\gamma)$, where $l\colon \Gamma \to \mathbf{Z}$ is the length function defined by the action of Γ on T_z.

3.4.5. What this makes clear is that the action of Γ on the tree T_z reflects the behavior of the functions I_γ near the ideal point z. Now the important thing about the functions I_γ is that they are invariant under the natural action by conjugation of $\mathrm{GL}_2(\mathbf{C})$ on $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$. So it is natural to expect that the I_γ can be defined on some sort of "quotient" of $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ under the action of $\mathrm{PGL}_2(\mathbf{C})$, and that the trees T_z are best understood in terms of this quotient. This is quite true, as I shall explain.

3.4.6. Some care is required in order to define the quotient of $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ under the action of $\mathrm{PGL}_2(\mathbf{C})$. This is already apparent in the case of an infinite cyclic group $\Gamma = \langle t \rangle$. For every $z \in \mathbf{C}$ we have a representation ρ_z which maps t to $\begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix}$. The points $\rho_z \in \mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C})), z \in \mathbf{C}^*$ constitute a single orbit under the action of $\mathrm{PGL}_2(\mathbf{C})$. But this orbit is not closed in $\mathrm{Hom}\ (\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ since its closure contains the trivial representation ρ_0. So the quotient in the category of topological spaces is not even a Hausdorff space.

The correct approach is to form the quotient in the category of algebraic varieties. This can be done using the machinery of geometric invariant theory, or it can be done from the very elementary point of view described in [CuS1]. In any event, the upshot is that to every finitely generated group Γ we can canonically associate an affine algebraic set $X(\Gamma)$ and a morphism[8] $\tau\colon \mathrm{Hom}\,(\Gamma, \mathrm{PSL}_2(\mathbf{C})) \to X(\Gamma)$ such that two points ρ_1 and ρ_2 have the same image under τ if and only if $I_\gamma(\rho_1) = I_\gamma(\rho_2)$ for every $\gamma \in \Gamma$. In particular, τ maps every $\mathrm{PGL}_2(\mathbf{C})$-orbit to a point. What is more, τ is surjective, and every fiber of τ either is exactly a $\mathrm{PGL}_2(\mathbf{C})$-orbit or consists entirely of representations which are reducible, i.e. are conjugate to representations of Γ by upper triangular matrices. (Thus if we stay away from the reducible representations, which are relatively degenerate examples, $X(\Gamma)$ behaves like a quotient in the naïve sense.) For any $\gamma \in \Gamma$ the function I_γ induces a function on $X(\Gamma)$ which I will also denote I_γ. The coordinate ring of $X(\Gamma)$ is generated by functions of the form I_γ: what this means in concrete terms is that we can take $X(\Gamma)$ to live in an affine space in such a way that the coordinate functions are of the form I_γ.

I like to think of $X(\Gamma)$ in terms of group characters. A representation $\rho\colon \Gamma \to \mathrm{PSL}_2(\mathbf{C})$ determines a function $\chi\colon \Gamma \to \mathrm{PSL}_2(\mathbf{C})$, its *character*, given by $\chi(\gamma) = (\text{trace } \rho(\gamma))^2$. The condition for two points ρ_1 and ρ_2 to have the same image under τ , that $I_\gamma(\rho_1) = I_\gamma(\rho_2)$ for every $\gamma \in \Gamma$, can be paraphrased by saying that ρ_1 and ρ_2 have the same character. So we can identify the points of $X(\Gamma)$ with characters of representations in $\mathrm{PSL}_2(\mathbf{C})$, and τ is then the map that takes every representation to its character. Furthermore, for any $\gamma \in \Gamma$, the function $I_\gamma\colon X(\Gamma) \to \mathbf{C}$ is then simply the evaluation map $\chi \mapsto \chi(\gamma)$.

3.4.7. Once one has defined the space of characters $X(\Gamma)$ it is not hard to adapt the construction of 3.4.3 to curves in $X(\Gamma)$. To each ideal point x of an arbitrary curve $C \subset X(\Gamma)$, one we can associate an action of Γ on a $\mathbf{Z}$-tree T_x. Again, for any $\gamma \in \Gamma$, the extension $\hat{I}_\gamma$ of $I_\gamma|C$ to $\hat{C}$ is finite-valued at x if and only if γ fixes some point of T_x; and if $\hat{I}_\gamma$ has a pole at x then the order of the pole is $-2l(\gamma)$, where $l\colon \Gamma \to \mathbf{Z}$ is the length function defined by the action of Γ on T_x.

The construction of the action is a slight variant of the one described in 3.4.3. Since the natural map $\mathrm{Hom}\,(\Gamma, \mathrm{PSL}_2(\mathbf{C})) \to X(\Gamma)$ is surjective, it maps some irreducible subvariety W of $\mathrm{Hom}\,(\Gamma, \mathrm{PSL}_2(\mathbf{C}))$ onto a dense subset of C, and there is an induced monomorphism from the function field F of v to the function field K of W. Let us identify F with a subfield of K. The valuation v of F defined by the ideal point z can be extended–after possibly enlarging the value group by finite index, so that it is still isomorphic to $\mathbf{Z}$–to a valuation w of K. As in 3.4.3 we have a tautological representation $P\colon \Gamma \to \mathrm{PSL}_2(K)$. The valuation w defines a $\mathbf{Z}$-tree$T$, and $\mathrm{PSL}_2(K)$ has a natural action on T. Pulling back this action via P we get an action of Γ on T.

One clear advantage of working with a curve $C \subset X(\Gamma)$, as opposed to a curve in $\mathrm{Hom}\,(\Gamma, \mathrm{PSL}_2(\mathbf{C}))$, is that the action of Γ on T_z is non-trivial for every ideal point z of C . Indeed, if the action were trivial, the function I_γ would be finite-valued at z for every $\gamma \in \Gamma$. In particular the coordinate functions would be finite-valued at z; this is impossible since $z \notin C$.

3.5. A finiteness theorem. As a first application of the theory that I have been discussing in 3.4, let me give a proof of a result that was first proved from a somewhat different point of view by Hyman Bass in [Bas1,2]. *Let Γ be a finitely generated group which admits no splitting in the sense of 1.3.3: that is, Γ is not an amalgamated free product of two proper subgroups, and admits no homomorphism onto $\mathbf{Z}$. Then up to conjugacy there are only finitely many irreducible representations of Γ in $PSL_2(\mathbf{C})$.*

To prove this we observe that the character space $X(\Gamma)$ is 0-dimensional. Indeed, if the dimension of $X(\Gamma)$ were > 0 then $X(\Gamma)$ would contain a curve C. Any ideal point of C would define a non-trivial action of Γ on a $\mathbf{Z}$-tree, which by the Bass-Serre theory would give a non-trivial splitting of Γ and a contradiction to the hypothesis. Thus $X(\Gamma)$ is a finite set. Since an irreducible representation is determined up to conjugacy by its character, the assertion follows.

3.6. The 3-manifold connection. The theory that I described in 3.4 is particularly well-adapted to the study of a 3-manifold M through its fundamental group (and was introduced in [CuS1] for this purpose). This is because, on the one hand, representations of $\pi_1(M)$ in $\mathrm{PSL}_2(\mathbf{C})$ are related to geometric structures on M, while on the other hand, actions of $\pi_1(M)$ on trees are related to the topology of M.

3.6.1. The work of W. Thurston shows that hyperbolic geometry plays a central role in 3-manifold topology. A *hyperbolic metric* (or hyperbolic structure) on an n-manifold M is a complete Riemannian metric of constant curvature -1. A hyperbolic metric on M gives an identification of the universal cover of M with

[8]A map between affine algebraic sets is a *morphism* if it is defined by polynomials in the affine coordinates.

hyperbolic n-space H^n, which up to isometry is the unique simply connected hyperbolic n-manifold. In the following discussion I will be thinking of H^n concretely as the (open) upper half-space $\mathbf{R}^{n-1} \times (0, \infty) \subset \mathbf{R}^n$ with the metric $x_n^{-2} \sum_{i=1}^n dx_i^2$.

Hyperbolic n-space is highly symmetric. Its group of isometries $\mathrm{O}(n, 1)$ acts transitively on points, and the stabilizer of any point acts transitively on orthogonal frames in the tangent space. Any isometry $\gamma \in \mathrm{SO}(n, 1)$ extends to a homeomorphism $\bar{\gamma}$ of $\bar{H}^n$, the one-point compactification of the closed upper half-space $\mathbf{R}^{n-1} \times [0, \infty)$, which is topologically a closed ball. The boundary $S^(n-1)_\infty$ of $\bar{H}^n$ is the one-point compactification of $\mathbf{R}^{n-1}$.

We can identify S^1_∞ with the real projective line, and S^2_∞ with the complex projective line or Riemann sphere. If $n = 2$ or 3 and if γ belongs to the orientation-preserving subgroup $\mathrm{SO}(n, 1)$ of $\mathrm{O}(n, 1)$, then the restriction of $\bar{\gamma}$ to S^{n-1}_∞ is a homography (linear fractional transformation) $z \mapsto \frac{az+b}{cz+d}$, where a, b, c, d are real if $n = 2$ and complex if $n = 3$, and $ad - bc = 1$. Thus $\mathrm{SO}(2, 1)$ and $\mathrm{SO}(3, 1)$ are respectively isomorphic to the groups $\mathrm{PSL}_2(\mathbf{R})$ and $PSLC$ of real and complex homographies.

3.6.2. If M^n is hyperbolic and orientable then $\pi_1(M)$ acts on the universal cover H^n by deck transformations which belong to $\mathrm{SO}_(n, 1)$, and thus with the hyperbolic metric on M there is associated a representation pf $\pi_1(M)$ in $\mathrm{PSL}_2(\mathbf{C})$. This representation is discrete, in the sense that it is an isomorphism of $\pi_1(M)$ onto a discrete subgroup of $\mathrm{SO}_(n, 1)$.

These observations allow one to define a natural bijective correspondence between conjugacy classes of representations of a group Γ in $\mathrm{SO}_(n, 1)$ and n-dimensional *homotopy-hyperbolic structures* on an aspherical space $K = K(\Gamma, 1)$ with fundamental group Γ. A homotopy-hyperbolic structure is defined by a pair (M, ϕ) where M is a hyperbolic n-manifold and $\phi : K \to M$ is a homotopy equivalence. Two such pairs (M, ϕ) and (M', ϕ') define the same homotopy-hyperbolic structure if there is a homotopy equivalence $h\colon M \to M'$ such that $h \circ \phi$ is homotopic to ϕ'.

3.6.3. Let us consider an orientable 3-manifold M with a hyperbolic metric of finite volume. Then M contains a compact 3-manifold-with-boundary M_0, its *compact core*, such that every component of $M - \mathrm{int}\ M_0$ is isometric to the quotient of the *horoball* $\Omega = \mathbf{R}^2 \times [1, \infty) \subset H^3$ by a rank-2 free abelian discrete subgroup of $\mathrm{PSL}_2(\mathbf{C})$ consisting of elements of the form $\pm \begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix}$. (The group consisting of all such elements clearly leaves Ω invariant.) These components are called *cusps*, or maybe I should say neighborhoods of cusps. Each of them is homemorphic to $T^2 \times [0, \infty)$. It follows that M is homeomorphic to the interior of M_0; in particular, $\pi_1(M)$ and $\pi_1(M_0)$ are identified in a natural way.

3.6.4. Mostow's rigidity theorem asserts that if a manifold M of dimension > 2 admits a hyperbolic structure of finite volume, then the homotopy-hyperbolic structure on M is unique; that is, the discrete representation of $\pi_1(M)$ in $\mathrm{PSL}_2(\mathbf{C})$ is unique up to conjugation by an element of $\mathrm{PSL}_2(\mathbf{C})$. However, if $n = 3$ and M is non-compact, i.e. has at least one cusp, then $\pi_1(M)$ always has useful *non-discrete* representations in $\mathrm{PSL}_2(\mathbf{C})$. In fact, Thurston has shown that if M has m cusps then the character χ_0 of the discrete representation of $\pi_1(M)$ is a smooth point of $X(\Gamma)$ and that the irreducible component of $X(\Gamma)$ containing χ_0 has dimension m. In particular, if M has finite volume but is not compact then $X(\Gamma)$ always contains curves, and the construction that I described in 3.4 can be applied to any ideal point of any of these curves.

3.6.5. As I said above, the second reason why the theory described in 3.4 is especially well-adapted to the study of a 3-manifold M is that actions of $\pi_1(M)$ on trees are related to the topology of M. This connection involves the same ideas of Stallings's that I invoked in 2.4.1. (In this 3-dimensional context these ideas were also discussed in [Sh].)

For simplicity let us suppose that M is compact—but possibly with boundary—and orientable. For example, M may be the compact core of an orientable hyperbolic 3-manifold of finite volume. If $\Gamma = \pi_1(M)$ acts on a tree T, there is a Γ- equivariant map $\tilde{f}$ from the universal cover $\tilde{M}$ to T. In the case of a simplicial action on a simplicial tree T, we can choose $\tilde{f}$ so that the inverse image under $\tilde{f}$ of the set of midpoints of edges of T is a 2-manifold $\tilde{\Sigma}$. Again, $\tilde{\Sigma}$ may have a boundary, but it is properly embedded in M, in the sense that $\partial\Sigma = \Sigma \cap \partial M$. The equivariance of Σ implies that Σ is invariant under Γ. So we get a properly embedded compact 2-manifold $\Sigma = \tilde{\Sigma}/\Gamma \subset M$. If Γ acts without inversions we can take Σ to be orientable. Using fundamental results due to Papakyriakopoulos, one can also arrange that

(i) the fundamental group of each component of Σ maps injectively to $\pi_1(M)$.

Furthermore, we can choose $\tilde{f}$ so that Σ is non-degenerate in the sense that

(ii) no component of Σ is the boundary of a 3-ball, and

(iii) no component Σ_i of Σ is the frontier of a subset A of M such that $\pi_1(\Sigma_i)$ maps onto $\pi_1(A)$.

A properly embedded, compact, orientable 2-manifold $\Sigma \subset M$ that satisfies conditions (i)–(iii) is said to be *incompressible.* (At least that's my terminology. Some other authors use the term in a slightly weaker sense.)

If the given action of $\Gamma = \pi_1(M)$ on T is non-trivial, then any incompressible surface Σ obtained from the above construction is non-empty. This is because if Σ and hence $\tilde{\Sigma}$ were empty, $\tilde{f}$ would map $\tilde{M}$ to the star of some vertex s in the first barycentric subdivision of T; the vertex s would then be fixed by the entire group Γ.

3.7. Separating suraces in knot manifolds. The interaction among hyperbolic structures on int M_0, representations of $\pi_1(M)$ in $\mathrm{PSL}_2(\mathbf{C})$, actions of $\pi_1(M)$ on trees, and incompressible surfaces in M has had a number of applications in 3-manifold theory. An amusing early application was made in [CuS2] to the problem of finding interesting incompressible surfaces in knot complements.

3.7.1. If k is a non-trivial knot in the 3-sphere then there is always a connected orientable surface $F \subset S^3$, called a *Seifert surface*, whose boundary is k. We can take F to meet a tubular neighborhood N of k in an annulus. Now $M = S^3 - \mathrm{int}N$ is a compact 3-manifold bounded by a torus. The surface $F \cap M$ is properly embedded in M. Its boundary is a *longitude*, i.e. a simple closed curve in ∂M whose homology class represents a generator of $\ker H_1(\partial M, \mathbf{Z}) \to H_1(M, \mathbf{Z})$. It follows from the work of Papakyriakopoulos that one can always choose F so that $F \cap M$ is incompressible. For example this is always the case if we take F to have minimal genus among all Seifert surfaces.

So if M is the complement of a non-trivial knot in S^3, the obvious connected incompressible surfaces in M are those whose boundary consists of a single longitude; these are essentially Seifert surfaces. The question arises whether M always contains incompressible surfaces other than the obvious ones.

3.7.2. If Σ is an *arbitrary* incompressible surface with non-empty boundary in M, then $\partial\Sigma$ consist of a certain number of homotopically non-trivial simple closed curves in ∂M. As ∂M is a torus, the components of $\partial\Sigma$ must all represent the same homology class $\delta \in H_1(\partial M)$. Since δ is represented by a simple closed curve, it is an indivisible (or unimodular) element of $\pi_1(\partial M)$, and it is defined only up to sign since we have not specified any orientations. An indivisible element of $H_1(\partial M)$, defined up to sign, is often called a *slope.* A slope δ determined by a bounded incompressible surface in the way that I have described is called a *boundary slope.*

In [CuS2] it was shown that for every knot in S^3 (or in any rational homology sphere) there exists a boundary slope which is distinct from the (slope of a) longitude. In other words, there is always a closed incompressible surface whose boundary components are not longitudes. It is easy to show that such a surface always separates M into two pieces, in contrast to the Seifert surfaces which are always non-separating.

3.7.3. One consequence is a result originally conjectured by L. Neuwirth:

Every knot group is a non-trivial free product with amalgamation in which the amalgamated subgroup is free.

3.7.4. To prove the main result of [CuS2] one uses Thurston's results to reduce to the case in which int M admits a hyperbolic structure. In this case, it follows from 3.6.4 that the component of $X(\pi_1(M))$ containing the character of the discrete faithful representation is a curve C.

Let us fix a base point in ∂M. Since $\pi_1(\partial M)$ is abelian, each element $\alpha \in H_1(\partial M)$ corresponds to a unique element in $\pi_1(\partial M)$, whose image in $\pi_1(M)$ I will denote by $e(\alpha)$.

Each ideal point x of C determines a non-trivial action of $\pi_1(M)$ on a tree, which can be used to construct a non-empty incompressible 2-manifold Σ_x. One wishes to show that for some ideal point x, the boundary of Σ_x is non-empty and the boundary slope that it defines is not a longitude. To show this one considers the function $I_{e(\lambda)}$ on C, where λ denotes the longitudinal slope. By chasing through the definitions one discovers that if $\partial\Sigma_x = \emptyset$, or if the boundary slope defined by $\partial\Sigma_x$ is a longitude, then $I_{e(\lambda)}$ is finite-valued–i.e. does not have a pole–at x. Now one of the basic properties of the curve C provided by 3.6.4 is that for any non-zero element α of $H_1(\partial M)$, the function $I_{e(\alpha)}$ is non-constant. In particular, $I_{e(\lambda)}$ is non-constant; since it is finite-valued on C it must have a pole at some ideal point of C. This proves the Neuwirth conjecture.

3.8. Surgery on knots. A more ambitious application of the techniques described in 3.4 and 3.6 occurs in the proof of the Cyclic Surgery Theorem, which was proved by Culler, Gordon Luecke and myself [CuGLS]. Before explaining the statement I must give a bit of background.

3.8.1. A *solid torus* is a 3-manifold N homeomorphic to $D^2 \times S^1$. The simple closed curves $\{0\} \times S^1 \subset \text{int } N$ and $(\partial D^2) \times S^1 \subset \partial N$ are well-defined up to isotopy and are called the *core* and *meridian* of N. If M is a 3-manifold whose boundary is a 2-torus, the operation known as *Dehn filling* consists of attaching a solid torus N to M by some homeomorphism between ∂M and ∂N. The topological type of the resulting closed manifold is determined if one specifies the slope α (see 3.7.2) of the simple closed curve in N which is attached to the meridian of N. I'll write M_α for the manifold obtained by the Dehn filling. The group $\pi_1(M_\alpha)$ is isomorphic to $\pi_1(M)/\langle e(\alpha)\rangle$, where $\langle\ \rangle$ denotes normal closure, and $e(\alpha)$ is defined as in 3.7.4.

Let M be a 3-manifold whose boundary is a torus, and α, β two indivisible elements of $H_1(\partial M)$. The 3-manifolds M_α and M_β are said to be related to each other by a *Dehn surgery.* Thus Dehn surgery is the operation of removing a solid torus from the interior of a manifold and sewing it back in a different way.

It is a classical result that any closed orientable 3-manifold can be obtained from the 3-sphere by a finite sequence of Dehn surgeries. A good deal of attention has been focused on the manifolds obtained from S^3 by a single Dehn surgery. Thus one considers Dehn fillings of manifolds $M = S^3 - N$, where $N \subset S^3$ is a solid torus. We can think of N as a tubular neighborhood of its core, which is a knot K. There is a natural basis of $H_1(\partial M)$ consisting of the meridian λ of N and the longitude μ of K (see 3.7.1). (One must specify some orientations for the signs of λ and mu to be well-defined.)

It is convenient to parametrize slopes in ∂M by elements of $\mathbf{Q} \cup \{\infty\}$, by letting $\frac{a}{b}$ correspond to the slope $a\mu + b\lambda$.

One often writes $K(\frac{a}{b}) = M_{a\mu+b\lambda}$. The number $\frac{a}{b}$ is called a surgery coefficient. Note that $K(\infty) = M(\mu) = S^3$: this is the *trivial* surgery.

3.8.2. Any surgery on the trivial knot gives a lens space, i.e the quotient of S^3 by a cyclic group acting freely by isometries (in the round metric). A lens space clearly has a cyclic fundamental group. The cyclic surgery theorem deals with the question of which non-trivial surgeries on non-trivial knots can give manifolds with cyclic fundamental group.

3.8.3. Elementary examples are provided by torus knots. A *torus knot* is a knot in S^3 which can be isotoped into a standard torus. If we think of S^3 as the unit sphere $|z|^2 + |w|^2 = 2$ in $\mathbf{C}^2$, the standard torus Q is defined by $|z| = |w| = 1$. Torus knots arise in algebraic geometry in connection with singularities of algebraic curves: a plane curve of the form $z^p w^q = 1$, where p and q are relatively prime integers, has a singularity at the origin, and its intersection with a small Euclidean sphere $S \subset \mathbf{C}^2$ is a torus knot in S. For any torus knot K, there are infinitely many Dehn surgeries on K that give lens spaces.

Torus knots often play an exceptional role in knot theory because their complements are *Seifert fibered spaces*, i.e. they admit C^∞ foliations by 1-spheres. Seifert fibered spaces form a manageable and pleasant, but rather degenerate, class of 3-manifolds.

Rolfson [R] showed that certain *iterated torus knots*— which, like torus knots, arise as links of singularities of plane algebraic curves—also admit non-trivial Dehn surgeries that give lens spaces. This may regarded as a partial generalization of what happens for torus knots; however, for an iterated torus knot which is not a torus knot there is at most one surgery (as opposed to infinitely many) that can give a lens space.

3.8.4. A remarkable example was given by Fintushel and Stern [FiS]. They showed that the surgeries on the so-called (-2,3,7)-pretzel knot with surgery coefficients 18 and 19 both yield lens spaces. As we shall see, the Cyclic Surgery Theorem sheds light on this example.

3.8.5. The general version of the theorem is best stated in terms of Dehn filling. A 3-manifold M is said to be *irreducible* if every smooth 2- sphere in M bounds a ball. A classical theorem due to Alexander implies that every knot complement in S^3 is irreducible.

> *Let M be an irreducible, compact, orientable 3-manifold whose boundary is a torus. Suppose that M is not a Seifert fibered space. Let α and α' be two slopes in $H_1(\partial M)$. Suppose that $\pi_1(M_\alpha)$ and $\pi_1(M_{\alpha'})$ are both cyclic. Then the homological intersection number of α and α' (which is defined up to sign) has absolute value at most 1.*

3.8.6. This has the following formal consequences regarding non-trivial surgery on a knot K in S^3.

(3.8.6.1) *Suppose that K is not a torus knot. For any $r \in \mathbf{Q}$, if $\pi_1(K(r))$ is cyclic then r is an integer. There are at most two integers r for which $\pi_1(K(r))$ is cyclic, and if there are two they must be consecutive integers (as in the Fintushel-Stern example). Only for $r = 1$ or $r = -1$–and not both–can $K(r)$ possibly be simply connected.*

Combining this last fact with a result due to Bleiler and Scharlemann, one can show:

(3.8.6.2) *If K is a non-trivial knot which is invariant under a non-trivial periodic homeomorphism of S^3, then there is no $r \in \mathbf{Q}$ for which $K(r)$ is simply connected.*

3.8.7. I will briefly sketch the proof of the Cyclic Surgery Theorem in the case where the manifold M contains no closed incompressible surfaces. In this case Thurston's results imply that int M has a hyperbolic metric, and as in 3.6.4 and 3.7.4 we have a curve $C \subset X(\pi_1(M))$.

As I mentioned in 3.7.4, $I_{e(\alpha)}$ is non-constant whenever $1 \neq \alpha \in \pi_1(M)$; thus the degree of $I_{e(\alpha)}$ is a positive[9] integer in this case. We can interpret deg $I_{e(\alpha)}$ as the number of poles of $I_{e(\alpha)}$. Since $I_{e(\alpha)}$ has poles only at ideal points, we have deg $I_{e(\alpha)} = \sum_z P_{z,\alpha}$, where z ranges over the ideal points of C and and $P_{z,\alpha}$ denotes the order of the pole of $I_{e(\alpha)}$ at z, or 0 if $I_{e(\alpha)}$ does not have a pole at z.

By 3.4 and 3.6, an ideal point z determines an action of $\Gamma = pi_1(M)$ on a tree, with which one can associate an incompressible surface $\Sigma_z \subset M$. Since we are assuming that M contains no closed incompressible surfaces, Σ_z must have non-empty boundary, and it therefore determines a boundary slope δ_z. One can show that there is a homomorphism $l_z \colon H_1(\partial M) \to \mathbf{Z}$, whose kernel is generated by β_z, such that $P_{z,\alpha} = |l_z(\alpha)|$ for every $\alpha \in H_1(\partial M)$. We can define a norm on the 2-dimensional vector space $V = H_1(\partial M; \mathbf{R})$ by $\|a\| = \sum_z l_z(a)$, where z ranges over the ideal points of C. We have deg $I_{e(\alpha)} = \|\alpha\|$ for every $\alpha \in H_1(\partial M)$. The unit ball of this norm is a compact convex polygon in V which is balanced (i.e. symmetric about 0) and whose vertices lie on lines spanned by boundary slopes.

Let's set $m = \min_{0 \neq \alpha \in L} \|\alpha\|$. Then the ball B of radius m with respect to our norm is again a convex balanced polygon whose vertices lies on lines spanned by boundary slopes. By definition, int B contains no non-zero points of the lattice $L = H_1(\partial M)$. This implies, by a well-known elementary argument due to Minkowski, that the area of B is at most 4. (Here I am measuring area in V in such a way that V/L has area 1. If we identify V with $\mathbf{R}^2$ in such a way that $L = \mathbf{Z}^2$, we are looking at ordinary area on R^2.)

The key step in the proof of the Cyclic Surgery Theorem, in the case we are considering, is to show that

(3.8.7.1) for any slope α such that $\pi_1(M_\alpha)$ is cyclic, we have $\|\alpha\| = m$, so that $\alpha \in \partial B$; and furthermore α is not a vertex of B.

Once (3.8.7.1) has been established the theorem follows easily. For if α and α' are two distinct slopes such that $\pi_1(M_\alpha)$ and $\pi_1(M_\alpha)$ are cyclic, then the four points $\pm\alpha, \pm\alpha'$ are the vertices of a parallelogram Π whose area is $2I$, where I denotes the absolute value of the homological intersection number of α and α'. It follows from (3.8.7.1) that $\Pi \subset B$ and hence that $I \leq \frac{1}{2}$Area $B \leq 2$, and that equality holds only if $B = \Pi$. But in the latter case α and α' would be vertices of B, contradicting the second assertion of (3.8.7.1). So we must have $I \leq 1$.

The first assertion of (3.8.7.1) is equivalent to saying that if α is a slope such that $\pi_1(M_\alpha)$ is cyclic, then for any non-zero element β of $H_1(\partial M; \mathbf{Z})$ we have deg $I_{c(\alpha)} \leq$ deg $I_{c(\beta)}$. This is proved by showing that at every point of $\tilde{C}$ (or more accurately of its de-singularization) where I_α takes the value 4, the function I_β also takes the value 4, and with at least the same multiplicity. This is in turn proved in two cases, depending on whether the given point lies in C or an ideal point. First let's consider the case of a point $x \in C$.

It may be shown that every point of C is the character of some representation $\rho \colon M \to \mathrm{PSL}_2(\mathbf{C})$ with a non-cyclic image. For such a ρ we must have $\rho(\alpha) \neq 1$; for otherwise ρ would factor through a representation of $\pi_1(\Sigma_\alpha) = \pi_1(M)/\langle e(\alpha)\rangle$ with a non-cyclic image, and this is impossible since $\pi_1(\Sigma_\alpha)$ is cyclic. Hence for every point $x \in C$ for which $I_{e(\alpha)} = 4$, there is a representation ρ with character x such that $\rho(e(\alpha)) = \pm\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$. Now since $\pi_1(\partial M)$ is abelian, any element $e(\beta)$ of $\pi_1(\partial M)$ is represented by an element that commutes with $\pm\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and hence has the form $\pm\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}$. In particular $I_{c(\beta)}(x) = 4$. By refining this argument—and doing a fair amount of hard technical work—one can show that $I_{c(\beta)}$ takes the

[9] in ze English sense

value 4 at x with at least the same multiplicity as $I_{c(\beta)}$; and one can make the argument work on points of the de- singularization of $\tilde{C}$ that correspond to points of C.

The case of an ideal point is quite different. In this case one actually shows that the assertion is vacuously true; that is, one shows that if $\pi_1(M_\alpha)$ is cyclic, then $I_{c(\alpha)}$ cannot take the value 4, or any finite value for that matter, at an ideal point. This is one point at which the theory of 3.4 and 3.6 is crucial. If $I_{c(\alpha)}$ is finite-valued at an ideal point z then we have $\alpha = \delta_z$, so that α is a boundary slope. But one can show that $\pi_1(M_\alpha)$ cannot be cyclic when α is a boundary slope.

This is done by considering an incompressible surface Σ_0 which has boundary slope α and which has the smallest possible number of boundary components among all such surfaces. In particular Σ_0 is connected. Let's write $M_\alpha = M \cup N$, where N is a solid torus. Since Σ_0 has boundary slope α, its boundary components bound disjoint disks in N. The union of Σ_0 with these disks is a closed surface $\hat{\Sigma}_0 \subset M_\alpha$. If Σ_0 has positive genus, one proves that $\hat{\Sigma}_0$ is incompressible in M_α. In this case, $\pi_1(M_\alpha)$ contains the fundamental group of a closed orientable surface of positive genus, and is therefore non-cyclic. In the case where $\hat{\Sigma}_0$ is a sphere, one shows that this sphere decomposes M_α as a connected sum of two non- trivial lens spaces. So in this case $\pi_1(M_\alpha)$ is a free product of two non-trivial cyclic groups, and is therefore not cyclic.

The second assertion of (3.8.7.1) is now easy. I pointed out in defining B that the vertices of B lie on lines spanned by boundary slopes. So if $\alpha \in L \cap \partial B$ were a vertex of B then α would itself be a boundary slope. But we just saw that this implies that $\pi_1(M_\alpha)$ is non-cyclic.

Section 4. Higher-dimensional varieties of characters and degenerations of hyperbolic structures

4.1. Compactifying character varieties. In 3.4 I explained how an ideal point z of a curve in the space $X(\Gamma)$ of $\mathrm{PSL}_2(\mathbf{C})$-characters of a finitely generated group Γ defines a non-trivial action of Γ on a $\mathbf{Z}$-tree. I pointed out that for any $\gamma \in \Gamma$ the function $\hat{I}_\gamma$ is finite-valued at z if and only if γ fixes some point of T_z, and that if $\hat{I}_\gamma$ has a pole at z then the order of the pole is $-l(\gamma)$, where $l \colon \Gamma \to \mathbf{Z}$ is the length function defined by the action of Γ on T_z.

These properties have a convenient restatement in terms of the formalism of projectivized length functions introduced in 1.7.1. Let $[l_z] \in \mathcal{PL}(\Gamma) \subset \mathcal{P}^\Gamma$ be the projectivized length function defined by the action associated to z. (By 1.7.2, the integer-valued length function defined by the action coincides with the real-valued length function defined by its real completion.) On the other hand, let us define a continuous map $\Theta : X(\Gamma) \to \mathcal{P}^\Gamma$ by defining $\Theta(x)$ to be the image in $\mathcal{P}^\Gamma$ of $(|\mathrm{Re}\ \mathrm{arccosh}\ (\frac{1}{2}(I_c(x))^2 - 1)|)_{c \in \mathcal{C}(\Gamma)} \in [0,\infty)^{\mathcal{C}(\Gamma)}$. Here the curve C is understood to have the topology induced by the usual topology of $\mathbf{C}$. The multi-valued function arccosh has a well-defined real part up to sign. Geometrically[10] , if x is the character of a representation in $\mathrm{PSL}_2(\mathbf{C}) = \mathrm{SO}(3,1)$, and $\gamma \in \Gamma$ is an element representing the conjugacy class c, the expression $|\mathrm{Re}\ \mathrm{arccosh}\ (\frac{1}{2}(I_c(x))^2 - 1)|$ gives the hyperbolic translation length of $\rho(\gamma)$ in H^3. In terms of the map θ, the properties of the action recalled above translate into the following fact: for any sequence (x_i) of points in C which converges to z in $\hat{C}$, the sequence $(\Theta(x_i))$ converges to $[l_z]$ in $\mathcal{P}^\Gamma$. So we can extend Θ to a continuous map from $\hat{C}$ to $\mathcal{P}^\Gamma$ by mapping each ideal point z to $[l_z]$.

4.1.1. This picture was generalized by Morgan and myself in [MSh1]. Instead of considering sequences of points on a fixed curve, we considered arbitrary sequences (x_i) tending to infinity in the locally compact space $X(\Gamma)$. It is not hard to show that any such sequence has a subsequence whose image under Θ converges to some point $[l] \in \mathcal{P}^\Gamma$. Morgan and I showed that the limit point $[l]$ always belongs to $\mathcal{PL}(\Gamma)$.

4.1.2. In contrast to the case where the sequence lies on a curve, the projectivized length function $[l]$ is not integer-valued in general. Instead, it takes values in a subgroup Λ of $\mathbf{R}$ whose $\mathbf{Q}$-rank $\dim_{\mathbf{Q}} \mathbf{Q}\Lambda$ is at most the dimension n of V. I will express this briefly by saying that the action has rank at most n.

An action on an $\mathbf{R}$-tree will be said to have rank $\leq n$ if its length function has rank $\leq n$. It follows from 1.7.2 that an action has rank $\leq n$ if and only if it is the real completion of a non-trivial action on a Λ-tree for some subgroup Λ of $\mathbf{R}$ whose $\mathbf{Q}$-rank is at most n.

---

[10]In place of $|\mathrm{Re}\ \mathrm{arccosh}\ (\frac{1}{2}(I_c(x))^2 - 1)|$ one could use any globally defined expression which goes to infinity like a constant multiple of $\log I_c(x)$ as $I_c \to \infty$. In [MSh1], Morgan and I used $\log(|I_c(x)| + 2$. I have used the more complicated expression above because it is more directly related to the geometric applications that I'll be talking about later.

4.1.3. Using 4.1.1 one can define a natural compactification $\hat{X} = \hat{X}(\Gamma)$ of $X = X(\Gamma)$, such that $\hat{X} - X$ is identified with a subset of $\mathcal{PL}(\Gamma)$. As a set, $\hat{X}$ is the disjoint union of X with a set $\mathcal{B} \subset \mathcal{PL}(\Gamma)$. The set $\mathcal{B}$ consists of all points which are limits in $\mathcal{P}^\Gamma$ of convergent sequences of the form $(\Theta(x_i))$ where $(x_i) \to \infty$ in X. The subsets X and $\mathcal{B}$ of $\hat{X}$ have respectively the complex topology and the topology inherited from $\mathcal{P}^\Gamma$. A sequence (x_i) in X converges in $\hat{X}$ to a point $[l] \in \mathcal{B}$ if and only if $\Theta(x_i) \to [l]$ in $\mathcal{P}^\Gamma$. These conditions characterize a compact topology on $\hat{X}$ in which X is an open dense subset.

4.1.4. Actually one can define a compactification of this type not only for the $\mathrm{PSL}_2(\mathbf{C})$ character space of a group, but for any complex affine algebraic set. (This is carried out in detail in [MSh1]. The construction depends on the choice of a countable set of generators for the coordinate ring as a $\mathbf{C}$-algebra; for the case of a character space $X(\Gamma)$ one takes the set to consist of all the functions $I_c, c \in \mathcal{C}(\Gamma)$.) In the case where the given algebraic set is a curve C, the compactification defined in this way is canonically isomorphic to the projective completion $\hat{C}$ described in 3.4.3.

However, the point to be emphasized is that when the given algebraic set is a character space of a group, the *ideal points*, i.e. the points of $\mathcal{B} = \hat{X} - X$, correspond to non- trivial actions of the group on $\mathbf{R}$-trees. This is what generalizes the theory of [CuS1], and–like the theory of [CuS1]–is useful for applications.

4.1.5. Before I talk about applications of this generalized theory, it will be best to say a few words about the proofs of the assertions of 4.1.1, and to discuss some further generalizations.

The proof of the main assertion of 4.1.1, while technically much more involved than the proof of the special case discussed in 3.4, is philosophically very similar. Suppose that (x_i) is a sequence tending to infinity in X, and that $\Theta(x_i)$ converges in $\mathcal{P}^\Gamma$. After passing to a subsequence we can assume that the x_i all lie in the same irreducible component V of X. The coordinate ring $\hat{\mathbf{Q}}[V]$, where $\hat{\mathbf{Q}} \subset \mathbf{C}$ denotes the algebraic closure of $\mathbf{Q}$, is an integral domain which contains all the functions $I_c, c \in \mathcal{C}(\Gamma)$. Its field of fractions is the function field $\hat{\mathbf{Q}}(V)$. After approximating the x_i by nearby generic points (without changing their limit) and again passing to a subsequence, we can achieve a nice situation in which for every function $f \in \hat{\mathbf{Q}}(V)$, the sequence $(f(x_i))$ has a limit in the extended complex line $\mathbf{C} \cup \{\infty\}$. In this situation there is a Krull valuation v of $\hat{\mathbf{Q}}(V)$ whose valuation ring $\mathcal{O}_v$ consists of all functions $f \in \hat{\mathbf{Q}}(V)$ such that $\lim_{i \to \infty}(f(x_i)) \neq \infty$. This valuation will of course play the role of the valuation associated to the ideal point in the case discussed in 3.4. The assumption that $x_i \to \infty$ in V implies that for some $c \in \mathcal{C}(\Gamma)$ we have $I_c \notin \mathcal{O}_v$. We have $v(\hat{\mathbf{Q}}^*) = 0$.

Since the natural map Hom $(\Gamma, \mathrm{PSL}_2(\mathbf{C})) \to X(\Gamma)$ is surjective, it maps some irreducible component W of $HGPSLC$ onto a dense subset of V, and there is an induced monomorphism from $\hat{\mathbf{Q}}(V)$ to the function field K of W over $\hat{\mathbf{Q}}$. Let us identify $\hat{\mathbf{Q}}(X)$ with a subfield of K. The valuation v can be extended–after possibly enlarging the value group by finite index–to a valuation w of K. As in 3.4.3 we have a tautological representation $P\colon \Gamma \to \mathrm{PSL}_2(K)$. The valuation w defines a Λ-tree T, where Λ is the value group of w, and $\mathrm{PSL}_2(K)$ has a natural action on T. Pulling back this action via P we get an action of Γ on T. One can check that this pulled-back action is non-trivial by checking that any element $\gamma \in \Gamma$ such that $I_\gamma \notin \mathcal{O}_v$ acts without a fixed point on T.

If we are lucky, Λ will be order-isomorphic to a subgroup of $\mathbf{R}$. In this case the action of Γ on T extends to an action on the real completion $\mathbf{R}T$. The identification of Λ with a subgroup of $\mathbf{R}$ is unique modulo a multiplicative constant, and hence the projectivized length function $l \in \mathcal{PL}(\Gamma)$ defined by the action of Γ on $\mathbf{R}T$is uniquely determined by the valuation w. In this case one can check that the given sequence (x_i) has limit $[l]$.

In general Λ need not be order-isomorphic to a subgroup of $\mathbf{R}$, so one has to work harder. What turns out to be true in general is that there are convex subgroups (see 2.6.1) $\Lambda_0 \subset \Lambda_1$ of Λ such that Λ_1/Λ_0 is order-isomorphic to a subgroup of $\mathbf{R}$, and such that we have $w_-(I_c) = -min(0, w_(I_c) \in \Lambda_1$ for every $c \in \mathcal{C}(\Gamma)$, but $w_-(I_c) \notin \Lambda_0$ for some $c \in \mathcal{C}(\Gamma)$. The abelian group Λ/Λ_0 inherits an order from Λ, and by composing w with the projection $\Lambda \to \Lambda/\Lambda_0$ we get a valuation $\bar{w}\colon F^* \to \Lambda/\Lambda_0$. This gives an action of Γ on a Λ/Λ_0-tree, which may be shown to contain a Γ-invariant Λ_1/Λ_0-tree T. Since Λ_1/Λ_0 is order-isomorphic to a subgroup of $\mathbf{R}$, we can complete T to an $\mathbf{R}$-tree and proceed as before.

4.1.6. It is not hard to see from this construction why the points of $\mathcal{B} = \hat{X} - X$ are defined by actions of rank at most $n = \dim X$, as asserted in 4.1.2. In fact, for any component V of X, the transcendence degree of $\hat{\mathbf{Q}}(V)$ over $\hat{\mathbf{Q}}$ is dim $V \leq n$, and from this it is a matter of elementary commutative algebra to deduce that any valuation of $\hat{\mathbf{Q}}(V)$ which is trivial on V has a value group Λ whose $\mathbf{Q}$-rank $\dim_{\mathbf{Q}}(\Lambda \otimes \mathbf{Q})$ is at most n. Enlarging a group by finite index and passing to a subgroup or quotient group do not increase the $\mathbf{Q}$-rank.

4.1.7. Notice that although the statement given in 4.1.1 involves only $\mathbf{R}$-trees, it is natural to prove it using Λ-trees for more general Λ. Indeed, the Λ-trees that arise in the proof contain important information. They describe the relative growth rates of the hyperbolic translation lengths of elements of Γ, in a sense that I will make precise in a moment.

If Λ is an ordered abelian group of $\mathbf{Q}$-rank $\leq n$, then the order-rank r of Λ (see 2.6.1) is also at most n. Let $\{0\} = \Lambda_0 \leq \ldots \Lambda_r = \Lambda$ be the convex subgroups of Λ. Each of the quotient groups Λ_k/Λ_{k-1} has order- rank 1 and hence admits an order-preserving embedding in $\mathbf{R}$. For any non-negative element λ of Λ, let us define the *height* of λ to be the least index k such that $\lambda \in \Lambda_k$. If λ and λ' are positive elements with the same height k, we define their *quotient* λ/λ' to be the real number $J(\lambda)/J(\lambda')$, where $J\colon \Lambda_k \to \mathbf{R}$ is the composition of the quotient map $\Lambda_k \to \Lambda_k/\Lambda_{k-1}$ with an embedding of Λ_k/Λ_{k-1} in $\mathbf{R}$. Since the embedding is unique up to a multiplicative constant, the quotient is well-defined. Let us set $\lambda/\lambda' = 0$ if height $\lambda <$ height λ', and $\lambda/\lambda' = \infty$ if height $\lambda >$ height λ'.

Now let (x_i) be a sequence tending to infinity in $X(\Gamma)$. After passing to a subsequence we can assume that for any two conjugacy classes c, c' in Γ, the sequence of quotients of hyperbolic translation lengths

$$\frac{|\mathrm{Re}\ \mathrm{arccosh}\ (\frac{1}{2}(I_c(x))^2 - 1)|}{|\mathrm{Re}\ \mathrm{arccosh}\ (\frac{1}{2}(I_{c'}(x))^2 - 1)|}$$

has a limit in $[0, \infty]$. The construction described in 4.1.5 gives a Λ-valued length function l defined by an action of Γ on a Λ-tree, where Λ is an ordered abelian group of finite $\mathbf{Q}$- rank, such that for any two elements $\gamma, \gamma' \in \Gamma$, the limit of the above sequence is equal to $l(c)/l(c')$. (in the quotient notation introduced above).

4.1.8. In [M1], Morgan generalized the theory described in 4.1.1–4.1.7. For any finitely generated group Γ and any $n \geq 2$ one can define a variety $X_n(\Gamma)$ of characters of representations of Γ in the isometry group $\mathrm{SO}(n, 1)$ of hyperbolic n-space. Since by 3.6.1 the groups $\mathrm{SO}(2, 1)$ and $\mathrm{SO}(3, 1)$ are respectively isomorphic to $\mathrm{PSL}_2(\mathbf{R})$ and $\mathrm{PSL}_2(\mathbf{C})$, we can identify $X_3(\Gamma)$ with $X(\Gamma)$, and $X_2(\Gamma)$ with the subset $X_{\mathbf{R}}(\Gamma)$ of $X(\Gamma)$ consisting of all characters of representations in $PSLR \subset \mathrm{PSL}_2(\mathbf{C})$. In [M1] it is shown that for any $n \geq 2$ the space has a natural compactification $\hat{X}_n(\Gamma)$ by projectivized length functions, which specializes for $n = 3$ to the compactification $\hat{X}(\Gamma)$ described in 4.1.3, and for $n = 2$ to the closure of $X_{\mathbf{R}}(\Gamma)$ in $\hat{X}(\Gamma)$.

4.2. Degeneration of hyperbolic structures. From the geometric point of view there is a particularly interesting subset of $X_n(\Gamma)$, namely the set $\mathcal{D}_n(\Gamma)$ of characters of discrete representations in the sense of 3.6.2. It follows from the discussion in 3.6.2 that the points of this set are in bijective correspondence with homotopy-hyperbolic structures on the space $K = K(\Gamma, 1)$. As I mentioned in 3.6.4, the Mostow rigidity theorem implies that when Γ is isomorphic to the fundamental group of a finite-volume hyperbolic n-manifold, $\mathcal{D}_n(\Gamma)$ is a single point.

4.2.1. The closure of $\mathcal{D}_n(\Gamma)$ in $X_n(\Gamma)$ is a compactification $\hat{\mathcal{D}}_n(\Gamma)$ of $\mathcal{D}_n(\Gamma)$. It was shown in [MSh1] for $n = 2, 3$, and in [M1] for all $n \geq 2$ that the points of $\hat{\mathcal{D}}_n(\Gamma) - \mathcal{D}_n(\Gamma)$ are small projectivized length functions; that is, we have $\hat{\mathcal{D}}_n(\Gamma) - \mathcal{D}_n(\Gamma) \subset \mathcal{SPL}(\Gamma)$.

Furthermore, the map $\Theta|\mathcal{D}_n(\Gamma)\colon \mathcal{D}_n(\Gamma) \to \mathcal{P}^{\Gamma}$ has direct geometric meaning in terms of the hyperbolic manifolds M_x. For any $x \in \mathcal{D}_n(\Gamma)$, the homogeneous coordinate of $\Theta(x)$ corresponding to a conjugacy class c in Γ is the length of the closed geodesic in M_x corresponding to c. The small length functions in $\hat{\mathcal{D}}_n(\Gamma) - \mathcal{D}_n(\Gamma)$ contain information about the growth of lengths of closed geodesics as a hyperbolic structure degenerates.

4.2.2. A famous example occurs when Γ is the fundamental group of a closed orientable surface Σ of genus $g > 2$, and $n = 2$. In this case $\mathcal{D}_2(\Gamma)$ is the set of 2-dimensional homotopy-hyperbolic structures on Σ, known to analysts as Teichmüller space and denoted $\mathcal{T}_g$. The points of $\hat{\mathcal{T}}_g - T_g$ are small projectivized length functions on $\pi_1(\Sigma)$, which by Skora's theorem are all defined by measured foliations on Σ. Thus $\hat{\mathcal{T}}_g$ is a natural compactification of $\mathcal{T}_g$ in which the ideal points are parametrized by measured foliations.

4.2.3. Just as the points of $\hat{\mathcal{D}}_n(\Gamma) - \mathcal{D}_n(\Gamma)$, where Γ is a finitely generated group, are defined by small actions on **R**-trees, so the actions on Λ-trees associated as in 4.1.7 with sequences of characters of discrete representations are small actions. These actions contain finer information about how hyperbolic structures degenerate.

The small actions on Λ-trees that are defined by sequences in Teichmüller space have been studied by Morgan and Otal. They are associated to generalized measured foliations[11] in which the transverse measure takes values in Λ. So these generalize measured foliations contain interesting asymptotic information about Teichmüller space.

4.2.4. The compactification described in 4.2.2 was first discovered by Thurston from a quite different point of view, and bears his name. Thurston used the space $\hat{\mathcal{T}}_g$ to study outer automorphisms of surface groups. Every outer automorphism α of $\pi_1(\Sigma)$ is known to be induced by a self-homeomorphism of Σ. There is an analysis of the action of α on $\hat{\mathcal{T}}_g$ which is similar to the analysis of the action of a real Möbius transformation on the compactified upper half-plane.

In particular, one shows that α always has a fixed point in $\hat{\mathcal{T}}_g$. In the case where the fixed point is in $\mathcal{T}_g$ the automorphism has finite order and can be completely understood. On the other hand, a fixed point in $\hat{\mathcal{T}}_g - \mathcal{T}_g$ is a projectivized length function which is defined by a measured foliation and is invariant under α. Using such invariant foliations one can describe the action of α.

The most interesting case is the one in which α has *two* fixed points in $\hat{\mathcal{T}}_g - \mathcal{T}_g$. In this case there are two mutually transverse, projectively invariant measured foliations whose projectivized length functions are α-invariant. This information can even be realized geometrically in the strongest imaginable way: α is induced by a homeomorphism η which leaves each of the two transverse foliations invariant and pulls backs each transverse measure to a constant multiple of itself. Such a homeomorphism η is called a *pseudo-Anosov map.* (Such maps were first studied by Thurston, who I believe named them in honor of Professor Ludwig von pseudo-Anosov, played by Sid Caesar.) The behavior of a pseudo-Anosov map with respect to the associated pari of transverse foliations leads to a rich theory of its dynamic behavior, which has become an exciting area of research.

4.2.5. There is an analogy between the compactification $\hat{Y}_n$ of outer space that I discussed in 2.1.1 and the Thurston compactification of Teichmüller space. The explanation for this analogy will become clear in 4.3. One goal of Culler and Vogtmann's program for studying $\mathrm{Out}(F_n)$ is to obtain a analysis for the action of an element of $\mathrm{Out}(F_n)$ on $\hat{Y}_n$ similar to Thurston's analysis for elements of the outer automorphism group of a surface group. The results of Bestvina and Handel that I discussed in 2.2.6 provide a step in this direction.

4.2.6. Since in general we have $\hat{\mathcal{D}}(\Gamma) - \mathcal{D}(\Gamma) \subset \mathcal{SPL}(\Gamma)$, the classification of small actions of an arbitrary finitely generated group Γ is a central question in the study of degenerations of hyperbolic structures, particularly in the case where Γ admits discrete faithful representations on $\mathrm{SO}(n, 1)$ for some n.

We saw in 2.3 that the fundamental group of a closed orientable surface Σ_g has a wealth of non-trivial small actions on **R**-trees. And we just saw that the length functions defined by these actions all appear in the Thurston boundary $\hat{\mathcal{T}}_g - \mathcal{T}_g = \hat{X}_2(\pi_1(\Sigma) - X_2(\pi_1(\Sigma)$. The opposite extreme occurs for a group Γ which admits *no* small non-trivial action on an **R**-tree. For such a group Γ, and for any $n \geq 2$, we have $\hat{\mathcal{D}}_n(\Gamma) - CalD_n(\Gamma) \subset \mathcal{SPL}(\Gamma) = \emptyset$; hence in this case the set $\mathcal{D}_n(\Gamma)$ is *compact.*

It follows that if Conjecture 2.5.5 is true, then for any finitely presented group Γ which admits no non-trivial splitting over a small subgroup, and for any integer $n \geq 2$, the space $\mathcal{D}_n(\Gamma)$ is compact. The condition in 2.5.5 that every small subgroup of Γ is finitely generated. is a harmless restriction here, because it is automatically satisfied whenever $\mathcal{D}_n(\Gamma) \neq \emptyset$, or more generally whenever Γ is isomorphic to a discrete subgroup of a Lie group. Of course if Question 2.5.3 had an affirmative answer in general, one could replace the hypothesis that Γ is finitely presented by the more natural and satisfactory hypothesis that it is finitely generated. This was the original motivation for Question 2.5.3.

4.2.7. It is worth pointing out that, since by 4.1.2 and 4.1.6 the set $\hat{X}(\Gamma) - X(\Gamma)$ consists of projectivized length functions of finite rank, for the application discussed in 4.2.6 it would be enough to prove Conjecture 2.5.5 for the finite-rank case; that is, to show that if Γ satisfies the hypotheses of 2.5.5 then every finite-rank

[11] Actually their results are stated in terms of laminations, not foliations. See the footnote to 2.4.

point of $\mathcal{SPL}(\Gamma)$ is a limit of points of $\mathcal{SPL}(\Gamma)$ defined by small integer-valued length functions. I shall return to this case of Conjecture 2.5.5 in 5.1 and 5.2.

4.2.8. The program outlined in 4.2.6 has been largely carried out in the case $n = 3$. Let Γ be a finitely generated group such that $X_3(\Gamma) \neq \emptyset$, i.e. Γ is isomorphic to a discrete subgroup of $\mathrm{SO}(3,1) = \mathrm{PSL}_2(\mathbf{C})$. For simplicity suppose that the group Γ is torsion-free, so that Γ is isomorphic to the fundamental group of a hyperbolic 3-manifold M. Again for simplicity, suppose that M is orientable.

In [MSh1], Morgan and I showed that Question 2.5.3 has an affirmative answer whenever Γ is a finitely generated group which arises as the fundamental group of an orientable 3-manifold M (possibly with boundary). Furthermore, in this case the result has topological meaning in terms of M.

Consider an arbitrary orientable 3-manifold M with finitely generated fundamental group. To simplify the language I'll assume that M is irreducible (3.8.5); this is automatically true if M is hyperbolic. According to a theorem first proved in complete generality by Scott [Sc], M has a *compact core*[12], i.e. there is a compact irreducible 3-manifold-with-boundary $M_0 \subset M$ such that the inclusion homomorphism $\pi_1(M_0) \to \pi_1(M)$ is an isomorphism. (In particular this implies that $\Gamma = \pi_1(M)$ is finitely presented, a result proved independently by Scott and myself.) Using the techniques of Stallings's that I referred to in 3.6.5, one can show that Γ admits a non-trivial small action on a $\mathbf{Z}$-treeif and only if M_0 contains a connected incompressible surface with a small fundamental group, or what is the same thing, with a non-negative Euler characteristic. What Morgan and I showed in [MSh2] is that if $\pi_1(M)$ admits a small non-trivial action on an $\mathbf{R}$-treethen M_0 contains such a surface. As I explained in [Sh], the proof involves extending Stallings's techniques to $\mathbf{R}$-trees, using codimension-1 *measured laminations* in place of surfaces, and applying a polynomial-vs.-exponential growth argument (similar to the one I described in 2.4.2) to approximate a lamination whose leaves have small fundamental groups by a surface of non-negative Euler characteristic.

When M is hyperbolic, the only possible connected incompressible surfaces of non-negative Euler characteristic in M_0 are disks and annuli. So the upshot, as far as hyperbolic geometry is concerned, is that if M is an orientable hyperbolic 3-manifold whose compact core contains no incompressible annuli or disks, then the space $D_3(\pi_1(M))$ is compact. By 3.6.2, this conclusion may be re-interpreted as saying that the space of all 3-dimensional homotopy-hyperbolic structures on M is compact. This result was first proved from an entirely different point of view by Thurston.

Note that although the project carried out in [MSh1] involved answering Question 2.5.3 for the case of a 3-manifold group, it did not lead to a proof of Conjecture 2.5.5 in this case. We showed that if $\pi_1(M)$ admits a non-trivial small action on an $\mathbf{R}$-treethen it admits a non-trivial small action on a $\mathbf{Z}$-tree; but we did not show that the given action is a limit of small simplicial actions. On the other hand, Morgan and Otal have good partial results on Conjecture 2.5.5 for 3-manifold groups.

4.2.9. Both the results of Thurston's that I have mentioned above—his classification of outer automorphisms of surface groups and his criterion for the compactness of the space of homotopy-hyperbolic structures on a 3-manifold— played central roles in his celebrated work on the existence of hyperbolic structures on 3-manifolds, which I discussed in 3.6. As I have explained, important components of both these theorems can be recovered through the study of group actions on $\mathbf{R}$-trees. This was my original excuse for getting interested in $\mathbf{R}$-trees.

4.3. The metric space approach; outer automorphism groups. There is a different approach to the theory that I described in 4.1 and 4.2.1. In place of algebro-geometric valuations and Tits buildings, this approach uses metric space geometry. The possibility of such an approach was suggested by Gromov and by Thurston. It has been carried out by Bestvina [Be] and, using a somewhat different point of view, by Paulin [Pau2]. Paulin's approach, which I will be describing, provides generalizations that apply not only to the study of hyperbolic manifolds, but to the theory of hyperbolic groups in the sense of Cannon and Gromov, which is discussed extensively elsewhere in this volume.

4.3.1. A representation of a group Γ in $\mathrm{SO}(n,1)$ can be interpreted as an action of Γ by isometries on H^n. So in the compactification $\hat{X}_n$, the points of X_n are described by actions on H^n, whereas the points of $\hat{X}_n - X^n$ are described by actions of Γ on $\mathbf{R}$-trees. The common feature is that they are actions of Γ on metric spaces. Paulin's approach to the compactification is based on some general considerations involving Γ-metric spaces,

[12]This is a more general notion than the one I referred to in 3.6.3.

where Γ is a given group; here by a Γ-metric space I mean a (real) metric space equipped with an action of Γ by isometries. Let's restrict attention to Γ-metric spaces having, say, at most the cardinality of the continuum, so that all equivariant isometry classes of Γ-metric spaces form a set, which I'll denote $\mathcal{U} = \mathcal{U}(\Gamma)$. Paulin begins by defining a topology on $\mathcal{U}$. The definition of this topology was suggested by F. Bonahon and is based on ideas due to Gromov and Thurston.

Let Y be a Γ-metric space. Given a finite set $K \subset Y$, a finite set $P \subset \Gamma$ and a positive number ϵ, we define a set $V(K, P, \epsilon) \subset \mathcal{U}$ as follows. A Γ-metric space Y belongs to $V(K, P, \epsilon)$ if and only if there is a map $f\colon K \to Y)$ such that

(i) For any two points $x_1, x_2 \in K$ we have $|\text{dist}\ (f(x_1), f(x_2)) - \text{dist}\ (x_1, x_2)| < \epsilon$, and

(ii) For any point $x \in K$, and any element $\gamma \in P$ such that $\gamma(x) \in K$, we have $\text{dist}\ (f(x), \gamma(f(x))) < \epsilon$.

If we let K and P vary over all finite subsets of Y and Γ, and ϵ over all positive numbers, the sets $V(K, P, \epsilon)$ satisfy the axioms for a basis of neighborhoods of Y in $\mathcal{U}$. In this way one defines a topology on $\mathcal{U}$.

In general this topology is pretty nasty: it is not even Hausdorff. But certain interesting subsets of $\mathcal{U}$ inherit nice subset topologies.

4.3.2. As a first example, consider the set $\mathcal{T}_0 = \mathcal{T}_0(\Gamma) \subset \mathcal{U}(\Gamma)$ consisting of all (equivariant isometry classes of) minimal, non-abelian actions of Γ on $\mathbf{R}$-trees. I have already implicitly described a topology on this set: in 1.5.4 I stated the result of Culler-Morgan and Alperin-Bass giving a bijective correspondence between this set and the set $\mathcal{L}_0(\Gamma) \subset \mathcal{L}(\Gamma)$ consisting of all non- abelian length functions. Of course $\mathcal{L}_0(\Gamma)$ inherits a topology from the product topology of $(0, \infty)^{\mathcal{C}(\Gamma)}$. Paulin shows in [Pau1] that the pull-back of this topology to $\mathcal{T}_0$ coincides with the subspace topology on $\mathcal{T}_0 \subset \mathcal{U}$. Paulin uses this to give an alternative proof of Culler and Morgan's result (see 2.1.2) that when Γ is not small, the space $\mathcal{SPL}(\Gamma)$, which is a subset of the image of $L_0(\Gamma)$ in $\mathcal{PL}(\Gamma)$, is compact.

One can also consider a slightly larger set than $\mathcal{T}_0(\Gamma)$, namely the set $\mathcal{T}(\Gamma) \subset \mathcal{U}(\Gamma)$ consisting of all minimal actions of Γ on $\mathbf{R}$-trees which are *semi-simple* in the sense that they are not exceptional abelian actions (1.6). It follows from 1.5.4 that there is a natural bijection between $\mathcal{T}(\Gamma)$ and $\mathcal{L}(\Gamma)$. I believe that the methods of [Pau1] also allow one to show that the topology induced on $\mathcal{T}(\Gamma)$ by this bijection coincides with the subspace topology on $\mathcal{T} \subset \mathcal{U}$, and to give a proof from this point of view of Culler and Morgan's result (see 1.7.1) that $\mathcal{PL}(\Gamma)$ is compact.

4.3.3. These compactness arguments are based on principles of which the applicability extends far beyond the case of $\mathbf{R}$-trees. From the point of view of Gromov's theory, $\mathbf{R}$-trees are 0-hyperbolic spaces. In [Pau2,3], a general compactness criterion is established for subsets of $\mathcal{U}(\Gamma)$ consisting of spaces that are hyperbolic in Gromov's sense. It is most conveniently stated (and proved) as a *sequential* compactness criterion, and this seems to cover all interesting applications.

Let (Y_i) be a sequence of Γ-metric spaces. Let $(\delta_i)_{i\geq 0}$ be a convergent sequence of non-negative numbers, and set $\delta = \lim \delta_i$. Suppose that Y_i is δ_i-hyperbolic for each i. Let $y_i \in Y_i$ be a base point for each i. Suppose that for every finite set $P \subset \Gamma$ and every $\epsilon > 0$ there exists an integer $N > 0$, such that for each $i \geq 0$ the closed convex hull of the set $P \cdot y_i \subset Y_i$ can be covered by at most N balls of radius ϵ. Then there is a subsequence of the Y_i that converges in $\mathcal{U}(\Gamma)$, and the limit is 50δ-hyperbolic.

(Here the closed convex hull of a subset $S \subset Y_i$ is defined to be the smallest closed convex subset of Y_i containing S. To say that a subset is convex means that every geodesic whose endpoints lie in the subset is itself contained in the subset.)

4.3.4. Using the compactness criterion 4.3.3, it is possible to recover the compactification of $X_n(\Gamma)$ that I discussed in 4.1. Let Γ be a finitely generated group with generators $u_1, \ldots, u_m$. Let (x_i) be an unbounded sequence of points in $X_n(\Gamma)$ for some $n \geq 2$. Each x_i is determined by a representation $\rho_i\colon \Gamma \to \mathrm{SO}_(n, 1)$. For each i and each $z \in H^n$ set $A_i(z) = \max_{1\leq j\leq m} \text{dist}\ (z, \rho_i(u_j)(z))$. One can show that for each i there is a point $y_i \in H^n$ where the function A_i takes a smallest value λ_i. Since (x_i) is unbounded, we can assume after passing to a subsequence that $\lambda_i \to \infty$. Now let X_i denote the metric space whose underlying set is H^n, with the distance function obtained by multiplying hyperbolic distance by λ_i^{-1}. Then X_i is λ_i^{-1}-hyperbolic for each i. For each i the representation ρ_i gives X_i the structure of a Γ-metric space. Using the base points y_i, and making strong use of hyperbolicity, it is possible to verify the hypotheses of the compactness criterion of 4.3.3. This means that after passing to a subsequence we can arrange that the X_i converge in $\mathcal{U}(\Gamma)$ to a

Γ-metric space T which is 0-hyperbolic, i.e. is an $\mathbf{R}$-tree. The action of Γ on T defines a length function l which is easily seen to be non-zero, so that we have a point $l \in \mathcal{PL}(\Gamma)$. One can then show that $(\Theta(x_i))$ converges to $[l]$ in $\mathcal{P}^\Gamma$.

One can also show from this metric space picture that if the x_i belong to $\mathcal{D}_n(\Gamma)$ then the action defining $[l]$ has small segment stabilizers (see 4.2.1). So the main properties of the compactifications $\hat{X}_n$ and $\hat{\mathcal{D}}_n$ can be established from this alternative point of view.

4.3.5. Because the metric-space approach to the compactification of the character variety sidesteps the use of valuations, it can be applied in situations where no algebraic variety is present. This was done by Paulin in his work [Pau3] on the outer automorphism group of a Gromov-hyperbolic group. He proved that if Γ is Gromov-hyperbolic, and if Γ admits no non-trivial small action on an $\mathbf{R}$-tree, then $\mathrm{Out}(\Gamma)$ is finite. It follows that if Conjecture 2.5.5 is true, then any finitely presented Gromov-hyperbolic group which does not split over a small subgroup has a finite outer automorphism group. (The hypothesis from 2.5.5 that the small subgroups of Γ are finitely generated is automatically satisfied by a hyperbolic group. Indeed, Gromov showed that every small subgroup of a hyperbolic group is cyclic-by-finite.)

Paulin's proof of the above finiteness theorem is very similar to his approach to the compactification of $X_n(\Gamma)$ that I outlined in 4.3.4. If $\mathrm{Out}(\Gamma)$ is infinite, it contains a sequence (α_i) of distinct elements. Let Y denote the Cayley graph of Γ with respect to generators $u_1, \ldots, u_m$, and for each i let $\rho_i \colon \Gamma \to \mathrm{Out}(\Gamma)$ be defined by $\rho_i(\gamma)(y) = \alpha_i(\gamma) \cdot y$. As in 4.3.4, for each i one can find a point $y_i \in Y$ where the function $A_i(z) = \max_{1 \le j \le m} \mathrm{dist}\ (z, \rho_i(u_j)(z))$ takes a smallest value λ_i; here dist denotes the word metric on Y. Using the fact that the α_i are all distinct one can show that Since (x_i) is unbounded, we can assume after passing to a subsequence that $\lambda_i \to \infty$. Now let Y_i denote the metric space whose underlying set is Y, with the distance function obtained by multiplying the word metric on Y by λ_i^{-1}. For each i, the representation ρ_i gives X_i the structure of a Γ-metric space. Using hyperbolicity one checks that the conditions of the compactness criterion 4.3.3 hold, so that some subsequence of the (Y_i) converges to an $\mathbf{R}$-tree with an action of Γ. Again this action can be shown to be non-trivial and to have small segment stabilizers.

4.3.6. It is clear from the above argument that if Γ is a hyperbolic group such that $\mathrm{Out}(\Gamma)$ is infinite, then the study of $\mathrm{Out}(\Gamma)$ is closely related to the study of actions of Γ on $\mathbf{R}$-trees. This explains the role of $\mathbf{R}$-trees in the study of outer automorphisms of a free group (see 2.1) and of a surface group (see 4.2.4). Furthermore, the analogy between the Culler-Vogtmann compactification of outer space and the Thurston compactification of Teichmüller space is made clear by 4.3.3–4.3.5: in both cases actions on $\mathbf{R}$-trees arise as limits of actions of a group on δ-hyperbolic spaces as $\delta \to 0$.

4.3.8. Other intriguing approaches to the theory of compactifying character varieties described in 4.1 and 4.2.1 have been worked out by Basarab [Bab] from the point of view of model theory, and by Chiswell [Ch] using non-standard methods. These approaches use a logical perspective to clarify or simplify the approaches that I have described above. Still another very elegant approach, due to Brumfiel and based on the theory of ordered fields, combines some features of the valuation approach of 4.1 and the metric space approach of 4.3, and to some extent clarifies the relationship between them.

But I think a mystery remains. The theory of Bruhat-Tits buildings for algebraic groups over valued fields and the theory of Gromovian convergence of metric spaces both have wide applicability. When two theories have a common special case they often have a common generalization. I wonder if there is some general picture that includes both the Bruhat-Tits tree over an arbitrary valued field (such as the p-adic numbers), and the theory described in 4.3, as special cases. My feeling is that this might be conceptually useful, and could perhaps even lead somewhere.

4.4. Further thoughts. The material presented in this section leads to lots of interesting research questions. For one thing, of course, it provides additional motivation for the questions and conjectures of 2.5, which I will be discussing in Section 5. Here I would like to mention a few other natural directions for further research.

4.4.1. I explained in 4.2.8 how Morgan and I proved Thurston's compactness theorem in [MSh2] using $\mathbf{R}$-trees. In his work on the existence of hyperbolic structures, Thurston used a generalization of this compactness theorem which applies to manifolds that do contain incompressible annuli. He gave an ingenious argument which reduced it to the compactness theorem stated in 4.2.8, or more precisely to a relative version of the latter theorem which is proved by the same method (and was proved in [MSh2] using **R**-trees). However, for reasons that I shall explain, it is very natural to try to proved the generalization directly using trees.

The generalized compactness theorem is stated in terms of the characteristic submanifold theory of [Jo] and [JaS]. The latter theory provides a picture of the incompressible annuli in a compact 3-manifold M. When M is the compact core of a hyperbolic 3-manifold, M contains a canonical submanifold Σ, each component of which is either an interval bundle meeting ∂M in the associated 0-sphere bundle or a solid torus meeting ∂M in a family of disjoint annuli. The components of the frontier of Σ are incompressible annuli, and every incompressible annulus in M is isotopic to one contained in Σ.

The generalized compactness theorem asserts that if A is any component of $M - \Sigma$, the restriction map $X_3(\pi_1(M)) \to X_3(\pi_1(A))$ maps $\mathcal{D}_3(\pi_1(M))$ to a set with compact closure.

In order to prove this, one has to show that if (x_i) is a sequence ending to infinity in $\mathcal{D}_3(\pi_1(M))$, then (x_i) has a subsequence whose image in $X_3(\pi_1(A))$ converges. By the construction of 4.1.7 we can associate with (x_i) a small action of $\mathcal{D}_3(\pi_1(M))$ on a Λ-tree, where Λ is some ordered abelian group of finite **Q**-rank. To prove the theorem it suffices to show that the restricted action of $\mathcal{D}_3(\pi_1(M))$ is trivial.

What is tantalizing is that Morgan and I proved in [MSh1] that for any small action of $\pi_1(M)$ acts on an **R**-tree, and any component A of $M - \Sigma$, the restricted action of $\pi_1(A)$ does have a fixed point. This is proved by a refinement of the argument that I mentioned in 4.2.8 and discussed in [Sh]. But our proof does not work for Λ-trees when Λ has order-rank greater than 1. In fact the general statement about actions seems to become false in this case. Nevertheless, it seems that there ought to be some way to adapt this approach to give a direct proof of the generalized compactness theorem in terms of trees.

4.4.2. The Cyclic Surgery Theorem, which I discussed in 3.8, gives strong information about how surgery on a knot in S^3 can yield a 3-manifold with a cyclic fundamental group. As I mentioned in 3.8.1, every closed orientable 3-manifold can be obtained from S^3 by a sequence of Dehn surgeries. Each surgery involves a solid torus which is the tubular neighborhood of some knot. An equivalent point of view is to think of all the surgeries as being done simultaneously by removing a finite union of disjoint solid tori from S^3 and sewing them back differently. Thus any closed orientable manifold can be obtained by Dehn surgery on a *link*, i.e. a finite disjoint union of knots in S^3.

If one could formulate and prove an analogue of the Cyclic Surgery Theorem for links, it might be useful in connection with the difficult problem of classifying orientable 3-manifolds with cyclic fundamental group. (This problem includes the Poincaré Conjecture, which asserts that any closed simply-connected 3-manifold is homeomorphic to S^3.)

Recall from 3.8.7 that the proof of the Cyclic Surgery Theorem, in the crucial case of a hyperbolic knot, involves looking at the curve in the $\mathrm{PSL}_2(\mathbf{C})$-character variety of the knot group containing the character of the discrete representation; and that one associates actions of the knot group on **Z**-trees, and hence incompressible surfaces in the knot manifold, with the ideal points of the curve, via the theory described in 3.4 and 3.6. The actions that arise in this way are not small, since the curve contains only one point which is the character of a discrete representation. Correspondingly, the incompressible surfaces that come up do not have small fundamental groups.

If we replace the hyperbolic knot by a hyperbolic link with n components, the irreducible component of the character variety containing the character of the faithful representation becomes n-dimensional. We have already seen that the theory described in 4.3 is the natural generalization to higher-dimensional varieties of characters of the theory described in 3.4, and that it involves considering **R**-trees in place of **Z**-trees. In [MSh2], the connection between trees and surfaces described in 3.6 was largely generalized to **R**-trees. In place of incompressible surfaces one uses incompressible measured laminations. While the immediate goal of [MSh2] was to apply the machinery to the small actions that arise as limits of discrete representations, the machinery is in principle of wider applicability. So one possible approach to generalizing the Cyclic Surgery Theorem to links would be to use this machinery.

I don't know whether this can be done. I mention it as an example of what a potentially rich subject dendrology seems to me to be.

Section 5. Free and small actions on R-trees

In 2.5.1 and 2.5.5 I stated two conjectures about actions of groups on **R**-trees. In 4.2.6 and 4.3.5 I illustrated the implications of Conjecture 2.5.5 for the geometry of hyperbolic manifolds and the study of outer automorphism groups. In this section I will summarize some of the existing evidence for Conjectures 2.5.1 and 2.5.5. This evidence, as I shall argue, actually suggests a stronger conjecture, which would imply both of the Conjectures 2.5.1 and 2.5.5. This stronger conjecture has only recently occurred to me, and its present form is somewhat tentative; I will state it in 5.5.6.

5.1. Some known results. As I explained in 4.2.8, Morgan and I showed in [MSh2] that Question 2.5.3 has an affirmative answer when the finitely generated group Γ is the fundamental group of an orientable 3-manifold. We also proved Conjecture 2.5.1 for the case of a 3-manifold group; alternatively, this can be deduced via the results of [MSk] from the affirmative answer to Question 2.5.3 in the 3-manifold group case, using the fact that a finitely generated 3-manifold group is finitely presented (see 4.2.8). This special case of Conjecture 2.5.1 is far from vacuous, because any free product of surface groups and free abelian groups of rank ≤ 3 is in fact the fundamental group of an orientable 3-manifold M. Indeed, we can construct M as a connected sum of finitely many 3- manifolds, each of which is either an interval bundle over a surface or a product of circles and arcs.

The results of Morgan and Skora in [MSk], which I discussed in 2.5.4, not only relate Conjectures 2.5.1 and Question 2.5.3, but also give direct evidence for Conjecture 2.5.1, since they show that it is true for any group that splits over a small subgroup.

Another source of evidence for the Conjectures 2.5.1 and 2.5.5 is provided by the results in my joint papers [GiS1] with Gillet and [GiSSk] with Gillet and Skora. These papers deal with actions on Λ-trees, where Λ is an arbitrary subgroup of $\mathbf{R}$ whose $\mathbf{Q}$-rank is at most 2. The $\mathbf{R}$-completion of such an action is an action of Γ on an $\mathbf{R}$-tree; this completed action has rank ≤ 2 in the sense of 4.1.2.

It follows from the results of [GiS1] and [GiSSk] that Conjectures 2.5.1 and 2.5.5 are true if one restricts attention to actions of rank at most 2. More precisely, if a finitely generated group Γ admits a free action of rank ≤ 2 on an $\mathbf{R}$-tree, then Γ is a free product of surface groups and infinite cyclic groups. (If the $\mathbf{Q}$-rank is 1 then Γ is actually free.) If Γ is a finitely presented group whose small subgroups are all finitely generated, then any projectivized length function on Γ defined by a small action of rank 2 is the limit of a sequence of projectivized $\mathbf{Z}$-valued length functions. We saw in 4.2.7 that the finite-rank case of Conjecture 2.5.5 is particularly important for applications, so it seems encouraging that the case of rank ≤ 2 is true.

5.2. Strong convergence, standard actions and the ascending chain condition. The proofs of Conjectures of 2.5.1 and 2.5.5 in the rank-2 case use the main result of [GiS1], which is a structure theorem for a large class of actions on Λ- trees, where $\Lambda \leq \mathbf{R}$ has $\mathbf{Q}$-rank ≤ 2. Some recent evidence suggests that a similar structure theorem may hold with no restriction on the rank of the action. The ultimate goal of this section is to formulate the appropriate conjecture. I will begin by explaining the statement of the structure theorem[13] that is proved in [GiS1].

5.2.1. The theorem applies to actions of a group Γ on a Λ-tree (where $\Lambda \leq \mathbf{R}$ has $\mathbf{Q}$-rank ≤ 2) that satisfy the following *ascending chain condition.*

> *If $\sigma_1, \sigma_2 \dots$ is a monotone decreasing sequence of segments in T with a common midpoint, and if Γ_i denotes the stabilizer of σ_i in Γ, then for all sufficiently large i we have $\Gamma_i = \Gamma_{i+1}$.*

5.2.2. The form of the theorem is that if $\Lambda \leq \mathbf{R}$ has $\mathbf{Q}$-rank ≤ 2, then any action on a Λ-tree is a limit in a strong sense—much stronger, in the case of an action of a finitely presented group, than the sense of 2.2.7—of actions of a standard type. These standard actions constitute a common generalization of two types of actions that I have discussed in previous sections: (i) the polyhedral actions discussed in 1.2.10, and (ii) the actions described in 2.3.7 which are associated to length systems on triangulated surfaces (or equivalently are dual to measured foliations on surfaces).

Note that actions of both types (i) and (ii) are constructed from 1-connected simplicial complexes in which a length is assigned to every 1-simplex. In case (i) the complex is a simplicial tree. Thus the link of every vertex is 0-dimensional. In case (ii) the complex is the universal covering $\tilde{\Sigma}$ of the given surface Σ; thus Σ is itself a triangulated surface, so that the link of every vertex is a 1- sphere. In general the actions

[13] In [Sh] I stated some of the consequences of the structure theorem. However, the statement of the structure theorem itself was not in final form when [Sh] was written.

that are to be taken as standard are defined in terms of length systems on 1-connected complexes of a type that in [GiS1] are called singular surfaces.

5.2.3. A *singular surface* is, by definition, a simplicial complex $\tilde{\Sigma}$ of dimension 1 or 2 in which the connected components of the link of every vertex are points and combinatorial 1-manifolds. These 1-manifolds may be homemorphic to either S^1 or $\mathbf{R}$. If the link of every vertex of $\tilde{\Sigma}$ is a connected combinatorial 1-manifold, $\tilde{\Sigma}$ is called a *surface with points at infinity.* In this case, the vertices whose links are non-compact, i.e. are homeomorphic to $\mathbf{R}$, are called *points at infinity.* If $\tilde{\Sigma}$ is a surface with points at infinity, the complement of the set of points at infinity in $\tilde{\Sigma}$ is a topological 2-manifold.

5.2.4. Any discrete subgroup Γ of $\mathrm{PSL}_2(\mathbf{R})$ leads to a natural example of a simply-connected surface $\tilde{\Sigma}$ with points at infinity. As a topological space, we define $\tilde{\Sigma}$ to be the union of H^2 with the set of all fixed points in S^1_∞ of parabolic elements of Γ. There always exists a Γ-invariant triangulation of $\tilde{\Sigma}$ in which every 1-simplex is an arc contained in the closure in $\bar{H}^2 = H^2 \cup S^1_\infty$ of a hyperbolic geodesic. With respect to this triangulation, $\tilde{\Sigma}$ is a surface with points at infinity, and its points at infinity are precisely the parabolic fixed points of elements of Γ. If Γ is a lattice, i.e. if $\mathrm{PSL}_2(\mathbf{R})/\Gamma$ has finite volume, then $\tilde{\Sigma}$ is finite modulo Γ, i.e. there are only finitely many Γ-orbits of simplices.

5.2.5. If X is an arbitrary 1-connected triangulated space, the *branches* of X are defined to be the connected components of the complement of the 0-skeleton in $\tilde{\Sigma}$. If $\tilde{\Sigma}$ is a 1-connected singular surface, each branch of $\tilde{\Sigma}$ is either a closed 1-simplex or a surface with points at infinity. Thus we think of 1-connected singular surfaces as complexes obtained by gluing together 1-dimensional branches which are arcs and 2-dimensional branches which are surfaces with points at infinity, in some simply-connected pattern. If we use only 1-dimensional branches, we get a simplicial tree.

5.2.6. Just as in the case of a surface, we define a *length system* on a singular surface $\tilde{\Sigma}$ to be a family (x_τ) of positive real numbers indexed by the 1-simplices of Σ, with the property that for each 2-simplex σ of Σ, we can label the edges of σ as τ, τ' and τ'' in such a way that $x_\tau = x_{\tau'} + x_{\tau''}$. Again we call τ the *long edge* of σ. Note that if $\tilde{\Sigma}$ is a tree, there is no restriction on the assignment of lengths to 1-simplices; thus a polyhedral tree can be thought of as a simplicial tree with a length system.

As in the non-singular case, any piecewise-linear path in $\tilde{\Sigma}$ has a well-defined length with respect to any given length system on $\tilde{\Sigma}$.

To generalize the non-degeneracy condition (ii) of 2.3.5 to length systems on singular surfaces, a smidgen of care is required. In the case where $\tilde{\Sigma}$ is a surface with points at infinity, the definition proceeds much as in 2.3.5: we define the order o_v of a vertex to be the number of 2-simplices incident to v whose long edges are not incident to v. Since v may be a point at infinity, the cardinal o_v may be finite or infinite. We define the given length system on Σ to be *non- degenerate* if $o_v \geq 2$ for every vertex v of Σ. Now let Σ be an arbitrary 1-connected[14] singular surface. Any length system on $\tilde{\Sigma}$ restricts to a length system on each branch of $\tilde{\Sigma}$. A length system on $\tilde{\Sigma}$ is said to be *non-degenerate* if its restriction to every 2-dimensional branch is non-degenerate.

5.2.7. As I illustrated in 2.3.9 for the non-singular case, a non-degenerate length system λ on a singular surface $\tilde{\Sigma}$ need not define a tree. As in the non-singular case, a we need an extra condition, namely that there is a group of simplicial homeomorphisms Γ of Σ such that λ is invariant under Γ and $\tilde{\Sigma}$ is finite modulo Γ. If this condition holds then, as in the non-singular case, any two points of $\tilde{\Sigma}$ are joined by a path of minimal length; again, this defines a pseudo-distance on $\tilde{\Sigma}$, and the associated metric space is an $\mathbf{R}$-tree$T$. Furthermore, the action of Γ on $\tilde{\Sigma}$ induces an action on T. For the present exposition I will define an action of a group Γ on an $\mathbf{R}$-treeto be *standard* if it is defined in this way.

5.2.8. It is natural to define a corresponding class of standard actions on Λ-trees, where Λ is any subgroup of $\mathbf{R}$. Let Γ act on a singular surface $\tilde{\Sigma}$ in such a way that $\tilde{\Sigma}$ is finite modulo Γ. Let $\tilde{\Sigma}$ be equipped with a length system which is Λ-valued in the sense that $x_\tau \in \Lambda$ for every 1-simplex τ of $\tilde{\Sigma}$. The action of Γ on $\tilde{\Sigma}$ defines a standard action on an $\mathbf{R}$-tree$T$, and by the definition of T there is a natural Γ-invariant map

[14]The definitions can be extended to the non-simply- connected case with some extra work, but the statement of the structure theorem involves only the simply-connected case.

$\chi\colon \tilde{\Sigma} \to T$. Let $X \subset T$ denote the image under χ of the 0-skeleton of Σ. It may be shown that the distance between any two points x and y of X is an element of Λ, so that the set

$$[x, y]_\Lambda = \{z \in T | \text{dist } (x, z) \in \Lambda\}$$

is isometric to an interval in Λ. It may also be shown that $T_0 = \bigcup_{x,y\in X}[x, y]_\Lambda$ is a Λ-tree; it is clearly invariant under the action of Γ. So if an action of Γ on $\tilde{\Sigma}$, and a length system, satisfy the above conditions, they determine an action of Γ on a Λ-tree T_0. I'll say that an action of a group on a Λ-tree is *standard* if it can be constructed in this way. The completion of a standard action on a Λ-tree is a standard action on an $\mathbf{R}$-tree.

5.2.9. As I have said, the gist of the main theorem of [GiS1] is that every action satisfying the hypotheses is a limit, in a strong sense, of standard actions. To define the appropriate notion of limit we define a *category of group actions on* Λ*-trees* where Λ is any ordered abelian group. The category of Λ-trees was defined in 1.7.3. Let's think of a group action on a Λ-tree as a triple $\mathcal{T} = (T, \Gamma, \rho)$, where T is a Λ-tree, Γ is a group and ρ is a homomorphism from Γ to the group of automorphisms of T. A *morphism* from an action (T, Γ, ρ) to an action (T', Γ', ρ') is a pair $\phi = (f, h)$, where $f\colon T \to T'$ is a morphism of Λ-trees and h is a group homomorphism, such that $\rho'(h(\gamma))(f(x)) = f(\rho(\gamma)(x))$ for every $x \in T$ and every $\gamma \in \Gamma$.

There is also a notion of standard morphism between standard actions. Suppose that $\mathcal{T}_1$ and $\mathcal{T}_2$ are standard actions. Then $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$ is defined by a 1-connected singular surface $\tilde{\Sigma}_i$, a non-degenerate λ-valued length system on $\tilde{\Sigma}_i$, and an action of Γ_i on $\tilde{\Sigma}_i$, such that $\tilde{\Sigma}$ is finite modulo Γ. Let χ_i denote the natural map from the 0-skeleton Y_i of $\tilde{\Sigma}_i$ to T_i. A morphism $phi = (f, h)\colon \mathcal{T}_1 \to \mathcal{T}_2$ is said to be *standard* if there is a continuous map $F\colon \tilde{\Sigma}_1 \to \tilde{\Sigma}_2$ such that (i) $F(\gamma \cdot x) = h(\gamma) \cdot F(x)$ for every point $x \in \Sigma$ and every $\gamma \in \Gamma$, (ii) $F(Y_1) \subset Y_2$ and (iii) $\chi_2 \circ (F|Y_1) = f \circ \chi_1$.

5.2.10. Now suppose that $(\mathcal{T}_i; \phi_{ij})$ is a direct system in the category of actions on Λ-trees. This means, first, that we have a family of actions $\mathcal{T}_i$ indexed by some filtered ordered set I. (For most real-life applications we can take I to be the natural numbers.) Second, whenever $i < j$ we have a morphism $\phi_{ij}\colon \mathcal{T}_i \to \mathcal{T}_j$, and we have $\phi_{jk} \circ \phi_{ij} = \phi_{jk}$ whenever $i < j < k$. For each i let's write $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)_{i\in I}$. The system $(\mathcal{T}_i; \phi_{ij})$ is said to *converge strongly* if for every $i \in I$ and for every segment $S \subset T_i$ there is an index $j \geq i$ such that the set $f_{ij}(S)$ is mapped isometrically into T_k by f_{jk} for every $k \geq j$. (Thus the segment S may be crumpled up (see 1.7.3) to a certain stage, but beyond some stage the crumpling stops.) In particular, for any $i \in I$ and any two points $x, y \in T_i$, the distance between $f_{ij}(x)$ and $f_{ij}(y)$ is independent of j for all sufficiently large $j \geq i$.

If the direct system $(\mathcal{T}_i; \phi_{ij})$ converges strongly, then there exist a tree T, and morphisms of Λ-trees $f_i\colon T_i \to T$ for all $i \in I$, such that (i) $f_j \circ f_{ij} = f_j$ whenever $i < j$, (ii) $T = \bigcup_{i\in I} f_i(T_i)$, and (iii) for any $i \in I$ and any two points $x, y \in T_i$ we have dist $(f_{ij}(x), f_{ij}(y)) = \text{dist}(f_i(x), f_i(y))$ for all sufficiently large $j \geq i$. The tree T and the maps f_i are unique up to isometry making all imaginable diagrams commute. Furthermore, if Γ is the direct limit group of the system Γ_i, and h_i denotes the natural homomorphism from Γ_i to Γ, there is a unique $\rho\colon \Gamma \to \text{Aut }(T)$ such that for each i the pair $\phi_i = (f_i, h_i)$ is a morphism from $\mathcal{T}_i$ to the action $\mathcal{T} = (T, \Gamma, \rho)$. The action $\mathcal{T}$ is called the *limit* of the strongly convergent system $(\mathcal{T}_i, \phi_{ij})$.

5.2.11. The limit of a strongly convergent direct system is in particular a direct limit in the category of actions. However, a direct system may well have a direct limit without converging strongly. An example is given by the sequence $(\mathcal{T}_i)_{i\geq 1}$, where $\mathcal{T}_i$ is the trivial action of the trivial group on a tree T_i; topologically, T_i is a cone on a 3-element set $\{x_i, y_i, z_i\}$, with cone point t_i, and the edges joining t_i to x_i, y_i and z_i have lengths $1 - \frac{1}{i}$, $\frac{1}{i}$ and $\frac{1}{i}$ respectively. See Figure 5.2.11.1.

Figure 5.2.11.1

For $i < j$ there is a unique morphism from T_i to T_j mapping x_i, y_i and z_i to x, y and z respectively. This defines a morphism $\phi_{ij}: T_i \to T_j$. The direct system $(T_i; \phi_{ij})$ is not strongly convergent since the distance from y_i to z_i approaches but never equals 0. However, the direct limit exists and is isometric to the unit interval.

5.2.12. Here is the main result of [GiS1]:

Let Λ be a subgroup of $\mathbf{R}$ and let $\mathcal{T}$ be an action on a Λ-tree. Suppose that either Λ has $\mathbf{Q}$-rank 1, or that Λ has $\mathbf{Q}$-rank 2 and the action $\mathcal{T}$ satisfies the ascending chain condition 5.2.1. Then $\mathcal{T}$ is the limit of a strongly convergent direct system $(\mathcal{T}_i, \phi_{ij})$, where the $\mathcal{T}_i$ are standard actions and the ϕ_{ij} are standard morphisms.

If the given action is on a countable tree we can take the strongly convergent direct system to be indexed by the natural numbers. (This is the case that is explicitly done in [GiS1], but the proof in general is essentially the same.)

5.2.13. Let me now sketch the arguments used in [GiS1] and [GiSSk] to deduce from the above structure theorem that Conjectures 2.5.1 and 2.5.5 are true in the case where the given action has rank at most 2. To deduce Conjecture 2.5.1 in this case, we observe that the given group admits a free action $\mathcal{T} = (T, \Gamma, \rho)$ on a Λ- tree, where $\Lambda \leq \mathbf{R}$ has $\mathbf{Q}$-rank at most 2. A free action automatically satisfies the ascending chain condition. The theorem therefore exhibits $\mathcal{T}$ as the limit of a strongly convergent system $(\mathcal{T}_i)$, where the standard actions $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$ must themselves be free. It then follows from the definition of a standard action that eahc group Γ_i admits a free action, with compact quotient, on a 1-connected singular surface. By convering space theory, Γ_i is the fundamental group of a compact singular surface, and is therefore a free product of infinite cyclic groups and surface groups. The group Γ is the direct limit of the Γ_i. It may be shown that a direct limit of groups, each of which is a free product of infinite cyclic groups and surface groups, is itself a free product of infinite cyclic groups and surface groups.

Note that this argument appears to establish a stronger result than Conjecture 2.5.1 in the case of $\mathbf{Q}$-rank ≤ 2, because in place of the free abelian factors predicted by 2.5.1 we have infinite cyclic factors. This is an illusion, however, because if Λ has $\mathbf{Q}$-rank at most n then it is elementary to show that any free abelian group which admits a free action of rank $\leq n$ must have rank at most n. Thus when $n = 2$ one expects to have free abelian factors of ranks 1 and 2. As a free abelian group of rank 2 is the fundamental group of a 2-torus, it can appear among the surface group factors.

To establish Conjecture 2.5.5 in the case of rank ≤ 2 one must show that if $\mathcal{T} = (T, \Gamma, \rho)$ is a small action on a Λ-tree, where $\Lambda \leq \mathbf{R}$ has $\mathbf{Q}$-rank at most 2 and Γ is a finitely presented group with the property that all its small subgroups are finitely generated, then the action is a limit, in the sense of 2.2.7, of small actions of Γ on $\mathbf{Z}$-trees. The hypothesis that the action $\mathcal{T}$ is small, together with the restriction on the small subgroups of Γ, guarantees that $\mathcal{T}$ satisfies the ascending chain condition. Using the countability of Γ we can reduce to the case where T is countable, so that the theorem exhibits $\mathcal{T}$ as the limit of a strongly convergent direct system $(\mathcal{T}_i, \phi_{ij})$ of small standard actions indexed by the positive integers. Let us write $\mathcal{T}_i = (T_i, \Gamma_i, \rho_i)$. Using the hypothesis that Γ is finitely presented, we can modify the system $(\mathcal{T}_i, \phi_{ij})$ so as to arrange that the Γ_i are equal to Γ and the group homomorphisms involved in the morphisms ϕ_{ij} are all the identity. It is then easy to show that T is the limit of the actions $\mathcal{T}_i$ in the sense of 2.2.7. This reduces the proof to the case where the given action $\mathcal{T}$ is standard. In this case the proof uses the ideas sketched in 2.3.11: one exhibits the length system defining the action $\mathcal{T}$ as a limit of constant multiples of integer-valued length systems. Each of these defines a small action on a $\mathbf{Z}$-tree, and by doing the approximation with some care one can prove that $\mathcal{T}$ is the limit of these actions.

5.2.14. In [Sh] I discussed a direct proof of Conjecture 2.5.1 for rank-2 actions, extracted from a preliminary version of [GiS1]. The proof of the structure theorem is a refinement of the latter argument. We are given an action $\mathcal{T} = (T, \Gamma, \rho)$, where T is a Λ-tree for some $\Lambda \subset \mathbf{R}$ of $\mathbf{Q}$-rank at most 2. In the case where Λ is free abelian of rank 2 one uses the contractible complex $K = K(T)$, defined in [Sh, Section 6], on which the automorphism group of the Λ-tree T acts.

Let us consider any connected subcomplex Y of K which is Γ-invariant and finite modulo Γ. One can use the ideas explained in [Sh, Section 6] to show that Y can be equivariantly deformed into a Γ-invariant singular surface Σ, also finite modulo Γ. Now by the definition of K, every 1-simplex of K corresponds to a segment of T; by associating to each 1-simplex of Σ the length of the corresponding segment in T one can define a Λ-valued length system on Σ. One can choose Σ so that this length system is non- degenerate. In general Σ need not be simply connected, but its universal cover $\tilde{\Sigma}$ inherits a Λ-valued length system. The action of Γ on Σ induces an action on $\tilde{\Sigma}$ of some group $\tilde{\Gamma}$ which is an extension of Γ by $\pi_1(\Sigma)$: that is, $\tilde{\Gamma}$ maps homomorphically onto Γ with kernel $\pi_1(\Sigma)$. Furthermore, $\tilde{\Sigma}$ is finite modulo $\tilde{\Gamma}$ and has an induced $\tilde{\Gamma}$-invariant, non-degenerate length system. By 5.2.7 and 5.2.8, the 1-connected singular surface $\tilde{\Sigma}$, with this length system and this action, defines a standard action of $\tilde{\Gamma}$ on a Λ-tree.

Let us write K as a monotone union[15] of subcomplexes Y_i that are finite modulo Γ. With each Y_i we can associate a standard action $\mathcal{T}_i$ by the above construction. This construction is not quite canonical, but by choosing the T_i with a bit of care one can arrange that there are natural morphisms that make the $\mathcal{T}_i$ into a direct system, and that this system converges strongly to $\mathcal{T}$.

The case where Λ has rank 2 is not free abelian requires a further refinement. In this case we need to write Λ as a monotone union of subgroups L_j that are free abelian of rank 2; with each of these subgroups we can associate a complex $K_j = K(T, L_j)$. These complexes are not contractible or even connected, but their homotopy-theoretic direct limit is contractible. One can then apply the above construction for each j; by doing this with some care one obtains a doubly indexed direct system of actions which converge strongly to $\mathcal{T}$.

5.3 Automorphisms of free groups revisited. Since the notion of strong convergence was introduced in [GiS1], it has gradually become clear that it arises naturally in other settings than that of rank-2 actions. I know of two different instances of this.

To explain the first instance, let me return to the example of the Bestvina-Handel theory that I discussed in 2.2.2–2.2.5, and re-examine it from a slightly different point of view. Recall from 2.2.2 that we thought of the Cayley graph T_0 of F_3 with respect to the generators x, y and z as a $\mathbf{Z}$-tree, and we considered the length function l_0 associated with the natural action of F on T_0. We associated a matrix A to the positive automorphism α, and considered the positive eigenvalue λ of A corresponding to its unique eigenvector v_0 in the first octant. We saw that the sequence $(\lambda^{-n} l_0 \circ \alpha^n)_{n\geq 0}$ converges to the length function of a non-polyhedral free action l_∞.

A slight variant of this approach is to regard the Cayley graph as a polyhedral tree T'_0 by assigning to each edge a length equal to one of the coordinates ξ_0, η_0 or ζ_0 of v_0, according to whether the given edge is labeled with the generator x, y or z. The action of F_3 on T'_0 defines a length function l'_0. The sequence $(\lambda^{-n} l'_0 \circ \alpha^n)_{n\geq 0}$ again converges to l_∞, but in a tamer way than $(\lambda^{-n} l_0 \circ \alpha^n)$: the arguments of 2.2.2–2.2.5 show that we have $l'0(W) = \lambda^{-n} l'_0 \circ \alpha^n(W) = l_\infty(W)$ for every cyclically legal word W and every $n \geq 0$, and that for an arbitrary $\gamma \in F$ we have $\lambda^{-n} l'_0 \circ \alpha^n(W) = l_\infty(W)$ for all sufficiently large n.

We can say more. For each $n \geq 0$ let T'_n denote the polyhedral tree obtained from T'_0 by multiplying the length of every edge by λ^{-n}. As a set, each T_i is identified with the Cayley graph of F_3, and it contains a copy X_i of the 0-skeleton of the Cayley graph, whose points are indexed by the elements of F_3. Let us write $X_i = \{x_i^\gamma \colon \gamma \in F_3\}$. For each $i \geq 1$ we can define a map $g_i \colon X_i \to X_{i+1}$ by $g_i(x_\gamma) = x_{i(\gamma)}$. It follows from our choice of lengths of edges in the polyhedral trees T_i that g_i extends to a morphism $f_i \colon T_i \to T_{i+1}$ which maps each edge homeomorphically onto a simplicial arc in T_i. We can now define a direct system $(\mathcal{T}_i)$ of actions indexed by the natural numbers: we have $\mathcal{T}_i = (T_i, F_3, \rho_i)$, where ρ_i is the natural action of F_3 on its Cayley graph, interpreted as an action on T_i; and the morphism from f_i to f_j when $j > i$ is $f_{j-1} \circ \cdots \circ f_i$. The arguments of 2.2.2–2.2.5 are easily adapted to show that this direct system is strongly convergent and that its limit is l_∞.

This is very striking, because the example of 2.2.2 gives the simplest example that I know of a non-polyhedral free action of a free group on an $\mathbf{R}$-tree, and it now turns out that this action is the limit of a strongly convergent direct system of polyhedral actions.

5.4. Contracting outer space and related spaces. The other situation in which strong convergence has arisen naturally is related to the contractibility of outer space, which I mentioned in 2.1.1. Several years after contractibility was proved by Culler-Vogtmann and Gersten, another proof was announced by Michael Steiner. Besides proving contractibility for Y_n he appears to have proved it for many other naturally

[15]If T is uncountable then K has uncountably many vertices. In this case we need a transfinite monotone union. Only the countable case was done explicitly in [GiS1].

defined subsets of $\mathcal{PL}(F_n)$, including $\hat{Y}_n$, $\mathcal{SPL}(F_n)$, and $\mathcal{PL}(F_n)$ itself. The method is to provide a geometric contraction of $\mathcal{PL}(F_n)$ which is so geometrically natural that it induces contractions of all these subsets.

I would like to describe briefly an elegant version of this theory which has been given by Richard Skora in [Sk2].

5.4.1. A key ingredient in this version of the theory is a construction for factoring any morphism between **R**-trees through a family of intermediate trees. Suppose that T and T' are **R**-trees and that $\phi\colon T \to T'$ is a surjective morphism. For each $t \in [0, \infty)$ we define a pseudo-distance $D_t\colon T \times T \to \mathbf{R}$ as follows. Let x and y be any two distinct points of T. Let us fix a homeomorphism $\alpha = \alpha_{xy}\colon [0,1] \to [x, y]$. Let $\mathcal{B}$ denote the set of all paths $\beta\colon [0,1] \to T'$ which are piecewise geodesics, i.e. are morphisms from the **R**-tree$[0, 1]$ to T', and which have initial point $\beta(0) = \phi(x)$ and terminal point $\beta(1) = \phi(y)$. Any $\beta \in \mathcal{B}$ is a rectifiable path in the metric space T and thus has a well-defined finite length. We have $\phi \circ \alpha \in \mathcal{B}$, and length $\phi \circ \alpha = \mathrm{dist}_T(x, y)$. Now let $\mathcal{B}_t$ denote the set of all paths $\beta \in \mathcal{B}$ which are uniform t-approximations to $\phi \circ \alpha$, i.e. which satisfy $\mathrm{dist}_{T'}(\beta(u), \phi \circ \alpha(u)) < t$ for every $t \in [0, 1]$. We define $D_t(x, y)$ to be the infimum of the lengths of all paths in $\mathcal{B}_t$. (This infimum can be shown to be realized as a minimum.) It is clear that $\mathcal{B}_0 = \{\phi \circ \alpha\}$ and hence that $D_0(x, y) = \mathrm{dist}_T(x, y)$. It is also clear that $\mathcal{B}_\infty = \mathcal{B}$ and hence that $D_\infty(x, y) = \mathrm{dist}_{T'}(\phi(x), \phi(y))$.

For each t, let T_t denote the metric space determined by the pseudo-distance D_t. Then T_0 and T_∞ are canonically identified with T and T'. Note also that whenever $0 \le s \le t \le \infty$ we have $D_t \le D_s$. Hence there is a natural distance-decreasing map $\phi_{st}\colon T_s \to T_t$. We have $\phi_{tu} \circ \phi_{st} = \phi_{su}$ whenever $s \le t \le u$, and $\phi = \phi_{0\infty}$.

In [Sk2] it is proved that the T_s are all **R**-trees and that the ϕ_{st} are all morphisms. If T and T' are equipped with (isometric) actions of a group Γ and if ϕ is Γ-equivariant, then by naturality we have an action of Γ on each T_s, and the ϕst are also Γ-equivariant.

5.4.2. In order to apply the above construction to prove contractibility results one needs a way of choosing a canonical base point in a given **R**-tree equipped with a non-abelian action of a group Γ with a given finite generating set S. In [Sk2] this is done by a construction similar to one that I mentioned in 4.3.4. Given an action of Γ on T and a point $x \in T$ we set $A(x) = \max_{\gamma \in S} \mathrm{dist}\ (x, \gamma \cdot x)$. It can be shown that the function $A(x)$ always takes a minimum value l_0 on T. The set X of all points $x \in T$ for which $A(x) = l_0$ is a subtree. If the action is non-abelian then X is finite. The barycenter (1.2.7) of X is a canonical base point of T.

5.4.3. Now suppose that we are given a point $[l] \in \mathcal{PL}(F_n)$. We can represent $[l]$ by a non-trivial minimal action of F_n on an **R**-tree$T$, and if $[l]$ is abelian we can take the action to be non-exceptional, i.e. we can take $T = \mathbf{R}$. If the action is non-abelian, then using the standard generators of F_n we can define a base point $x_0 \in T$ by the above construction. If the action is abelian and non-exceptional we take x_0 to be an arbitrary point of T. Now let T_0 denote the Cayley graph of F_n with respect to the standard generators. Each edge e of T_0 is labeled with a generator u_e from the standard generating set, and we can give T_0 the structure of a polyhedral **R**-tree by assigning to e the length $\mathrm{dist}_T(x_0, u_e(x_0)$. Now T_0 has a natural base point—the vertex labeled by the identity element—and a natural action of F_n defined by left multiplication. It follows from the definitions of the lengths of edges in T_0 that there is a unique F_n-equivariant morphism $\phi\colon T_0 \to T$. Applying the construction of 5.4.1 with this choice of ϕ, we get a tree T_t with an F_n-action for each $t \in [0, \infty)$, and a morphism $\phi_{st}\colon T_s \to T_t$ whenever $s \le t$. The action of F_n on T_t defines a point $[l_t] \in \mathcal{PL}(F_n)$. Note that $[l_\infty] = [l]$ and that $[l_0]$ is defined by the natural polyhedral action of F_n on T_0.

We can now define a map $H\colon \mathcal{PL}(F_n)[0, 1] \to \mathcal{PL}(F_n)[0, 1]$ by $H([l], t) = [l_{t^{-1}-1}]$. In [Sk2] it is shown that H is continuous. By construction we have $H([l], 0) = [l]$ and $H([l], 1) \in \Delta$, where $\Delta \subset \mathcal{PL}(F_n)$ consists of all projectivized length functions obtained from the standard action of F_n on its Cayley graph by assigning positive lengths to the (standard) generators. It is also straightforward to check that $H([l], 1) = [l]$ for each $[l] \in \Delta$, so that H is a deformation retraction of $\mathcal{PL}(F_n)$ to Δ. But by the definition of Δ there is a natural bijection between Δ and the quotient of the positive cone in $\mathbf{R}^n$ by homotheties. This bijection is a homeomorphism, so Δ is contractible. The contractibility of $\mathcal{PL}(F_n)$ follows. It is also not hard to show that the sets Y_n, $\hat{Y}_n$, $\mathcal{SPL}(F_n)$, and the set of all projectivized length functions defined by free actions of F_n are invariant under H; that is, if W denotes any of these sets we have $H(W \times [0, 1]) \subset W$. It follows that these sets are all contractible.

5.4.4. Now consider any small minimal action of F_n on an **R**-tree T. Such an action cannot be abelian. Hence the construction of 5.4.3 gives a family $(\mathcal{T}_t)_{t \ge 0}$ of actions of F_n on **R**-trees, and morphisms $\phi_{st}\colon \mathcal{T}_s \to \mathcal{T}_t$

for $s \leq t$. As I pointed out in 5.4.1, we have $\phi_{tu} \circ \phi_{st} = \phi_{su}$ whenever $s \leq t \leq u$; so the $\mathcal{T}_t$ for $t < \infty$ and the ϕ_{st} for $s \leq t \leq \infty$ constitute a direct system. It is easy to see from the construction of 5.4.1 that this direct system converges strongly and that its limit is the given action of F_n on T. If you prefer to think countably, you may think of the given action as the limit of the strongly convergent system $(\mathcal{T}_n)$, where n ranges over the natural numbers.

If the given action is a limit of free polyhedral actions, i.e. if its projectivized length function $[l_\infty]$ is in $\hat{Y}_n$, then for every t the projectivized length function $[l_t]$ defined by $\mathcal{T}_t$ lies in $\hat{Y}_n$. We have $[l_0] \in Y_n$. By the continuity of H, there is a smallest t, say $t = t_0$, for which $[l_t] \notin Y_n$. Thus the small action $\mathcal{T}_{t_0}$ is not polyhedral but it is the limit of a strongly convergent system of free polyhedral actions.

5.4.5. The construction in 5.3, which looked very different from the one just described in 5.4.4, also exhibited certain actions corresponding to points of $\hat{Y}_n - Y_n$ as limits of strongly converging systems of free polyhedral actions. The construction that I have just given certainly suggests that there should be a wealth of points in Y_n that correspond to actions which are limits of this type. I don't think every action corresponding to a point of $\hat{Y}_n - Y_n$ is such a limit. Some points of $\hat{Y}_n - Y_n$ are defined by actions dual to measured foliations, and such actions do not appear to be limits of strongly convergent systems of polyhedral actions. However, it may well be the case that for every action corresponding to a point of $\hat{Y}_n$, or more generally for every small minimal action of F_n, the system $(\mathcal{T}_t)_{0 \leq t < \infty}$ consists of actions that are standard in the sense of 5.2.7.

5.4.6. This leads to the conjecture that every small action of a free group is the limit of a strongly convergent direct system of actions which are standard in the sense of 5.2.7.

5.4.7. Some of Jiang's results in [Ji] are relevant to Conjecture 5.4.6. Jiang gives an intricate and ingenious proof that if a finitely generated free group acts freely and minimally on an $\mathbf{R}$-tree$T$, then there are only finitely many orbits of branch points under the action. Here a point $x \in T$ is called a *branch point* if there are two segments $S, S' \subset T$ whose interiors contain x but such that x is not an interior point of $S \cap S'$. Jiang also proves an estimate for the number of orbits of branch points; this estimate is best possible and can realized by a standard action. Skora has pointed out that if Conjecture 5.4.6 were true, one could deduce Jiang's finiteness result, and his estimate, as corollaries. So one may regard these results of Jiang's as additional evidence that Conjecture 5.4.6 is true.

5.5. A tentative conjecture. There is an obvious parallel between Conjecture 5.4.6 and the main theorem of [GiS1], which I stated in 5.2.12. This suggests that there may be a general structure theorem for group actions on $\mathbf{R}$-trees which satisfy the ascending chain condition. However, a simple-minded example shows that some care is needed formulating the appropriate statement. If a free action of a finitely generated group Γ on an $\mathbf{R}$-tree is the limit of a strongly convergent direct system of actions that are standard in the sense of 5.2.7, then the arguments that I sketched in 5.2.13 show that Γ is a free product of surface groups and infinite cyclic groups. But any free abelian group Γ is isomorphic to a subgroup of $\mathbf{R}$ and therefore acts freely on the $\mathbf{R}$-tree $T = \mathbf{R}$ by translations; and if Γ has rank > 2 then it is not a free product of surface groups and cyclic groups.

So we need to broaden the definition of "standard action." One may hope to prove that every group action on an $\mathbf{R}$-tree which satisfies the ascending chain condition is the limit of a strongly convergent direct system of actions which are standard in a suitably generalized sense. I would like to propose a tentative definition, which I have been worked out with the help of Marc Culler, Henri Gillet and Richard Skora.

5.5.1. Let X be a triangulated space. Let μ be a function that assigns a non-negative real number $\mu(\alpha)$ to every polyhedral path $\alpha\colon [0,1] \to X$. I'll call μ a *path measure* on X if (i) it is invariant under reparametrization, i.e. $\mu(\alpha \circ h) = \mu(\alpha)$ for any polyhedral path μ and any homeomorphism $h\colon [0,1] \to [0,1]$; and (ii) μ is additive under composition; that is, if α_1 and α_2 are polyhedral paths with $\alpha_1(1) = \alpha_2(0)$, their composition $\alpha_1 * \alpha_2$ satisfies $\mu(\alpha_1 * \alpha_2) = \mu(\alpha_1) + \mu(\alpha_2)$.

If Y is a subcomplex of X then any path measure on Y restricts to a path measure on Y.

5.5.2. If the triangulated space X consists of single closed edge, then for any positive real number λ there is a natural path measure on X: the length of a polyhedral path α is the ordinary length in $\mathbf{R}$ of the path $h \circ \alpha$, where h is an affine homeomorphism of X onto an interval of length λ in $\mathbf{R}$.

5.5.3. Let $\tilde{\Sigma}$ be a surface with points at infinity (5.2.3). Any non-degenerate length system on $\tilde{\Sigma}$ determines a path measure on $\tilde{\Sigma}$: for any polyhedral path α we define $\mu(\alpha)$ to be the length of α with respect to the given length system.

5.5.4. If X is a triangulated space and J is a piecewise-linear homeomorphism from X to a product $Y \times \mathbf{R}$, where Y is another triangulated space, then J defines a length function on X: for any path α we define $\mu(\alpha)$ to be the length in $\mathbf{R}$ of the path $pJ(\alpha)$, where $p: Y \times \mathbf{R} \to \mathbf{R}$ denotes projection to the second factor.

5.5.5. Let X be a 1-connected triangulated space. I will say that a path measure μ on X is *standard* if every branch (5.2.5) Y of X is either a closed edge or a surface with points at infinity, or is homeomorphic to a product; and the restriction of μ to Y is defined by one of the three corresponding constructions 5.5.2–5.5.4.

If every branch is of one of the first two types then X is a singular surface, and in this case a standard path measure is essentially the same thing as a non-degenerate length system.

5.5.6. Let X be a triangulated space with a standard path measure, and let Γ be a group of automorphisms of X such that X is finite modulo Γ. For any polyhedral path α, let's call $\mu(\alpha)$ the *length* of α. It should be easy to prove, generalizing 5.2.7, that any two points of X are joined by a path of minimal length, and that if we make X into a pseudo-metric space by defining the distance between two points to be the length of the minimal path between them, then the corresponding metric space is an $\mathbf{R}$-tree. The group γ then has an induced action on this tree. My tentative proposal is to define an action to be *standard* if it arises in this way.

As I have said, with this tentative definition comes a tentative conjecture, that every action satisfying the ascending chain condition is the limit of a strongly convergent direct system of standard actions. When I say that this is tentative, I mean that I would not be surprised if it were necessary to revise the definition of standard action slightly in order to make it true.

This tentative conjecture would imply Conjectures 2.5.1 and 2.5.2, via the same arguments that are used in [GiS1] and [GiSSk] (and sketched in 5.2.13 above) to deduce the rank-2 case of these conjectures from the main theorem of [GiS1].

5.6. Unsmall actions. Let me briefly discuss the question posed in 2.5.7, whether $\mathcal{PL}(\Gamma)$, where Γ is a finitely generated group, has a dense subset consisting of projectivized integer-valued length functions. What I believe to be true, at least, is that every action satisfying the ascending chain condition is a limit (in the sense of 2.2.7) of simplicial actions. In fact, in the rank-2 case this is proved in [GiSSk] by an argument almost identical to the one that I sketched in 5.2.13 for the small case. (In the rank-1 case, for example for a $\mathbf{Q}$-tree, one does not need to assume the ascending chain condition.) This argument would go through with no restriction on the rank if the tentative conjecture of 5.5.6 were known.

For actions not satisfying the ascending chain condition there seems to be rather little known about this question. Morgan and I proved in [MSh2] that Question 2.5.8 has an affirmative answer if Γ is a finitely generated fundamental group of a 3-manifold: that is, if such a group Γ admits a non-trivial action on an $\mathbf{R}$-tree, then it admits a non-trivial action on a $\mathbf{Z}$-tree. However, we did not show that the given action is a limit of simplicial actions.

In general there seems to be rather little known in general about actions on $\mathbf{R}$-trees not satisfying the ascending chain condition. In [GiS2], Gillet and I did prove the following result:

> *Let Γ be a group which acts without inversions on a Λ-tree, where Λ is a subgroup of $\mathbf{R}$. Then the cohomological dimension of Γ is at most $1 + r + d$, where r is the $\mathbf{Q}$-rank of Λ and d is the supremum of the cohomological dimensions of the stabilizers in Γ of points of T.*