Trees: with a view towards Outer space.

TCC Lectures, Hilary 2023

Ric Wade

February 24, 2023

Lecture 1

Simplicial trees and $\mathbb{R}$-trees

1.1 Overview

Housekeeping

If you are attending the course on MS Teams, I'd recommend keeping your camera on outside of the break. Personally, I find it very easy to drift off and lose concentration when I don't have the camera on, and I hope it will help you to feel more at home in the class and make it easier to ask questions. I won't, however, enforce this in any way (or judge if you just want to attend with the camera off!).

Also in the spirit of making everyone at home/welcome in the class, please feel free to ask questions. I think questions/answers are more important than necessary covering as much material as possible: I hope that these notes will partly compensate for any material that is not covered in the lectures.

After writing this I've noticed I keep using the rather Americanized word 'class' rather than lecture: this is mostly by accident, but I hope it does emphasize that this should be two hours a week where we talk to each other rather than I talk at you nonstop.

If you have any questions/comments about the course/notes please do drop me an email or you can ask after the lecture too.

What's the goal?

I initially started with the very vague goal of looking at 'Trees.' In order to give the course a bit more shape, I've added 'with a view towards Outer space,' as a subtitle. My viewpoint comes the world of Geometric Group Theory: the trees will typically be infinite and come equipped with an action by an infinite group, and the group will usually be a finitely generated discrete group. In this course, most of the trees will be simplicial, however some results will be given in the more general context of $\mathbb{R}$-*trees*, where the set of branch points can be dense, or the trees can even branch everywhere!

As we will see in lecture 2, group actions on trees arise from graphs of groups (which generalize amalgamated free products and HNN extensions). However they also appear 'naturally' in geometry as dual objects to codimension one subobjects. An example to think about is a simple close curve on a surface. This lifts to a collection of lines in the universal cover (which is the hyperbolic plane) and dual to this collection of lines is a tree which the fundamental group acts on (via deck transformation on the universal cover). Similar ideas work for surfaces in 3-manifolds, hyperplanes in cube complexes, etc.

Groups can act on trees in many different ways, which is where *spaces of group actions on trees* come into play. The most well-known of these is Outer space. Getting a basic understanding of the 'tree description' of Outer space, and proving it is contractible, is the main goal of the first few lectures. After that, I have some ideas: we could look at JSJ decompositions, for instance, but I'm also happy to steer the lectures in whatever direction the audience is most interested in.

1.2 Simplicial trees and $\mathbb{R}$-trees

Group actions on trees can either be approached in the simplicial/combinatorial setting, or in the more general setting of $\mathbb{R}$-trees. As the combinatorial setting is very often cleaner and easier to work in, we'll cover both.

Graphs á la Serre

A graph is a tuple

$$\Gamma = (E\Gamma, V\Gamma, {}^{-}\colon E\Gamma \to E\Gamma, \iota\colon E\Gamma \to V\Gamma),$$

where $E\Gamma$ is the edge set and $V\Gamma$ is the vertex set. There is an involution $^{-}$ on the edges such that $\bar{e} \neq e$ and $\bar{\bar{e}} = e$, and $\iota(e)$ is the *initial vertex* of every edge. Each edge also has a *terminal vertex* $\tau(e) = \overline{\iota(e)}$, which is useful even if it is not technically needed in the definition. If you have not seen this definition of a graph before, it has the surprising feature of having twice as many edges as one might expect - associated to each geometric edge in the picture we would draw of Γ is a pair $\{e, \bar{e}\}$ in $E\Gamma$. An *orientation* of a graph is therefore a subset $\mathcal{O} \subset E\Gamma$ containing exactly one element from each pair $\{e, \bar{e}\}$. A *subgraph* $\Delta \subset \Gamma$ is a pair $E\Delta \subset E\Gamma$, and $V\Delta \subset V\Gamma$ such that $E\Delta$ is preserved by edge inversion and $\iota(E\Delta) \subset V\Delta$.

An *edge path* is a either a single vertex (in which case the path is degenerate, or trivial), or a sequence

$$p = e_1, e_2, \ldots, e_n$$

such that $\tau(e_i) = \iota(e_{i+1})$ for all i. Initial vertices, terminal vertices, and inverses of paths are defined in the obvious way: i.e. $\iota(p) = \iota(e_1)$ and $\bar{p} = \overline{e_n}, \overline{e_{n-1}}, \ldots, \overline{e_1}$ (unless the path p is a single vertex v, in which case $\iota(p) = \tau(p) = \bar{p} = v$). A path is *reduced* if $e_{i+1} \neq e_i$ for all i, and is a *loop* if $\iota(p) = \tau(p)$. A graph is *connected* if there is an edge path between any two vertices.

Definition 1.2.1 (Combinatorial forests and trees)**.** A graph Γ is a *forest* if every reduced edge loop is trivial. A *tree* is a connected forest.

There is a unique reduced edge path between any two points in a tree (why?). In the lecture I defined a tree as a connected, simply connected, graph. I think it is a bit cleaner to use the above definition as then you don't have to define what the fundamental group is.

$\mathbb{R}$-trees

A *geodesic metric space* is a metric space (X, d) such that any points x and y are connected by an isometrically embedded arc (a *geodesic*) n X. One characterization/definition of $\mathbb{R}$-trees is of geodesic metric spaces where these arcs are unique and essentially the only way to get between x and y.

Definition 1.2.2 (Definition of $\mathbb{R}$-tree by the unique path property)**.** An $\mathbb{R}$-tree is a metric space (T, d) such that every pair of points x, y are connected by a unique geodesic $[x, y]$ and every embedded path f from x to y has image equal to $[x, y]$.

I'm not sure if this is the best definition. Here are some alternatives

- (Cut point/0-bottleneck property) An $\mathbb{R}$-tree is a geodesic metric space T that if $[x, y]$ is a geodesic in T and p is a point in the interior of $[x, y]$, then p separates x and y into distinct components of $T - p$.
- (0-slim) An $\mathbb{R}$-tree is a geodesic metric space T such that every geodesic traingle is 0-*slim*.
- (Relaxing the embedding condition in the first definition) An $\mathbb{R}$-tree is a geodesic metric space T such that geodesics are unqiue and any path (not necessarily embedded) from x to y contains the geodesic $[x, y]$ in its image.

To catch up on a definition: a geodesic triangle is δ-*slim* if any one side is contained in the δ-neighbourhood of the other two sides. In the context of trees, this means that all geodesic triangles are tripods.

Lemma 1.2.3 (Tripod lemma)**.** *If a, b, c are three points in an $\mathbb{R}$-tree there exists a unique point $m \in T$ such that*

$$[a, b] = [a, m] \cup [m, b]$$
$$[b, c] = [b, m] \cup [m, c]$$
$$[c, a] = [c, m] \cup [m, a]$$

In other words, the geodesic triangle given by the points a, b, c is a tripod with midpoint m.

Another immediate consequence of the definitions is that geodesics in trees have a strong local-to-global property: any arc that is locally a geodesic is a geodesic. One way to phrase this is as follows:

Lemma 1.2.4 (Local-to-global property)**.** *If $[x_0, x_1], [x_1, x_2], \ldots, [x_{n-1}, x_n]$ are nondegenerate geodesics and*

$$[x_{i-1}, x_i] \cap [x_i, x_{i+1}] = \{x_i\}$$

for all i, then their concatenation is equal to the geodesic from x_0 to x_n.

Lemma 1.2.5 (Closed subtree projection)**.** *It S is a closed subtree of T then every point $x \in T$ has a unique closest point $\pi_S(x) \in S$ satisfying $d(\pi_S(x), x) \leq d(s, x)$ for all $s \in S$.*

Sketch proof. As S is closed, points of minimum distance from S to x exist. Uniqueness is a consequence of the tripod lemma - if there are two points $a, b \in S$ that are closest to x, then by the tripod lemma, all points along the geodesic $[a, b]$ are closer to x than a and b, and belong to S as it is a subtree. It follows that $[a, b]$ must be degenerate and $a = b$. □

Morphisms and attempts at analysis

Combinatorial morphisms

A *combinatorial morphism* between two trees T and T' is a map sending each edge $e \in T$ to an edge path $f(e)$ in T', such that if $\iota(e) = \iota(e')$ then $\iota(f(e)) = \iota(f(e'))$. Note that the path $f(e)$ is allowed to be degenerate, in which case we say the morphism is degenerate. If all paths $f(e)$ are nondegenerate then we say the morphism is *nondegenerate*.

Metric morphisms

Doing the analogous thing for $\mathbb{R}$-trees is a little trickier. In this situation, a morphism is a map $f : T \to T'$ such that every arc $p \subset T$ there exists a covering $p_1, p_2, \ldots, p_n$ of p by subarcs and constants $\lambda_1, \lambda_2, \ldots, \lambda_n$ such that for any two points $x, y \in p_i$ we have

$$d_{T'}(f(x), f(y)) = \lambda_i d_T(x, y).$$

In other words, each subarc is stretched uniformly by f. Note that with this definition, morphisms need not be Lipschitz (are they even always continuous?), however in practice the λ_i's that appear will be uniformly bounded from above.

A morphism is *non-degenerate* if no nondegenerate arc in T is mapped to a point in T'.

Every morphism factors as a composition $T \to T''$ and $T'' \to T'$ where the first map is a forest collapse and the second map is nondegenerate.

Note that non-degenerate maps are not always injective. Some of the most important maps we will look at later on - *folds* will be nondegenerate but not injective.

Directions

The closest we have to a unit tangent bundle for a tree is the set/space of *directions* (although there aren't great topologies for the set of directions you can have a think about how to topologise this!). A *direction* d is a component of $T - x$ for a point $x \in T$. In this case we say that d is *based at* the point x. A point x is a *branch point* if there are at least three directions based at x, and is a *leaf* or an *end* of T if there is exactly one component of T based at x.

Any nondegenerate morphism from T to T' induces a map between the directions of T and the directions of T'.

Exercise 1.2.6. If the set of directions is the 'unit tangent bundle' of a tree, what should be the definition of the tangent bundle of a tree? Is there a reasonable way to topologise this?

Simplicial ℝ-trees

A simplicial ℝ-tree is one obtained from a simplicial tree by assigning a length λ_e to each edge such that the lengths are uniformly bounded below by some $\epsilon > 0$. Alternatively, one can ask that T is complete and the set of distances between branch points and/or leaves is bounded below by some $\epsilon > 0$. There is not really a gold standard definition of simplicial ℝ-tree in the literature but I hope this one is natural enough.

Three Tree Topologies

For any ℝ-tree, there are two topologies that commonly appear in the literature. The first is the *metric topology*, which is the topology induced by the metric on T. The second is the *visual topology*, which is defined to be the coarsest topology such that every direction in T is an open set.

For a simplicial tree, there is also the *simplicial topology*, where a set is defined to be open if an only if its intersection with every simplex of the tree (e.i. edge/vertex) is open. For historical reasons, this topology on simplicial/CW complexes is often called the weak topology. This naming has not aged well - many topologists would argue that in fact, this a rather strong topology:

Exercise 1.2.7. If T is a locally finite simplicial tree, show that all three tree topologies coincide on T. Furthermore, if T is not locally finite, show that visual topology is strictly coarser than the metric topology and that the metric topology is strictly coarser than the simplicial topology.

Lecture 2

Classification of isometries and group actions

In this lecture we want to discuss how individual isometries and how groups of isometries can act on a tree. We work in the more general setting of $\mathbb{R}$-trees for the duration. We mostly follow Culler–Morgan's very readable paper BLAH, although many of the ideas here appeared (in the simplicial setting) in Serre.

2.1 Displacement functions

Given a metric space X and an isometry $g\colon X \to X$, the *displacement* of g is defined by

$$\|g\|_X = \inf\{d(x, gx) : x \in X\}$$

The *characteristic set* of g is then the set of points where the minimal displacement is realized:

$$C_g = \{x \in X : d(x, gx) = \|g\|_X\}$$

In general, this set may be empty. If $C_g = \emptyset$ the isometry is called parabolic, and isometries of metric spaces fall into four classes:

- An isometry g is **elliptic** if $\|g\|_X = 0$ and $C_g = \text{Fix}(g) \neq \emptyset$;
- it is **neutral parabolic** if $\|g\|_X = 0$ and $C_g = \emptyset$;

- it is **hyperbolic/loxodromic** if $\|g\|_X > 0$ and $C_g \neq \emptyset$; and
- it is **non-neutral parabolic** if $\|g\|_X > 0$ and $C_g = \emptyset$

This language is not exactly standardized - in these notes as well as the broader literature hyperbolic and loxodromic are often used interchangeably. Non-neutral parabolic isometries are sometimes called *ballistic*. I will not focus at all on parabolic isometries, as the main goal of this section is to show that such isometries do not exist for group actions on trees.

2.2 The classification theorem

The following theorem has three important pieces:

- Isometries of trees are elliptic or hyperbolic
- $d(g, xg)$ depends only on $\|g\|_T$ and the distance from x to C_g
- If g is loxodromic then C_g is a line, called the *axis* of g.

Here is the full statement:

Theorem 2.2.1 (Classification of isometries). *If g is an isometry of a tree T then C_g is a nonempty closed subtree. In particular, every isometry is either elliptic or loxodromic. Furthermore*

$$d(x, gx) = \|g\|_T + 2d(x, C_g)$$

for all $x \in T$. If g is loxodromic then C_g is a line on which g acts by translation by $\|g\|_T$.

Proof. We first consider the case when g fixes a point. Note that C_g is always closed for an isometry of any metric space as isometries are continuous. If x and y are fixed by g then uniqueness of geodesics implies $[x, y]$ is fixed by g. Hence C_g is a subtree of T. As C_g is closed we can consider the closest point projection $\pi : T \to C_g$. The isometry sends $[\pi(x), x]$ to $[\pi(x), gx]$. If the union of these two lines was a nondegenerate tripod, then g would fix the midpoint of this tripod (as g permutes the endpoints). However, as $\pi(x)$ is the closest fixed point to x, the tripod is degenerate and these two arcs form a geodesic. Hence $d(x, gx) = 2d(x, C_g)$.

Now assume no point in T is fixed by g. Let x be arbitrary and consider the arcs $I = [x, gx]$ and $J = I \setminus (gI \cup g^{-1}I)$. The arc $g^{-1}I$ is an initial segment of I, and the arc gI is a terminal segment of I, and one can show that if gI and $g^{-1}I$ overlap then g fixes the midpoint of I, which contradicts are assumption that no point of T is fixed by g. Furthermore, J and gJ are nondegenerate and intersect only at the terminal endpoint of J. It follows that $J \cup gJ$ is a geodesic by the strong local-to-global property, and by induction the union

$$\cdots g^{-1}J \cup J \cup gJ \cup g^2 \cdots$$

is a line L, and g acts on this line by translation of length $|J|$. Now let π be the projection of T onto the line L. Note that $[x, \pi(x)]$ only intersects L in $\pi(x)$, so as L is g-invariant, the line $g[x, \pi(x)]$ intersects L only at $g\pi(x)$, from which we see that $g\pi(x) = \pi(gx)$, and the lines $[x, \pi(x)]$, $[\pi(x), g\pi(x)]$, and $[g\pi(x), gx]$ only intersect at $\pi(x)$ and $g\pi(x)$, so form a geodesic. As the length of the middle segment is J, it follows that $d(x, gx) = 2d(x, L) + |J|$, and therefore $L = C_g$. □

We have show that an isometry that is not elliptic acts by translation along a copy of $\mathbb{R} \subset T$. Hence:

Corollary 2.2.2. *Every finite order isometry of an $\mathbb{R}$-tree is elliptic.*

Directions and isometries

If $d \subset d'$ or $d' \subset d$ we say the directions d and d' are *coherent*. Otherwise, we say they are *incoherent*.

Lemma 2.2.3 (Coherence lemma)**.** *If d and gd are coherent directions based at x and gx respectively, then $[x, gx]$ is contained in C_g.*

Proof. If $x \notin C_g$, and d is a direction based at x, then either d is facing away from C_g (i.e. $C_g \cap d = \emptyset$), or d is facing towards C_g (i.e. d contains C_g). In the former case, d and d' are in different components of $T - C_g$ (as they are based at points in different components). In the latter case d contains gx but gd does not, and similarly gd contains the point x but d does not. This shows that if d is any direction based at a point not in C_g then d and gd are incoherent. □

The bridge lemma

If A and B are disjoint closed subtrees of a tree T, then every point in B has the same closest point $\pi_A(B)$ in A. Similarly, every point in A has the same closest point $\pi_B(A)$ in B. The *bridge* from A to B is the arc $[\pi_A(B), \pi_B(A)]$.

Lemma 2.2.4 (Bridge lemma)**.** *Suppose g and h are isometries of an $\mathbb{R}$-tree T with disjoint characteristic sets. Then gh is a hyperbolic isometry of translation length*

$$\|gh\|_T = \|g\|_T + \|h\|_T + 2d(C_g, C_h).$$

Furthermore, C_{gh} contains the bridge from C_g to C_h.

Skech proof. Let $[x, y]$ be the bridge from C_h to C_g, and let d be the direction at $h^{-1}x$ containing C_g. The direction $h(d)$ is based at x and does not contain C_g (why?), so $gh(d)$ does not contain C_g by g-invariance. As d contains C_g it contains every component of $T - C_g$ not containing $h^{-1}(x)$, so it contains $gh(d)$. Hence these directions are coherent and $[h^{-1}(x), g(x)]$ belongs to C_{gh}. One can check that

$$[h^{-1}(x), g(x)] = [h^{-1}(x), x] \cup [x, y] \cup [y, g(y)] \cup [g(y), g(x)],$$

which as $x \in C_h$ and $y \in C_g$ implies that $d(h^{-1}(x), g(x)) = \|g\|_T + \|h\|_T + 2d(C_g, C_h)$ (some of the full arguments above might require breaking up into cases where g, h are elliptic/hyperbolic, but everything stated above is true regardless). □

The bridge lemma is useful for constructing hyperbolic isometries with particular axes.

2.3 The boundary at infinity

The boundary of an $\mathbb{R}$-tree, which in terms of hyperbolic spaces is the *Gromov boundary* of the tree, can be defined in terms of equivalence classes of rays, where a ray is an isometrically embedded copy $r\colon [0, \infty) \to T$ of a half line. We say that two rays r and r' are equivalent if their intersection $r \cap r'$ is also a ray. This is an equivalence relation.

Definition 2.3.1 (The boundary of a tree)**.** The boundary ∂T of a tree is defined to be the set of equivalence classes of rays in T. We equip ∂T with the *topology induced by directions*, where we insist that the set of rays lying in a direction d is open in ∂T.

Exercise 2.3.2 (For those interested in δ-hyperbolic geometry)**.** Show that the boundary ∂T as above is homeomorphic to the Gromov boundary of the tree.

Defined in this way, the visual topology extends to a topology on $T \cup \partial T$, and if we take $\overline{T}$ to be the metric completion of T, extends to a topology on $\overline{T} \cup \partial T$.

Exercise 2.3.3. Show that if $x \in T$ then any open set in the visual topology containing x contains all but a finite number of directions at x. Assuming $X = \overline{T} \cup \partial T$ is separable, use this to show that X is compact with the visual topology (separability lets you assume that the open cover you take to prove compactness is countable).

Exercise 2.3.4. (If you like set theory/point set topology) Show that $X = \overline{T} \cup \partial T$ is compact with the visual topology for any $\mathbb{R}$-tree T.

Exercise 2.3.5. Suppose that T is a locally finite simplicial tree. Show that $T \cup \partial T$ is metrizable with the visual topology.

2.4 Classification of group actions

Minimal subtrees

In general, when studying group actions we would like the space to be as small as possible while still carrying all the necessary information about said action. This is covered by the notion of *minimality.*

Definition 2.4.1. An action of a group G on a tree T is *minimal* if there exists no proper, G-invariant subtree of T.

Proposition 2.4.2 (The minimal subtree for an action with a hyperbolic element)**.** *Suppose that the action of G on T contains a hyperbolic isometry. Then there is a unique nonempty G-invariant subtree $S \subset T$, which is the union of the axes of all hyperbolic elements.*

Proof. The proof of classification of isometries shows that if $x \in T$ and g is a hyperbolic element, then the convex hull of the orbit of x under $\langle g \rangle$ contains C_g. This implies that C_g is contained in every g-invariant subtree of T. Hence if $S \subset T$ is G-invariant, it contains C_g for every hyperbolic element G. It remains to show that the union of axes of hyperbolic elements is a G-invariant subtree of T. The G-invariance follows from the fact that if g is hyperbolic and $h \in G$ is any other element then hgh^{-1} is hyperbolic (with the same translation length as g) and $h(C_g) = C_{hgh^{-1}}$ (I would recommend proving this if you haven't seen this before - it applies to all characteristic sets for arbitrary group actions on metric spaces). We are left to prove that this union S of hyperbolic axes is a subtree (i.e. it is connected). Assume for a contradiction that g and h are hyperbolic elements such that C_g and C_h are in different components of S. In particular, the axes of g and h are disjoint. By the Bridge Lemma, the element gh is hyperbolic and C_{gh} contains the bridge from C_g to C_h in T. As $C_{gh} \subset S$, this contradicts C_g and C_h being in different components of S

□

It is crucial that a hyperbolic isometry exists in order for unique minimal subtrees to exist: for instance, if G is acting trivially on a tree T, then every point in T is a proper, nonempty, G-invariant subtree, so that minimal subtrees are far from unique. However, at least when G is finitely generated, this is the only thing that can happen:

Proposition 2.4.3. *If G is a finitely generated group and G acts on a tree T by elliptic isometries, then G has a global fixed point in T.*

Proof. Let X be a finite generating set of G. If $X = \{g_1, \dots, g_n\}$ then by induction we can assume that the subgroup $G_{n-1} = \langle g_1, \dots, g_{n-1} \rangle$ has a global fixed point in T. Let T_{n-1} be the subtree fixed by G_{n-1}. If $C_{g_n} \cap T_{n-1} = \emptyset$, then there exists $h \in G$ that does not fix any point in C_{g_n}, so that $C_h \cap C_{g_{n-1}} = \emptyset$. However in this case hg_n would be hyperbolic by the Bridge Lemma, which is a contradiction. As $T_{n-1} \cap C_{g_n}$ is nonempty, there exists a point fixed by all generators of G, hence fixed by the whole of G. □

More generally, we have the following result:

Exercise 2.4.4. If G acts on a tree T by elliptic isometries, then G fixes a point in $X = \overline{T} \cup \partial T$. HINT: Show that the set $\{X_g\}_{g \in G}$ of fixed point sets in X satisfies the finite intersection property, and apply compactness of X (proved in Exercise 2.3.4).

When T is simplicial the above exercise requires less topological acrobatics: see exercise 2 on page 66 of Serre.

The classification theorem for actions

Definition 2.4.5 (Irreducible action)**.** An action of a group G on a tree T is *irreducible* if there exist hyperbolic elements $g, h \in G$ with disjoint axes.

Theorem 2.4.6. *Let G be a group acting on a tree T with at least one loxodromic element. Then either*

- *G has a fixed point in ∂T*
- *There exists a line $\mathbb{R} \subset T$ that is invariant under the action.*
- *The action is irreducible.*

Sketch proof. If all the hyperbolic isometries share either two common ends or one common end then we get put in cases two and one respectively. Therefore we may assume there exist hyperbolic isometries g, h whose axes have distinct ends. If they are disjoint, we are happy. Otherwise, $C_g \cap C_h$ is a compact arc. Note that $g^k h g^{-k}$ is hyperbolic, and has axis $g^k(C_h)$. Choose a large enough power of g so that the arcs $C_g \cap C_h$ and $C_g \cap C_{g^k h g^{-k}}$ are disjoint. Then $C_{g^k h g^{-k}} \cap C_h$ is empty (as otherwise this would contradict Helly's theorem). □

Exercise 2.4.7 (Ping-pong lemma)**.** Suppose G acts on a space X and $g, h \in G$. Suppose there exist disjoint, nonempty subsets $U, V \subset X$ such that $g^k(U) \subset V$ and $h^k(V) \subset U$ for all $k \neq 0$. Show that $\rangle g, h \langle$ is a free group of rank 2.

Exercise 2.4.8 (Irreducible actions have free subgroups)**.** Show that if G acts irreducibly on a tree T then G contains a free nonabelian subgroup. HINT: Take $X = T$, take a point p on the bridge between C_g and C_h. Let U be the direction at p containing C_h and let V be the direction at p containing C_g.

Lecture 3

A whistle-stop introduction to Bass–Serre theory

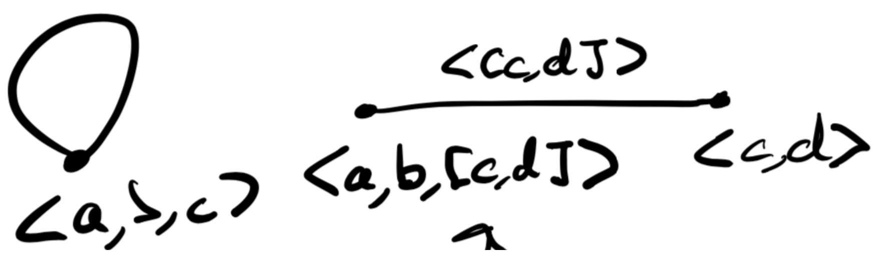

Figure 3.1: Two actions of the free group $F_4 = F(a, b, c, d)$ on trees, written as graphs of a groups.

The goal of this section is to give an overview of Bass–Serre theory, focusing on notation, definitions, and examples, rather than proofs. There are two excellent references for proofs: Serre [6] is a complete classic, and Scott–Wall [5] gives the topological point of view. The notation here is mostly taken from a paper of Cohen–Lustig [1]. In an attempt to compensate for the complete lack of proofs,

In short, Bass–Serre theory tells us how to translate between objects called *graphs of groups* on the one hand, and *group actions on simplicial trees* on the other. While I much prefer working with the geometry that is afforded to us by working with group actions, graphs of groups have the advantage of having very compact notation, and allow us to describe these infinite objects and group actions with a small picture (see Figure 3.1).

These methods only work for simplicial actions - there is not a clean general theory for $\mathbb{R}$-trees in general, but there are some results we shall look into a bit later on.

3.1 A first example (Figures 3.2 and 3.3)

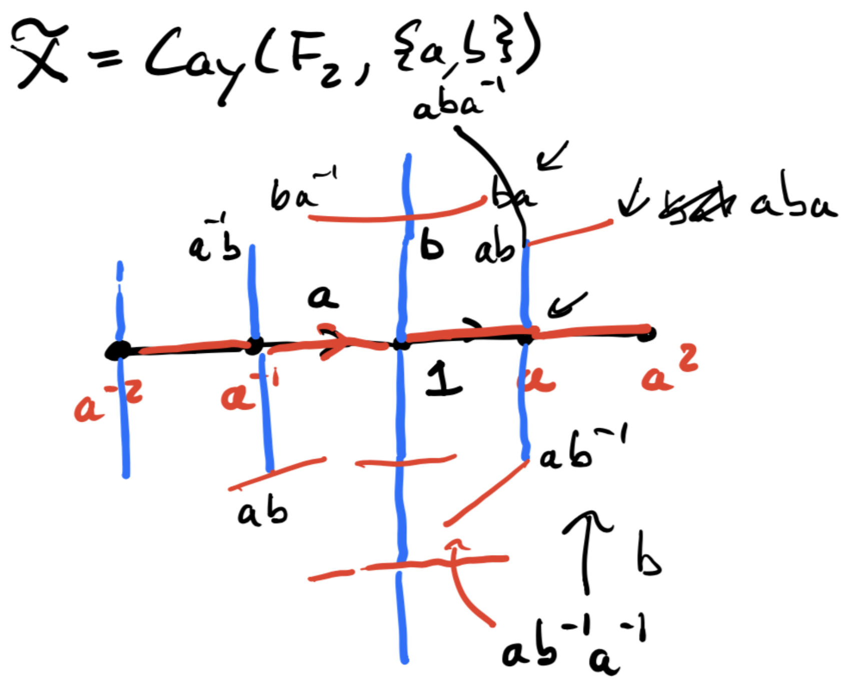

Figure 3.2: A first example, with some of the vertices labelled. Thank you to those in the lecture for helping me get them the right way round.

Let T be the Cayley graph of a free group of rank 2 with respect to a generating set given by two elements called a and b. Every vertex is labelled by a group element, and each edge is labelled by a basis element. Recall

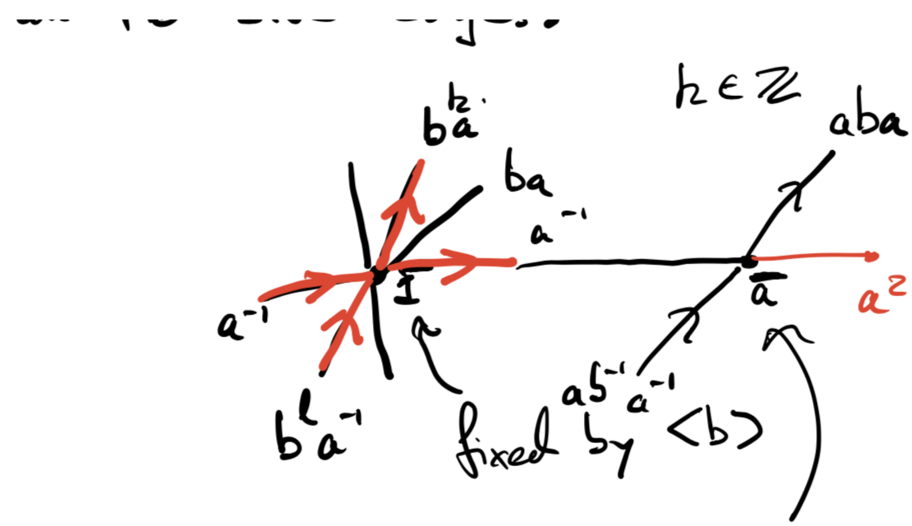

Figure 3.3: The tree after collapsing the blue edges. The red edges still have the same trivial stabilizers as before, however the vertices now have stabilizers conjugate to $\langle b \rangle$ (the group stabilizing the line).

that if s is a basis element then the elements g and gs are linked by an edge labelled s. In this simple example we have a tree with blue and red edges, and F_2 acts freely on this tree.

What happens if you collapse the blue edges? (Figure 3.3)

This example is really worth playing with if you haven't thought about it before (or in fact, even if you have). The blue lines give an F_2-equivariant forest in the Cayley graph. The stabilizer of each of these lines is a conjugate of $\langle b \rangle$, so collapsing then gives a new action of F_2 on a tree. Now the vertices have infinite valence and infinite stabilizers. For instance, if $\overline{1}$ is the image of the identity element under the collapse map, there is an infinite set of red edges going out to vertices of the form $\overline{b^k a}$, and infinitely many red edges coming in from vertices labeled by $\overline{b^l a^{-1}}$. The stablizer of $\overline{1}$ is $\langle b \rangle$, and this stabilizer acts locally by permuting the outgoing/incoming edges at $\overline{1}$.

We shall see that this tree with this very interesting action has a simple description as a graph of groups with a single loop with a trivial edge group and the vertex group labeled by $\mathbb{Z}$.

3.2 Graphs of groups

A *graph of groups* $\mathcal{G}$ is a tuple

$$\mathcal{G} = (\Gamma, (G_v)_{v \in V(\Gamma)}, (G_e)_{e \in E(\Gamma)}, (f_e)_{e \in E(\Gamma)})$$

such that:

- Γ is a connected graph in the sense of Serre (cf. I §2.1 in [6]) with vertex set $V(\Gamma)$ and edge set $E(\Gamma)$.
- Each G_e, G_v is a group.
- If $\tau(e)$ is the terminal vertex of an edge e, we have an injective *edge homomorphism* $f_e \colon G_e \to G_{\tau(e)}$.
- For any edge e, we have $G_e = G_{\bar{e}}$, where $\bar{e}$ denotes the edge e with reversed orientation.

We let $\iota(e) = \tau(\bar{e})$ denote the initial vertex of an edge e.

The path group

The *path group* of $\mathcal{G}$, denoted $\Pi(\mathcal{G})$, is defined by taking the free group F generated by the letters $(t_e)_{e \in E(\Gamma)}$ and quotienting out the free product $(*_{v \in V(\Gamma)} G_v) * F$ by the relations:

- $t_e = t_{\bar{e}}^{-1}$ for all $e \in E(\Gamma)$,
- $t_e f_e(a) t_e^{-1} = f_{\bar{e}}(a)$ for all $e \in E(\Gamma)$ and $a \in G_e$.

In our simple example of a loop with trivial edge group and vertex group $G_v = \mathbb{Z}$, the path group is isomorphic to F_2. However this is a feature of the associated graph being a single loop. More general, we either have to pass to a subgroup or a quotient group of this path group.

The fundamental group of the graph of groups

As a subgroup of the path group

We say that an element $g \in \Pi(\mathcal{G})$ is *connected* if there exists a (possibly trivial) path $e_1, \ldots, e_k$ in Γ starting from a vertex v_0 and elements $g_0, g_1, \ldots, g_k$ such that $g_0 \in G_{v_0}$, $g_i \in G_{\tau(e_i)}$ for each $i \geq 1$ and:

$$g = g_0 t_{e_1} g_1 t_{e_2} \cdots g_{k-1} t_{e_k} g_k.$$

We define $\pi_1(\mathcal{G}, v, w)$ to be the set of elements of $\Pi(\mathcal{G})$ represented by connected words whose underlying paths start at v and end at w. If $v = w$, the set forms a subgroup of $\Pi(\mathcal{G})$ – the *fundamental group of the graph of groups* – and is denoted $\pi_1(\mathcal{G}, v)$.

Given any element x of a group G, let ad_x be the inner automorphism given by the map $g \mapsto xgx^{-1}$. If $W \in \pi_1(\mathcal{G}, v, w)$ then the restriction of $\mathrm{ad}_W \colon \Pi(\mathcal{G}) \to \Pi(\mathcal{G})$ to $\pi_1(\mathcal{G}, w)$ induces an isomorphism between $\pi_1(\mathcal{G}, w)$ and $\pi_1(\mathcal{G}, v)$.

As a quotient of the path group

Alternatively, given a maximal tree S of the graph Γ, the fundamental group of the graph of groups can be defined as a quotient of the path group by

$$\pi_1(\mathcal{G}, S) = \Pi(\mathcal{G}) / \langle\langle t_e : e \in S \rangle\rangle.$$

For any tree T and any $v \in \Gamma$, the composition

$$\pi_1(\mathcal{G}, v) \hookrightarrow \Pi(\mathcal{G}) \twoheadrightarrow \pi_1(\mathcal{G}, S)$$

is an isomorphism.

Useful results

- The vertex groups G_v embed into $\pi(\mathcal{G}, S)$ (this implies G_w is a subgroup of $\pi(\mathcal{G}, v)$ but we have to think a bit more carefully about what this embedding looks like). The edge groups also embed via the maps f_e.
- This is a special case of the fact that connected words in the path group have a *normal form theorem* (see **??**). I won't write this out in full, but in particular it says that if two connected words are equal in $\Pi(\mathcal{G})$ they have the same underlying edge paths. This gives a surjective homomorphism $\pi_1(\mathcal{G}, v) \twoheadrightarrow \pi_1(\Gamma, v)$.

3.3 Building the Bass–Serre tree

We have defined the fundamental group of a graph of groups. We now need to give this group a tree $T_{\mathcal{G}}$ to act on. This will have the following properties:

- Edge stabilizers in $T_{\mathcal{G}}$ are conjugates of the edge groups G_e in $\pi_1(\mathcal{G}, S)$.
- Vertex stabilizers in $T_{\mathcal{G}}$ are conjugates of the vertex groups G_v in $\pi_1(\mathcal{G}, S)$.
- The adjacent edges at a vertex are in a one-to-one correspondence with the cosets of the edge groups in the vertex group (really this is just a restatement of the orbit-stabilizer theorem, but it's a useful fact to know).

The definition

We need to define vertices, edges, and how the edges are glued. As above, we let $\mathcal{G}$ be the graph of groups and Γ the underlying graph. Let $\mathcal{O}$ be an orientation of Γ, and let $S \subset \Gamma$ be a maximal tree in Γ.

- $V(T_{\mathcal{G}})$ is the set of cosets gG_v of the vertex groups $G_v \subset \pi_1(\mathcal{G}, S)$.
- $E(T_{\mathcal{G}})$ is the set of cosets gG_e of the edge groups $G_e \subset \pi_1(\mathcal{G}, S)$.
- Edge inversion is given by $\overline{gG_e} = gG_{\overline{e}}$

To define the ends of edges, if $e \in \mathcal{O}$ we set:

$$\iota(gG_e) = gG_{\iota(e)}$$
$$\tau(gG_e) = gt_eG_{\tau(e)}$$

otherwise if $e \notin \mathcal{O}$ we set

$$\iota(gG_e) = gt_e^{-1}G_{\iota(e)}$$
$$\tau(gG_e) = gG_{\tau(e)}$$

3.4 Example: Baumslag–Solitar groups

3.5 The topological picture

Lecture 4

How Bass–Serre theory fails for general $\mathbb{R}$–trees

4.1 Rips' theorem

Lecture 5

Deformation spaces I: Outer space and the main definition

5.1 Outer space served three ways

As a space of graphs

As a space of sphere systems

As a space of F_n-trees

Lecture 6

Deformation spaces II

We continue to work through Guirardel and Levitt's paper on deformation spaces [3].

6.1 WTF_n is going on????

It's 6 weeks in. We've had 10 hours of stuff and a busy term. Let's take a time-out for a recap.

- Groups (sometimes) act on trees.
- Two actions are *the same* if they are equivalent up to G equivariant isomorphism.
- If G acts on a tree T, then $\mathrm{Aut}(G)$ acts on the space $\mathcal{T}$ of G-actions on simplicial trees *on the right* (at least this most common in the literature) by precomposition. In other words, an action is given by a homomorphism $f\colon G \to \mathrm{Isom}(T)$, and if $\phi\colon G \to G$ is an automorphism we get a new action $f \circ \phi$.
- Understanding the topology, combinatorics, and geometry of $\mathcal{T}$ (or nice subspaces) and this action gives loads of information about the automorphism groups (this has been downplayed so far as we've been more focused on individual actions).
- If ad_g is the inner automorphism by $g \in G$ then $f \circ \mathrm{ad}_g$ is the same action as f (like really, I know I'm not pushy about exercises, but this is a good short one to do). Hence we get an action of $\mathrm{Out}(G)$ on $\mathcal{T}$

- Elements of a group are either elliptic or hyperbolic for an action $G \curvearrowright T$.
- If G is finitely generated and all elements are elliptic then G has a global fixed point (the action is *trivial*).
- If there is some hyperbolic isometry then action $G \curvearrowright T$ has a unique minimal subtree, denoted $T|_G$, which is the union of the hyperbolic axes.
- This (despite there being quite a few nontrivial things here) is often the starting point in the literature - a *G-tree* is often defined as a *minimal* action of G on an $\mathbb{R}$-tree.
- Actions are either *trivial*, *abelian*, *dihedral*, or *irreducible*.
- We mostly care about the *irreducible actions* (that have disjoint hyperbolic axes, contain free groups of hyperbolic elements, etc). These are particularly nice as, due to a theorem of Culler–Morgan [2], they are completely determined by their length functions.
- The *deformation space* $\mathcal{D}(T)$ determined by a tree T is the set of trees T' with the same elliptic subgroups as T. Alternatively two G-trees determine the same deformation space if there exist G-equivariant morphisms $f : T \to T'$ and $f' : T' :\to T$ (again, v v good exercise).
- All trees in the same deformation space are of the same type.

6.2 Making the space smaller and some simple examples

For this section, our main example will be the deformation space given by the tree T coming from the free product $A * B$ of two groups. We will see that this tree is not unique in $\mathcal{PD}$, but is unique (up to homothety) in a sub-object called the reduced deformation space $\mathcal{D}_r$. The need for this is highlighted as follows: if we take any chain $A = A_0 \supset A_1 \supset A_2 \cdots \supset A_k$ we can build a new tree with $k+1$ edge-orbits coming from the line:

The set of elliptic subgroups of this new splitting is the same as the set of elliptic subgroups in T, so they belong to the same deformation space.

Reduced trees

Definition 6.2.1 (Reduced trees)**.** A tree T is *reduced* if every tree T' obtained by collapsing an orbit of edges in T is in a different deformation space. Equivalently, for every edge $e \in T$, either e is a loop in the quotient graph T/G or the edge stabilizer G_e is a proper subgroup of both of its adjacent vertex groups.

Example 6.2.2. The example $T = A * B$ is reduced, but the tree T' we obtained by stretching out the A vertex is not.

In Outer space, a tree T is reduced if and only if T/F_n is a rose.

Proposition 6.2.3 (Arc stabilizers in reduced trees ([3], Proposition 4.6))**.** *Let T be reduced. A subgroup $H \subset G$ fixes a nondegenerate arc in T if and only if it is contained in a subgroup K of the form*

- $K = A \cap B$ *where* $A, B \in \mathcal{E}(T)$ *but* $\langle A, B \rangle \notin \mathcal{E}(T)$
- $K \in \mathcal{E}(T)$ *and there exists g hyperbolic such that* $K \subset gKg^{-1}$

As these conditions are phrased in terms of elliptic subgroups, the set of subgroups that fix nondegenerate arcs is the same for every reduced tree in a deformation space.

Proof. First note that any such K (and hence its subgroups) fixes an arc in T. In the first case, the conditions imply that A and B fix subtrees S_A and S_B, however, as $\langle A, B \rangle \notin \mathcal{E}(T)$, these subtrees are disjoint. It follows that $A \cap B$ fixes the bridge from S_A to S_B. In the second case, if v is a fixed point of K, then so is gv, so K fixes the arc from v to gv.

Now suppose T is reduced and H fixes a nondegenerate arc in T. Then H is contained in some edge stabilizer $G_e := K$. Let v, w be the endpoints of e. If G_e is a proper subgroup of both G_v and G_w. Then we can take $A = G_v$, and $B = G_w$. Otherwise, without loss of generality we can assume that $G_e = G_v$ and, as T is reduced, e maps to a loop in T/G. Then there exists hyperbolic $g \in G$ such that $gv = w$ (e.g. given by a stable letter for the loop in the graph of groups), and $gG_eg^{-1} = gG_vg^{-1} = G_w$ contains G_e. □

The general mantra is that although vertex and edge stabilizers are not always the same throughout a deformation space (even for reduced trees), *the set of subgroups of edge stabilizers* does not depend on a given reduced tree.

Definition 6.2.4 (Restricted deformation spaces, reduced deformation spaces)**.** If $\mathcal{A}$ is a family of subgroups of G then $\mathcal{D}_{\mathcal{A}}$ is the set of G-trees $T \in \mathcal{D}$ such that each edge stabilizer in T belongs to $\mathcal{A}$. For a deformation space $\mathcal{D}$, we let $\mathcal{A}_{\min}$ be the set of subgroups of edge groups in a reduced tree in $\mathcal{D}$. This is well-defined by Proposition 6.2.3. The *reduced deformation space* is defined to be $\mathcal{D}_r := \mathcal{D}_{\mathcal{A}_{\min}}$.

Remark 6.2.5 (This notation sucks)*.*

- Reduced Outer space in the sense of Culler–Vogtmann is a completely different notion to reduced deformation spaces.
- Not all trees in reduced deformation space are reduced (again look at the Outer space example).
- An irreducible action (in the sense of having independent loxodromics) need not be reduced, and conversely reduced trees need not be irreducible (e.g. the HNN tree for $BS(1,2)$).

Proposition 6.2.6. *Let T be the one-edge splitting given by a free product decomposition $G = A * B$. Then every tree in the reduced deformation space $\mathcal{D}_r(T)$ is equal to T, up to G-equivariant homothety.*

Proof. As all edge stabilizers in T are trivial, the reduced deformation space is the set of trees T' with the same elliptic subgroups as T and trivial edge stabilizers. Let v and w be the vertices of T fixed by A and B respectively, and let e be the edge between them. There exist (unique) vertices v' and w' in T' fixed by A and B respectively. Define $f(v) = v'$, $f(w) = w'$ and extend to a G-equivariant map $f : T \to T'$ (why can you do this?), mapping e to the unique reduced edge path from v' to w'. The map f is surjective, as the image is a G-invariant subtree of T', which must be T' itself by minimality. The map is also locally injective: suppose that d and d' are two directions at a vertex in T that are mapped to the same direction in T'. Without loss of generality, let us assume that this vertex is the one fixed by A (we called this v above). Note that A acts transitively on the set of directions at v, so that $ad = d'$ for some $a \in A$. This implies that a fixes some initial segment of the image of d in T', which is a contradiction as T' has trivial edge stabilizers. Therefore f is locally injective, but locally injective maps of trees are injective by the local-to-global property of geodesics. Hence f is a simplicial isomorphism (or a homothety in the metric setting). □

What happens with $A * B * C$?

6.3 Topologies on spaces of trees

Let G be a finitely generated group. Let $\mathcal{T}$ be the set of all simplicial G-trees. We previously defined three different topologies on a simplicial $\mathbb{R}$–tree. We will now define three different topologies on the space of simplicial $\mathbb{R}$–trees. The number three appearing twice here is mostly a coincidence, although both times there is a CW, or 'weak' topology, that happens in fact to be quite strong.

The 'weak', or *open simplicial* topology

Calling this the weak topology is quite misleading, but people do it. This is a 'simplicial' style topology on a deformation space, and as we've seen with trees earlier, when the spaces in question are not locally finite, the simplicial topology ends up being finer than other natural topologies.

Definition 6.3.1 (Open simplicial topology)**.** Given $T \in \mathcal{D}$, let $\mathcal{C}(T)$ be the *open cone* in $\mathcal{T}$ determined by T. This is the set of simplical $\mathbb{R}$-trees obtained by varying the edge lengths in T in $(0, \infty)$. We topologise this as a subspace of $\mathbb{R}^{E(T)/G}$. The *closed cone* $\overline{\mathcal{C}}(T)$ is the set of nontrivial simplicial $\mathbb{R}$-trees obtained by collapses of elements of $\mathcal{C}(T)$ (i.e., the space where we allow all-but-one edge lengths to go to zero). A set $V \subset \mathcal{T}$ is defined to be closed if the intersection of V with every closed cone is closed.

An unfortunate part of this definition is that open cones are not necessarily open in $\mathcal{T}$, although maximal open cones (i.e., coming from unrefinable trees) are open in this topology.

The equivariant Gromov-Hausdorff topology

In the *equivariant Gromov-Hausdorff topology*, or simply *Gromov topology*, we define neighbourhoods by approximating how finite subsets of G act on finite subtrees of T.

Definition 6.3.2 (Equivariant G–H topology)**.** Given $T \in \mathcal{T}$, a neighbourhood $V_T(X, A, \epsilon)$ of T in the Gromov–Hausdorff topology is determined by finite subsets $X \subset T$ and $A \subset G$ and a real number $\epsilon > 0$. A tree T' belongs to $V_T(X, A, \epsilon)$ if and only if there exists a map $f : X \to T'$ such that

$$|d_T(x, gx') - d_{T'}(f(x), gf(x'))| < \epsilon$$

for all $x, x' \in X$ and $g \in A$.

The length function topology

Let $F : \mathcal{T} \to \mathbb{R}^G$ be the map defined by

$$T \mapsto (\|g\|_T)_{g \in G}.$$

The *length function topology*, or *axes topology* is the coarsest topology on $\mathcal{T}$ that makes F continuous with respect to the product topology on $\mathbb{R}^G$. Equivalently, it is the coarsest topology on $\mathcal{T}$ such that every individual length function $T \mapsto \|g\|_T$ is continuous.

For an in-depth discussion of the relationship between these topologies, see Section 5 of [3]. However, with respect to irreducible actions, the picture is a much clearer:

Theorem 6.3.3 ([2])**.** *Let $\mathcal{T}_{irr}$ be the space of irreducible G–trees. The map*

$$F : \mathcal{T}_{irr} \to \mathbb{R}^G$$

given by $T \mapsto (\|g\|_T)_{g \in G}$ is injective.

Theorem 6.3.4 ([4])**.** *Let $\mathcal{T}_{irr}$ be the space of irreducible G–trees. The Gromov–Hausdorff and length function topologies agree on $\mathcal{T}_{irr}$.*

For our purposes, we will only make use of the following:

Proposition 6.3.5. *Let $\mathcal{D}$ be a deformation space. The Gromov topology and the weak topology induce the same topology on any finite union of cones in $\mathcal{D}$ (resp. of simplices of $\mathcal{PD}$)*

Misc. Exercises

These have been transported from some previous notes. As a result, we may not have seen all the definitions required for these, and I plan to re-home them as appropriate.

Exercise 6.3.6. Let T be a tree and $\overline{T}$ be the metric completion of T. Show that for every point $x \in \overline{T} - T$ there is only one direction in $\overline{T}$ at x.

Exercise 6.3.7. Suppose G acts on a tree T. Show that if N is a normal subgroup of G containing a hyperbolic element then the action of G on T is minimal if and only if the action of N on T is minimal.

Exercise 6.3.8. Suppose that T has trivial arc stabilizers, let H be a subgroup of G with trivial arc stabilizers and suppose that $g \in G - H$. Show that the set of nondegenerate trees of the form $Y_h = \{x : gx = hx\}$ forms a transverse family in T.

Exercise 6.3.9. Show that if $\mathcal{Y}$ is a transverse family in a tree T with dense orbits and finitely many orbits of branch directions, then $\mathrm{Stab}(Y)$ acts on Y with dense orbits.

Exercise 6.3.10. Show that if the action of G on T is minimal and irreducible with dense orbits then the branch points are dense in every arc of T.

Exercise 6.3.11. Show that if g and h fix a point in T and g and h commute, then the product gh also fixes a point.

Exercise 6.3.12. Suppose that G is a finitely generated group with a nontrivial action on a tree T with dense branch points. Show that the minimal subtree T_{min} of G is nowhere dense in the metric completion $\overline{T}$. Hint: Show any arc is nowhere dense in T_{min} and apply the Baire Category Theorem.

Exercise 6.3.13. Suppose that $G \curvearrowright T$ and g, h are (not necessarily hyperbolic) elements with $C_g \cap C_h = \emptyset$ and $C_{h^n} = C_h$ and $C_{g^n} = C_g$ for all non-identity powers of g and h, then $\langle g, h \rangle \cong \langle g \rangle * \langle h \rangle$

Exercise 6.3.14. Show that for any action $G \curvearrowright T$, if $g \in G$ and $x \in T$ then

$$\|g\|_T = \max\{d(x, g^2x) - d(x, gx), 0\}.$$

Use this to show that every open subset of X_{ne} with the axis topology is also open with respect to the Gromov-Hausdorff topology.

Exercise 6.3.15. Show that if $G \curvearrowright T$ is irreducible and $B(T, F, S, \epsilon)$ is a basic neightbourhood of T in the Gromov–Hausdorff topology then there exists a set F' of branch points in T and $\epsilon' > 0$ such that $B(T, F', S, \epsilon') \subset B(T, F, S, \epsilon)$.

Bibliography

[1] M. M. Cohen and M. Lustig. The conjugacy problem for Dehn twist automorphisms of free groups. *Comment. Math. Helv.*, 74(2):179–200, 1999.

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

[3] V. Guirardel and G. Levitt. Deformation spaces of trees. *Groups Geom. Dyn.*, 1(2):135–181, 2007.

[4] F. Paulin. The Gromov topology on R-trees. *Topology Appl.*, 32(3):197–221, 1989.

[5] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.

[6] J.-P. Serre. *Trees.* Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003. Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.