Outer Space

Contents

Introduction **1**

1 G-trees **3**

1.1 Isometries of trees . . . 3

1.2 G-trees . . . 3

1.3 Spaces of trees. Cones . . . 5

1.4 Topologies in the spaces of trees . . . 5

1.5 Projectivized spaces of trees . . . 6

2 Deformation spaces **8**

2.1 Domination of trees . . . 8

2.2 Deformation spaces . . . 9

2.3 Elementary deformations . . . 10

2.4 Topology of deformation spaces . . . 11

3 Outer Space **12**

3.1 Marked graphs and $\mathbb{F}_n$-trees . . . 12

3.2 Outer Space as a projectivized deformation space . . . 13

3.3 The deformation retract . . . 15

3.4 The case $n = 2$ . . . 16

4 The proof of contractibility **20**

4.1 Outline of the proof . . . 20

4.2 Skora's Deformation . . . 22

4.3 Choice of basepoint P_T . . . 23

4.4 Choice of f_T. Contractibility in the Gromov topology . . . 24

4.5 Contractibility in the weak topology . . . 26

5 Consequences for Out($\mathbb{F}_n$) **27**

References **30**

Introduction

In the highly cited paper [CV86], Culler and Vogtmann introduced Outer Space: a topological space whose elements are trees and which has a nice action by $\mathrm{Out}(\mathbb{F}_n)$, the outer automorphism group of the free group with n generators. They prove that Outer Space is contractible, and deduce some algebraic results about $\mathrm{Out}(\mathbb{F}_n)$ as a consequence. This proof is combinatorial in nature, and based on a detailed study of the action of $\mathrm{Out}(\mathbb{F}_n)$ on a deformation retract of Outer Space.

This essay aims to obtain the same results as [CV86], but putting the contractibility result in a more general context: Outer Space turns out to be a very simple example of a deformation space, that is, a space of trees with actions by a fixed group and which can be deformed in to each other, in some sense we will detail in section 2. In [GL06], Guirardel and Levitt give very mild conditions for deformation spaces to contractible, and that paper will be the main reference for this essay. The essay should be understandable for any part III student, with one possible exception: especially in section 4, basic knowledge about isometries of trees is assumed. The paper [CM87] is a great introduction to this topic, and sections $1-3$ of it are more than enough for the purposes of this essay.

Section 1: G-trees contains the information about G-trees (trees with an action of a given group G by isometries) necessary to introduce deformation spaces effortlessly. Almost all the results of this section are contained in sections 2 and 5 of [GL06]. Only simplicial trees will be useful for the purposes of the essay, but most of the results will be expressed for any $\mathbb{R}$-tree, as it is good to know which results and definitions make sense for any $\mathbb{R}$-tree and which ones require a simplicial structure.

Subsections 1.1 and 1.2 cover the notation we will use to talk about G-trees, similar to the one in [GL06]. Then 1.3 explains how the space of minimal metric simplicial G-trees can be divided into open cones obtained from changing the length of edges of trees. In 1.4, we introduce the three topologies we will usually consider in spaces of trees: the axes, Gromov and weak topologies, and some relations between them. Finally, 1.5 defines the projectivized space of trees $\mathrm{P}\mathcal{T}$ and the relation between its topology and the topology of $\mathcal{T}$.

Section 2: Deformation spaces explains everything we will need to know about deformation spaces for the purposes of this essay. Its results are mostly contained in sections 3 and 5 of [GL06].

Subsection 2.1 introduces the equivalence relation of domination between trees, and why it is interesting. This leads to the definition of deformation spaces in subsection 2.2, where we also explain what deformation spaces have to do with the different types of G-trees defined in 1.2, concluding that only genuine abelian or irreducible deformation spaces can be interesting. In 2.3 we talk about contractions, expansions and elementary deformations, which are different ways to continuously transform trees into other trees. The last subsection mentions some necessary properties of the topology of deformation spaces.

Section 3: Outer Space gives an introduction to Outer Space similar to the one in [CV86], aided by the tools from sections 1 and 2.

We begin in 3.1 defining marked graphs and showing that they essentially are the deformation space $\mathcal{D}_n$ of $\mathbb{F}_n$ with no non-trivial elliptic subgroups. Then in 3.2 we obtain some properties of $\mathcal{D}_n$ as a deformation space, we define Outer Space X_n as the space of marked graphs of volume 1, and we prove that it is essentially the projectivized version of $\mathcal{D}_n$. In subsection 3.3 we go on to study K_n, a deformation retract of X_n which is a simplicial complex and also has an action of $\mathrm{Out}(\mathbb{F}_n)$: this, along with the fact that K_n is contractible, will be useful for us in section 5. Finally, subsection 3.4 studies in detail the nice geometric structure of Y_2, a deformation retract of X_2 formed by its marked graphs without separating edges.

Section 4: The proof of contractibility has two main results: contractibility of any deformation space in the weak topology and contractibility of any deformation space with finitely generated vertex stabilizers in the Gromov topology. This section is based on section 6 of [GL06] and section 3 of [GL07].

Firstly, 4.1 gives a brief outline of the proof of contractibility of a deformation space $\mathcal{D}$: it explains how we will define a contraction of $\mathcal{D}$ as a composition of three functions, and proves that one of the three (the easiest) is continuous in the Gromov topology. Then 4.2 defines in detail the second function, Skora's Deformation, and 4.3 and 4.4 define and prove in detail the continuity of the third function in the Gromov topology, concluding the proof of contractibility in the Gromov topology. The lemmas involved are good examples of how to prove continuity of maps between spaces of trees, and how 'stable' systems of equations can be used to do it. Finally 4.5 explains how contractibility in the weak topology is an easy corollary of contractibility in the Gromov topology, modulo a technical result proved in [GL06].

Section 5: Consequences for Out$(\mathbb{F}_n)$ indicates how the contractible cell complexes on whose cells a group G acts freely are related to the cohomological dimension of G and whether it has the FL property. It also explains why, as $\mathrm{Out}(\mathbb{F}_n)$ has torsion, it has neither of these properties, so we have to consider a finite index torsion free subgroup of $\mathrm{Out}(\mathbb{F}_n)$. The section concludes proving that $\mathrm{vcd}(\mathrm{Out}(\mathbb{F}_n)) = 2n-3$ and $\mathrm{Out}(\mathbb{F}_n)$ is of type WFL. The main reference for this section is [Br82].

1 G-trees

This chapter explains the basic prerequisites needed to understand deformation spaces: the space of G-trees for a given group G.

1.1 Isometries of trees

In this essay we will mainly study **simplicial trees**, that is, 1-dimensional connected cell complexes which contain no loops. We will also consider **metric simplicial trees**, which consist on simplicial trees with a metric that makes them geodesic spaces. Notice that up to isometry, a metric simplicial tree only depends on the simplicial structure of the tree and the lengths of the edges.

The following is an important generalization of simplicial trees:

Definition 1.1. We say a metric space T is a $\mathbb{R}$**-tree** if for every two points $a, b \in T$ there is a unique arc joining a to b, and that arc is isometric to $[0, d(a,b)]$.

The concept of $\mathbb{R}$-tree is a very natural one, as any metric space is an $\mathbb{R}$-tree iff it is 0-hyperbolic. So $\mathbb{R}$-trees are more 'complete' than simplicial trees in the sense that any limit of $\mathbb{R}$-trees in the Gromov topology (which will be defined in section 1.4) is another $\mathbb{R}$-tree. Moreover, a lot of the properties of simplicial trees are true for all $\mathbb{R}$-trees. The basics of isometries of $\mathbb{R}$-trees are covered in detail in [CM87], here we will just clarify basic notation.

Let T be an $\mathbb{R}$-tree and let $f : T \to T$ be an isometry. We define $l(f)$ (or $l_f(T)$), the **length** of f, to be $\min\{d(x, f(x)); x \in T\}$. This minimum is always reached, and we define $C_f = \{x \in T; d(x, f(x)) = l(f)\}$. We can distinguish two possibilities for f depending on $l(f)$:

- **Elliptic**: $l(f) = 0$, that is, f fixes some point. In this case C_f, the fixed point set of f, is a subtree, and for every $x \in T$, $d(x, f(x)) = 2d(x, C_f)$.
- **Hyperbolic**: $l(f) > 0$. In this case C_f will be a line, which we call the **axis** of f, and f will act on C_f by a translation of length $l(f)$. For every $x \in T$, we will have $d(x, f(x)) = l(f) + 2d(x, C_f)$.

1.2 G-trees

During the whole document, except chapter 3, G will denote a finitely generated group. Given a simplicial tree T and a simplicial automorphism $f : T \to T$, we will say it is an **inversion** if $f(u) = v, f(v) = u$ for some edge (u, v).

Definition 1.2. We will say an $\mathbb{R}$-tree is a G**-tree** if it has an action of G by isometries. We will say a (metric) simplicial tree T is a (metric) simplicial G**-tree** if G has a simplicial action on T (by isometries) without inversions. We say T is **minimal** if it has no proper invariant subtree.

As the notation suggests, instead of fixing a tree T and considering the possible actions of groups on T, we are fixing G and considering the possible trees on which G acts.

Notation 1.3. Given points $a, b \in T$, we will denote by $[a, b]$ the segment from a to b. A **ray** of T is the image of an isometric imbedding of $[0, \infty)$ into T. An **end** of T is an equivalence class of rays, with 2 rays being equivalent if their intersection is a subray.

Proposition 1.4. *A minimal simplicial G-tree T has no nodes of valence 1 and has finitely many orbits of edges.*

Proof. If T had a vertex v of valence 1, then we could remove all the orbit Gv from T and the adjacent edges and we would be left with a proper invariant subtree of T. To prove the second part consider a finite set of generators of G, $\{g_1, \ldots, g_n\}$, and a point $p \in T$. As T is minimal, $T = \cup_{g,h\in G}[gp, hp] = \cup_{g\in G}[p, gp]$. So any $x \in T$ is contained in a segment $[p, hp]$ for some $h \in G$. If we express $h = g_{i_1} \ldots g_{i_N}$ and $h_j = g_{i_i} \ldots g_{i_j}$, then $x \in [h_{k-1}p, h_k p]$, for some k. Thus $h_{k-1}^{-1}p \in [p, g_{i_k}p]$. This means that T is contained in GX, where $X = \bigcup_{i=1}^{n}[p, g_i p]$. As X is a finite subtree, G contains only finitely many orbits of edges. □

Definition 1.5. Following the notation from [GL06], we can divide minimal G-trees into five types:

- Trivial: T is just 1 point.
- Linear abelian: T is isometric to $\mathbb{R}$ and G acts by translations.
- Dihedral: T is isometric to $\mathbb{R}$ and some elements of G act by symmetries.
- Genuine abelian: The action of T fixes exactly one end of T.
- Irreducible: It has two hyperbolic elements whose axes have compact intersection.

Any G-tree T has an associated length function $G \to \mathbb{R}$ which sends an element g to its length $l_g(T)$ (we will also call it $l_T(g)$ sometimes). In the two abelian types, this length function is the absolute value of a homomorphism from G to $\mathbb{R}$. In the non abelian cases, the length function determines the tree uniquely, as stated in theorem 3.7 of [CM87]:

Theorem 1.6. *If two non abelian G-trees have the same length function, then they are equivariantly isometric. If they are irreducible, then the equivariant isometry is unique.* □

This theorem completely fails in the abelian case, here is one counterexample: the length function of a genuine abelian G-tree[1] is given by $l_T(g) = |h(g)|$, where $h : G \to \mathbb{R}$ is an homomorphism. But G also acts on $\mathbb{R}$ by $g(x) = x + h(g)$, and this new action is linear abelian and has the same length function as T.

[1] Such trees exist, see example 3.9 of [CM87].

1.3 Spaces of trees. Cones

Let $\mathcal{T}_{\mathbb{R}}$ be the set of $\mathbb{R}$-trees with a minimal action of G by isometries, up to equivariant isometries, and let $\mathcal{T}$ be its subset formed by minimal metric simplicial G-trees.

We define an **open cone** of $\mathcal{T}$ to be a maximal subset of $\mathcal{T}$ whose trees are pairwise equivariantly homeomorphic. Open cones clearly form a partition of $\mathcal{T}$, and their name comes from the fact that non trivial open cones can be parametrized by the open cones $(0, \infty)^n$. In the linear abelian and dihedral case, giving the parametrization is easy: an element of $\mathcal{T}$ is equivariantly homeomorphic to a linear abelian/dihedral tree T iff its metric is a scalar multiple of the metric of T, so we can parametrize those cones by $(0, \infty)$.

If a G-tree T is genuine abelian or irreducible, it can be given a unique structure as a graph without nodes of valence 2. Let $e_1, \ldots, e_n$ be the orbits of edges of T. Then to each point $x = (x_1, \ldots, x_n) \in (0, \infty)^n$ we can associate a metric simplicial G-tree T_x with the same simplicial structure as T but with e_i having length x_i. This tree T_x is unique up to equivariant isometry, and $\{T_x; x \in (0, \infty)^n\}$ is the whole open cone. Moreover, if T is irreducible and $x \neq y$ then T_x and T_y cannot be equivariantly isometric: suppose we have an equivariant isometry $\phi : T_x \to T_y$. Then, as $x \neq y$, we will have $i(e_i) = e_j$ for some $i \neq j$. But then, changing the metrics in T_x and T_y so that all edges have length 1, ϕ becomes a non trivial isometry from the irreducible tree $T_{(1,1,\ldots,1)}$ to itself, which contradicts theorem 1.6.

We can extend the domain of this function to the closed cone $[0, \infty)^n$ by collapsing the edges in the orbit of e_i whenever $x_i = 0$ (the graph we obtain will still be a tree): for example, if $x_i = 0$ $\forall i$, we obtain the trivial tree. We denote the set of trees we obtain as $\overline{\mathcal{C}_T}$ and call it a **closed cone**.

1.4 Topologies in the spaces of trees

As in [GL06], we will consider 3 topologies on these spaces.

1. **Axes topology**

   It is the weakest topology on $\mathcal{T}_{\mathbb{R}}$ which makes the length functions $l_g : \mathcal{T}_{\mathbb{R}} \to \mathbb{R}$ continuous $\forall g \in G$.

2. **Gromov topology**

   In this topology, two G-trees T, T' are near each other if finite subsets of G have similar actions on finite subsets of T and T'. Specifically, neighborhoods of a tree T will be sets $V_T(X, A, \varepsilon)$, where $X \subseteq T$ and $A \subseteq G$ are finite and $\varepsilon > 0$, defined as the set of trees T' such that there is a function $f : X \to T'; x \mapsto x'$ such that $|d(x, ay) - d(x', ay')| < \varepsilon$ $\forall x, y \in X, a \in A$.

3. **Weak topology**:

   This one can only be defined in $\mathcal{T}$, it consists of the sets $O \subseteq \mathcal{T}$ such that, for every closed cone $\overline{\mathcal{C}}$, $O \cap \overline{\mathcal{C}}$ is open in $\overline{\mathcal{C}}$.

Proposition 1.7. *The axes topology is weaker than the Gromov topology, which is weaker than the weak topology.*

Proof. For the first part we have to prove that the length functions l_g are continuous respect to the Gromov topology $\forall g \in G$. If g is elliptic, choose a point p fixed by g and then for any $T' \in V_T(\{p\}, \{g\}, \varepsilon)$ we have that $d(p', gp') < d(p, gp) + \varepsilon = \varepsilon$, so $l_g(T') < \varepsilon$. If g is hyperbolic, the axis of g^2 is the same as the axis of g, thus $l_g(T) = d(p, g^2p) - d(p, gp)$. More in general, for any tree T' the inequality $l_g(T') \geq d(p, g^2p) - d(p, gp)$ holds (it is obvious in the elliptic case). Suppose you have $T' \in V_T(\{p\}, \{g, g^2\}, \varepsilon)$. Then $l_g(T') \geq d(p', g^2p') - d(p', gp') \geq d(p, g^2p) - d(p, gp) - 2\varepsilon = l_g(T) - 2\varepsilon$. So if you choose ε small, T' will be hyperbolic, and in that case, $l_g(T') = d(p', g^2p') - d(p', gp') \leq d(p, g^2p) - d(p, gp) + 2\varepsilon$, and we are done.

For the second part, we just have to prove that the functions $[0, \infty)^n \to \mathcal{T}$ associated to the closed cone of any tree T are continuous in the Gromov topology. This is true because for any two points $p, q \in T$, $d_{T_x}(p, q)$ depends linearly on $x \in [0, \infty)^n$. □

Corollary 1.8. *The restriction of the three topologies to non abelian trees is* $T2$.

Proof. It is a consequence of 1.7 and 1.6. □

The following theorem is not easy: it is the main result of [Pa89].

Theorem 1.9. *The Gromov topology and the axes topology have the same restriction to the set of non abelian trees.* □

1.5 Projectivized spaces of trees

Consider the action of $\mathbb{R}$ on $\mathcal{T}_\mathbb{R} \setminus \{0\}$ by homotheties. This action is continuous, and taking the quotient by the action we get the **projectivized space** of G-trees, $P\mathcal{T}_\mathbb{R}$. Similarly, we define $P\mathcal{T}$, the projectivized space of metric simplicial G-trees, as the image of $\mathcal{T} \setminus \{0\}$ under this quotient. We will call the images of open/closed cones of $\mathcal{T}$ in $P\mathcal{T}$ by this map **open/closed simplices**.

As one would expect, $\mathcal{T}_\mathbb{R}$ is homeomorphic to $P\mathcal{T}_\mathbb{R} \times \mathbb{R}$. To give an explicit homeomorphism we will need a lemma about isometries of trees. Take a set $g_1, \ldots, g_n$ of generators of G.

Lemma 1.10. *If T is a minimal G-tree and g_i and g_ig_j are elliptic $\forall i, j$, then T is a trivial tree.*

Proof. It is enough to prove that the fixed point sets $C_{g_1}, \ldots, C_{g_n}$ have non empty intersection. They certainly have non empty pairwise intersections, because if two elliptic elements have disjoint fixed point sets, then their composition is hyperbolic. So it is enough to prove that if a finite set of subtrees has nonempty pairwise intersections, then they have nonempty intersection. In the case $n = 3$, this is easy: given subtrees T_1, T_2, T_3 with points $p_1 \in T_2 \cap T_3, p_2 \in T_1 \cap T_3$ and $p_3 \in T_1 \cap T_2$, the central point of the tripod formed by P_1, P_2, P_3 is in $T_1 \cap T_2 \cap T_3$. If $n > 3$ and we are given subtrees $T_1, \dots, T_n$ with nonempty pairwise intersections, then the subtrees $T_1 \cap T_2, \dots, T_1 \cap T_n$ have non empty pairwise intersections due to the case $n = 3$, so we can conclude by induction. $\square$

Something similar to the following result is mentioned in section 5 of [GL06]:

Proposition 1.11. *The identification of $\mathcal{T}$ to $P\mathcal{T} \times \mathbb{R}$ via $T \mapsto ([T], \sum_{1 \le i < j \le n} l_{g_i g_j}(T) + \sum_i l_{g_i}(T))$ is an homeomorphism in all three topologies. Thus $P\mathcal{T}$ is homeomorphic to the set of trees $T \in \mathcal{T}$ with $\sum_{1 \le i < j \le n} l_{g_i g_j}(T) + \sum_i l_{g_i}(T) = 1$.*

In the Gromov and axes topologies, these results extend to $\mathcal{T}_{\mathbb{R}}$.

Proof. The given function is clearly continuous, as length functions are continuous in all topologies. We can prove the inverse is continuous in the three topologies separately

In the weak topology, notice that the projection $\mathcal{T} \times \mathbb{R} \to P\mathcal{T} \times \mathbb{R}$ is open, so it is a quotient map. So to prove that the inverse is continuous, we just need to prove that the map $\mathcal{T} \times \mathbb{R} \to \mathcal{T}; (T, x) \mapsto T_x$, where T_x is the scalar multiple of T with $\sum_{1 \le i < j \le n} l_{g_i g_j}(T_x) + \sum_i l_{g_i}(T_x) = 1$, is continuous. This map is continuous in every cone, and $\mathcal{T} \times \mathbb{R}$ has the weak topology (see theorem $A.2$ of [Ha00]), so we are done.

In the Gromov topology, we just need to use the fact that $\lambda T'$ will always be in $V_T(X, A, \varepsilon)$ whenever $T' \in V_T(X, A, \frac{\varepsilon}{2})$ and $\lambda \in [1 - \delta, 1 + \delta]$ for small enough δ.

In the axes topology it follows from the length functions $P\mathcal{T}_{\mathbb{R}} \times \mathbb{R} \to \mathbb{R}$ being continuous. $\square$

This proposition essentially means that we can translate topological results about $\mathcal{T}$ ($\mathcal{T}_{\mathbb{R}}$) to $P\mathcal{T}$ ($P\mathcal{T}_{\mathbb{R}}$) and vice-versa. The following is a good example of a result which can be more easily proved in $P\mathcal{T}$ than in $\mathcal{T}$:

Proposition 1.12. *The Gromov and the weak topologies coincide in finite unions of simplices (cones) formed by non abelian trees.*

Proof. It is enough to prove the result for finite unions of closed simplices, as we will prove in 2.7 that if an open simplex is non abelian, its closure contains only non abelian trees. These finite unions are compact in the weak topology and T_2 in the Gromov topology due to 1.8, so as the Gromov topology is weaker, both topologies coincide. $\square$

2 Deformation spaces

We proceed to study a partition of $\mathcal{T}$ into the main objects of study of this essay: deformation spaces. The results of this chapter are contained in sections 3 and 5 of [GL06], although the presentation is different and some proofs may differ.

2.1 Domination of trees

Definition 2.1. Given $T, T' \in \mathcal{T}$, we say T **dominates** T' if there is an equivariant map (continuous function) $f : T \to T'$.

Note that any equivariant map $f : T \to T'$ is surjective, because T' is minimal, hence the name 'dominates'. We say two trees T, T' are **equivalent**, $T \sim T'$, if T dominates T' and T' dominates T. This is clearly an equivalence relation.

To show why equivalence of trees is a useful and deep relation, let's reformulate it in other ways. Remember that a subgroup $H < G$ is elliptic in a G-tree T if H fixes some point of T. The following is a simpler version of theorem 3.8 of [GL06]:

Theorem 2.2. *Given $T, T' \in \mathcal{T}$, TFAE:*

1. *$T \sim T'$.*
2. *T and T' have the same elliptic subgroups.*
3. *There is an equivariant quasi-isometry between T and T'.*

Proof. If T dominates T' then any elliptic subgroup of T is elliptic in T', so $1 \implies 2$. To show $2 \implies 1$, suppose T, T' have the same elliptic subgroups and choose nodes $p_1, \dots, p_n$ representing the classes of nodes of T under the action of G. Let H_i be the elliptic subgroup that fixes p_i and pick a point $q_i \in T'$ fixed by H_i. Then we can define an equivariant map $f : T \to T'$ by setting $f(p_i) = q_i$, extending equivariantly to the set of vertices of T and then extending linearly to edges.

To prove $1 \implies 3$, let $f : T \to T'$ and $g : T' \to T$ be equivariant maps (so, they are surjective) and let's prove that they are in fact quasi-isometries. Let $e_1, \dots, e_n$ be representatives of the edges of G with lengths $m_1, \dots, m_n$ and $m = \min_i(m_i)$. Let M_i be the diameter of $f(e_i)$ and $M = \max_i(M_i)$. Then given two points $x, y \in T$, the path from x to y passes through $\leq \frac{d(x,y)}{m} + 2$ edges of T, thus $d(f(x), f(y)) \leq M\left(\frac{d(x,y)}{m} + 2\right)$. Similarly, $d(g(x), g(y)) \leq M'\left(\frac{d(x,y)}{m'} + 2\right)$ for some M', m'. Using this and the fact that the maps $g \circ f : T \to T$ and $f \circ g : T' \to T'$ are close to the identity (because they are equivariant and continuous), it is easy to deduce that f, g are quasi isometries.

Lastly we prove $3 \implies 2$, let $f : T \to T'$ be an equivariant quasi-isometry. We can suppose f is continuous: if not change it to a map $f' : T \to T'$ which coincides with f in nodes of T and is linear on edges. Every elliptic subgroup in T is elliptic in T'. Reciprocally, let H be an elliptic subgroup in T' fixing a point y. Then the action of H on T leaves the bounded set $f^{-1}(y)$ invariant. This implies that H fixes some point of T, namely the point x_0 which minimizes the minimum distance d_{x_0} such that $f^{-1}(y) \subseteq B(x_0, d_{x_0})$. To show that this point x_0 exists and is unique, we can consider $d = \inf\{d_x; x \in T\}$ and let $(x_n)_{n\geq 1}$ be a sequence with $d_{x_n} < d + \frac{1}{n}$. Now note that if z is the middle point of x and y, then $d_z \leq \max(d_x, d_y) - \frac{d(x,y)}{2}$. This not only proves the uniqueness of x_0, but also implies that x_n is a Cauchy sequence. As metric simplicial G-trees are complete, the x_n converge to the point x_0 we are looking for. □

The proof of 1 $\implies$ 3 can be rephrased as follows:

Proposition 2.3. *Any equivariant map between equivalent metric simplicial G-trees is a quasi-isometry.*

□

2.2 Deformation spaces

Definition 2.4. A **deformation space** $\mathcal{D} \subseteq \mathcal{T}$ is a equivalence class of G-trees of $\mathcal{T}$ under $\sim$. If we identify non trivial trees of $\mathcal{D}$ which have the same metric up to rescaling, we get a **projectivized deformation space** $P\mathcal{D}$.

Deformation spaces are unions of open cones. We will define **closed cones (simplices)** in $\mathcal{D}$ ($P\mathcal{D}$) as closures in the weak topology of open cones (simplices) of $\mathcal{D}$ ($P\mathcal{D}$). They may not be homeomorphic to $[0, \infty)^n$, because some trees of $\mathcal{T}$ in the closure of the open cone may not be inside $\mathcal{D}$.

Moreover, inside a deformation space all the elements have the same type (type as in 1.5). This is because the type of a tree can be expressed just in terms of the properties from the proposition 2.2; for example, trivial G-trees are the only ones such that the whole group G is elliptic. Non trivial G-trees can be classified as follows:

	Quasi isometric to $\mathbb{R}$	Not quasi isometric to $\mathbb{R}$
All elements of $[G, G]$ are elliptic	Linear abelian	Genuine abelian
Some elements of $[G, G]$ are not elliptic	Dihedral	Irreducible

Of course, the trivial G-tree is the only point of its own deformation space, so it is not a very interesting case. As noted in [GL06], the linear abelian and dihedral cases are not very interesting either:

Proposition 2.5. *If a deformation space $\mathcal{D}$ contains a linear abelian or an irreducible tree, then $P\mathcal{D}$ has just one point.*

Proof. Let T, T' be linear abelian trees in the same deformation space. We can think of the underlying tree as $\mathbb{R}$, and the actions of G on T, T' as homomorphisms $f, f' : G \to \mathbb{R}$. Moreover, as f, f' are simplicial actions, $\mathrm{Im}(f)$ and $\mathrm{Im}(f')$ are both homothetic to $\mathbb{Z}$, and as T, T' have the same elliptic subgroups, $\ker(f) = \ker(f')$. Thus f, f' are the same up to scaling, as we wanted.

If T, T' are dihedral trees in the same deformation space, we can again think of them as actions of G on $\mathbb{R}$ by isometries. As there is an equivariant isometry between T and T', any element $g \in G$ acts in T by a translation iff it does in T'. Let H be the subgroup of G consisting of translations. Then the actions of H on T and T' are linear abelian, so by the previous case they are equal up to scaling. By scaling T, we can suppose that they are given by the same homomorphism $H \to \mathbb{R}$. Now let $s \in G$ be a symmetry in T and T', and let $p, p' \in \mathbb{R}$ be the fixed points of s in T, T' respectively. Then $x \mapsto x + (p' - p)$ is an equivariant isometry from T to T'. $\square$

So, from now on, we will just consider linear abelian and irreducible deformation spaces. We will be especially interested in the irreducible ones, as Outer Space, the main example we will study in this essay, is an irreducible deformation space.

One can wonder why deformation spaces are named that way. It is because given T, T' in the same deformation space, we will be able to *deform* T into T', in a sense which will be made precise in the next section.

2.3 Elementary deformations

Given a G-tree $T \in \mathcal{T}$, we can create a new minimal G-tree $T' \in \mathcal{T}$ by collapsing some edge e and the rest of edges of its orbit to points. Of course, T dominates T'. We then say that T' is obtained from T by a **contraction**, or that T is obtained from T' by an **expansion**. If $T \sim T'$, we say the contraction/expansion is **elementary**, and the edge e is **collapsible**.

Theorem 2.2 allows us to give a simple characterization of collapsible edges. We will define G_v to be the subgroup of G which fixes a given vertex $v \in T$, and G_e the subgroup of G fixing an edge e of T. G_e will fix both vertices of e because the action of G on T is without inversions. The next results appear in section 3 of [GL06]:

Proposition 2.6. *An edge $e = (u, v)$ of T is collapsible if and only if u, v are in distinct orbits and either $G_u \subseteq G_v$ or $G_v \subseteq G_u$.*

Proof. If u and v are in the same orbit, say $v = gu$, then g is hyperbolic in T, as there are no inversions, but collapsing e makes the subgroup $< g >$ elliptic, thus $T \not\sim T'$.

Call $w \in T'$ the vertex obtained from collapsing e. If u and v are in distinct orbits, G_w contains both G_u and G_v. If neither $G_u \subseteq G_v$ nor $G_v \subseteq G_u$, then G_w, which contains G_v and G_u, is an elliptic subgroup of T' which is not elliptic in T, so $T \not\sim T'$. However if, for example, $G_u \subseteq G_v$, then $G_w = G_v$

(here we are using again that there are no inversions) so no new elliptic subgroups are added, and $T \sim T'$. □

Proposition 2.7. *If T' is obtained from T by a (maybe not elementary) contraction and T' is not trivial, then T and T' have the same type.*

Proof. First of all, if T is not isometric to the real line, then T' isn't either: that's because there is a point $p \in T$ which divides T into ≥ 3 components, so its image in T' divides T' into ≥ 3 components too (because each component contains edges from every orbit). So T is isometric to the real line iff T' is. And in that case, clearly T contains a reflection of the real line iff T' does. Thus we have proved the result in the linear abelian and dihedral cases.

If T is genuine abelian, as all elliptic subgroups of T (including all $[G, G]$) are elliptic in T', T' is genuine abelian. If T is irreducible, let $g \in G$ be hyperbolic in T', and let A_g be the axis of g in T. Taking some hyperbolic $h \in G$ with $A_g \cap A_h$ compact, hgh^{-1} is also hyperbolic and $A_g \cap A_{hgh^{-1}}$ is compact. So in T', g, hgh^{-1} are also hyperbolic and with axes having compact intersection, thus T' is irreducible. □

Definition 2.8. An **elementary deformation** is a finite sequence of elementary collapses and expansions.

So if two trees T, T' are related by an elementary deformation, then they are in the same deformation space. In fact, as we will prove in section 4.5 (in the same order as [GL06]), the converse is also true:

Theorem 2.9. *If T, T' are in the same deformation space, then they are related by an elementary deformation.*

2.4 Topology of deformation spaces

We will need some properties of the restrictions of the topologies we have considered to deformation spaces. Let $\mathcal{D}$ be a deformation space. First of all, proposition 1.12 implies that if $\mathcal{D}$ is irreducible, the weak and Gromov topologies coincide in finite unions of cones of $\mathcal{D}$. We will not prove it but this is also true for abelian deformation spaces (see [GL06] prop. 5.2):

Proposition 2.10. *The Gromov and weak topologies have the same restriction to finite unions of cones of any deformation space $\mathcal{D}$.* □

With some extra conditions, the Gromov and weak topologies will coincide in the whole deformation space. The following is proposition 5.4 of [GL06]:

Proposition 2.11. *If $\mathcal{D}$ consists of locally finite trees with finitely generated vertex stabilizers, then the Gromov topology and the weak topology coincide on $\mathcal{D}$ and $P\mathcal{D}$.* □

3 Outer Space

This chapter introduces the main example of deformation space we will be considering: Outer Space. The classical introduction to this space is the paper [CV86], which introduces the space and proves that it is contractible. In this chapter we will define the space, an action on it by $\mathrm{Out}(\mathbb{F}_n)$ and an equivariant deformation retract of it which is a cell complex.

During this chapter, the only groups we will use are $\mathbb{F}_n$ and $\mathrm{Out}(\mathbb{F}_n)$, so unlike in the rest of the chapters, the letter G will represent a graph instead of an arbitrary group.

3.1 Marked graphs and $\mathbb{F}_n$-trees

This section first introduces marked graphs in a similar fashion to [CV86], and then shows that they are essentially a deformation space of $\mathbb{F}_n$-trees.

The group we will consider in this chapter will be $\mathrm{Out}(\mathbb{F}_n)$, the group of **outer automorphisms** of the free group $\mathbb{F}_n$, for $n \geq 2$. This group is the quotient of the group $\mathrm{Aut}(\mathbb{F}_n)$ of automorphisms of $\mathbb{F}_n$ by the normal subgroup of inner automorphisms, that is, conjugations. So an element $f \in \mathrm{Out}(\mathbb{F}_n)$ can be seen as choosing an ordered basis for $\mathbb{F}_n$ up to conjugacy.

Let R_n be the wedge of n circles and let v_n be the node of R_n. Fix an identification $\mathbb{F}_n \equiv \pi_1(R_n, v_n)$.

A **marking** (g, G) is a graph G with no vertices of valence ≤ 2 together with a homotopy equivalence $g : R_n \to G$.

This is equivalent to giving an identification $g_* : \mathbb{F}_n \to \pi_1(G, g(v_n))$. Reciprocally, any isomorphism $\mathbb{F}_n \to \pi_1(G, g(v_n))$ can be achieved with such a map.

We say two markings (g, G) and (g', G') are **equivalent** if there is an isometry (in G, the distance between two points is the length of the shortest segment joining them) $\phi : G \to G'$ such that $\phi \circ g$ is homotopic to g'. A **marked graph** is an equivalence class of markings.

Of course, changing g by a homotopy doesn't change the marked graph it represents.

In terms of maps between fundamental groups, (g, G) and (g, G') being equivalent means there is an isometry $\phi : G \to G'$ such that $(g')_*$ and $(\phi \circ g)_*$ induce the same maps of fundamental groups up to conjugacy by some path.

Any marking (g, G) induces an action of $\mathbb{F}_n$ on the covering space of G, which I will denote by $\overline{G}$. If we give $\overline{G}$ the geodesic metric which makes the projection $\overline{G} \to G$ a local isometry, it becomes a simplicial $\mathbb{F}_n$-tree.

This covering space is a tree, and we give it the geodesic metric such that the projection $\overline{G} \to G$ is an isometry.

Proposition 3.1. *Two markings are equivalent if and only if they induce $\mathbb{F}_n$-equivariantly isometric actions in the fundamental covers of their graphs.*

Proof. Suppose two markings (g, G) and (g', G') are equivalent via a map $\varphi : G \to G'$, and let $\overline{G} \xrightarrow{\pi_G} G$ and $\overline{G'} \xrightarrow{\pi_{G'}} G'$ be their universal covering spaces. The lift $\overline{\varphi} : \overline{G} \to \overline{G'}$ is certainly an isometry, as it is locally an isometry and both trees have geodesic metrics. To check is is equivariant, let $g \in \mathbb{F}_n$ and $x \in \overline{G}$, we want to prove that $g(\overline{\varphi}(x)) = \overline{\varphi}(gx)$.

Notice that gx is $\overline{\alpha}(1)$, where $\overline{\alpha} : [0,1] \to \overline{G}$ is the path obtained from lifting the loop α based in $\pi_G(x)$ and representing the conjugacy class of g, and with $\overline{\alpha}(0) = x$. Moreover, the path $\overline{\varphi} \circ \overline{\alpha}$ satisfies $\overline{\varphi} \circ \overline{\alpha}(0) = \overline{\varphi}(x)$, and $\pi_{G'}(\overline{\varphi} \circ \overline{\alpha}) = \varphi(\pi_G \circ \overline{\alpha}) = \varphi(\alpha)$, which by assumption represents the conjugacy class of g in G'. Thus $\overline{\varphi} \circ \overline{\alpha}(1) = g\overline{\varphi}(x)$. However $\overline{\varphi} \circ \overline{\alpha}(1)$ is also $\overline{\varphi}(gx)$, so we are done.

Conversely, given two markings (g, G) and (g', G'), any equivariant isometry $\overline{\varphi} : \overline{G} \to \overline{G'}$ descends to an isometry $\varphi : G \to G'$. We can suppose $g(v_n) = g'(v_n)$ changing them by a homotopy, and then, if we choose the lift $\overline{\varphi}$ correctly, we have the following commuting triangle of equivariant maps:

$$\begin{tikzcd} & (\overline{R_n}, \overline{v_n}) \arrow[dl, "\overline{g}"'] \arrow[dr, "\overline{g'}"] & \\ (\overline{G}, \overline{g}(\overline{v_n})) \arrow[rr, "\overline{\varphi}"] & & (\overline{G'}, \overline{g'}(\overline{v_n})) \end{tikzcd}$$

A loop $\alpha \subseteq G$ based at $g(v_n)$ represents an element $w \in \mathbb{F}_n$ iff its lift beginning in $\overline{g}(\overline{v_n})$ ends in $w\overline{g}(\overline{v_n})$. In that case, $\overline{\varphi}\alpha$ begins in $\overline{g'}(\overline{v_n})$ and ends in $w\overline{g'}(\overline{v_n})$, so the loop $\pi_{G'}(\overline{\varphi}\alpha) = \varphi\alpha$ in G' represents the element w of $\mathbb{F}_n$. Thus φ makes the two markings equivalent. □

So we can see marked graphs as minimal simplicial $\mathbb{F}_n$-trees[2]. Moreover, as actions of fundamental groups are free, the only elliptic subgroup of marked graphs is the identity. Reciprocally, any minimal simplicial $\mathbb{F}_n$-tree T with its only elliptic subgroup being $\{e\}$ is the fundamental cover of the finite graph $G := \frac{T}{\mathbb{F}_n}$, and we can define a map $R_0 \to \frac{T}{\mathbb{F}_n}$ such that its lift $\overline{R_0} \to T$ is $\mathbb{F}_n$-equivariant. Thus, T comes from some marking. So we have proved the following:

Proposition 3.2. *Let $\mathcal{D}_n$ be the set of marked graphs with fundamental group of rank n. Then $\mathcal{D}_n$ is the deformation space of $\mathbb{F}_n$-trees with only elliptic subgroup the identity.* □

3.2 Outer Space as a projectivized deformation space

Proposition 3.3. *$\mathcal{D}_n$ is irreducible and the axis, Gromov and weak topologies coincide in $\mathcal{D}_n$.*

Proof. $\mathcal{D}_n$ is irreducible because, given a marked graph $(g, G) \in \mathcal{D}_n$, and two simple closed curves $\gamma_1, \gamma_2 \in \pi_1(G)$, they induce two hyperbolic isometries in $\overline{G}$ whose axes have compact intersection.

[2]One way to show that the fundamental covers are minimal $\mathbb{F}_n$-trees is checking that every edge is in the axis of some isometry.

As $\mathcal{D}_n$ is not abelian, the Gromov and axis topologies coincide in it due to 1.9. Moreover, the trees of $\mathcal{D}_n$ are locally finite and they have trivial vertex stabilizers, so 2.11 says that the Gromov and the weak topologies coincide in it. □

What dimension can the open cones of $\mathcal{D}$ have? Well, the dimension of a cone is just the number of edges of a graph representing some point of the cone.

Proposition 3.4. *If a marked graph is given by $g : R_n \to G$, then G has $\leq 3n-3$ edges and $\leq 2n-2$ nodes, and these maxima can be achieved in graphs without separating edges.*

Proof. First of all, G has finitely many edges by 1.4. Let E/V be the number of edges/nodes of G. Then, as every node has valence ≥ 3, $2E \geq 3V$. On the other hand, as the fundamental group of G is $\equiv \mathbb{F}_n$, $E = V + n - 1$. Combining these relations we obtain $V \leq 2n-2$, thus $E = V + n - 1 \leq 3n-3$.

To prove the second part of the proposition, begin with some marked graph G without separating edges (for example the rose of n petals, R_n). If G has some node v with valence ≥ 4, we can expand v into an edge in such a way that the new marked graph still has edges of degree ≥ 3 and no separating edges.

We can repeat this process until we have a graph G' with non separating edges and all nodes of valence 3, so all the inequalities above become equalities and G' has $3n-3$ edges and $2n-2$ nodes. □

It will be useful to know the relation between the lengths of the edges of a marked graph (g, G) and the lengths of the elements $w \in \mathbb{F}_n$ in $\overline{G}$.

Proposition 3.5. *Given $w \in \mathbb{F}_n$ and a marked graph (g, G), there is a unique cycle of minimal length in G representing the conjugacy class of w.*

Proof. Given a loop γ representing w in G, its lift $\overline{\gamma}$ in $\overline{G}$ goes from some point p to wp, so it has length $\geq l_{\overline{G}}(w)$. Equality will be achieved if and only if p is in the axis of w and γ is the image of the axis of w under the quotient $\overline{G} \to G$. □

Corollary 3.6. *The length of an element $w \in \mathbb{F}_n$ is the sum of lengths of the edges traversed in the minimal loop representing w, so for G inside a closed cone, $l_G(w)$ depends linearly on the lengths of the edges of G.*

Definition 3.7. The **volume** of a graph G, $\mathrm{vol}(G)$, is the sum of lengths of its edges.
Outer Space, X_n, is the set of marked graphs with fundamental group of rank n and volume 1.

We consider it as a subspace of $\mathcal{D}_n$ (see 3.2) and we give it the subspace topology.

Proposition 3.8. *Outer Space is homeomorphic to the projectivized deformation space $P\mathcal{D}_n$ of $\mathbb{F}_n$-trees whose only elliptic subgroup is the identity.*

Proof. The composition map $X_n \hookrightarrow \mathcal{D}_n \to P\mathcal{D}_n$ is a continuous bijection, we just need to see that its inverse is continuous, and for that it is enough to see that the map $\mathcal{D}_n \to \mathbb{R}; T \mapsto \mathrm{vol}\left(\frac{T}{\mathbb{F}_n}\right)$ is continuous. This is obvious considering the weak topology on $\mathcal{D}_n$, as $\mathrm{vol}\left(\frac{T}{\mathbb{F}_n}\right)$ is just the sum of lengths of the (orbits of) edges of T. □

So we can see Outer Space as an 'incomplete' cell complex of dimension $3n-4$. There are some missing open cells because if a marked graph has an edge e which is a loop, then the graphs obtained from contracting e to a point are not in X_n, because they do not have fundamental group of rank n.

Moreover, we can define an action of $\mathrm{Out}(\mathbb{F}_n)$ on the complex: For any $\phi \in \mathrm{Out}(\mathbb{F}_n)$ we can find a map $\overline{\phi} : R_n \to R_n$ which induces ϕ in the fundamental group; it is determined up to homotopy. Thus we can define the action of ϕ on X_n by sending a marked graph represented by (g, G) to $(g \circ \overline{\phi}, G)$.

3.3 The deformation retract

This section describes K_n, an equivariant deformation retract of X_n that actually is a cell complex.

First of all, consider the subspace $Y_n \subseteq X_n$ consisting of marked graphs without separating edges. Y_n is also a union of open cones, and by proposition 3.4, it also has dimension $3n-4$. Y_n also carries an action of $\mathrm{Out}(\mathbb{F}_n)$, in fact:

Proposition 3.9. *There is an equivariant deformation retraction from X_n to Y_n.*

Proof. We will define $\varphi : X_n \times [0,1] \to X_n$ with $\varphi_0 = Id_{X_n}$ and $\varphi_1(X_n) \subseteq Y_n$ in the following way: given a marked graph in which the sum of lengths of the separating edges is $k \in [0,1)$, define $\varphi_t(G)$ as the same marked graph but uniformly scaling the separating edges so that they have sum of lengths kt and uniformly scaling the non separating edges so that they have length $1-kt$. The length functions $\mathcal{C} \times [0,1] \xrightarrow{\varphi} X_n \xrightarrow{l_w} \mathbb{R}$ are continuous for any simplex $\mathcal{C}$ and $w \in \mathbb{F}_n$ due to 3.6, so φ is continuous. □

However Y_n is still not a cell complex. We will now describe a retraction from Y_n to a subspace of it, K_n, which is a cell complex. Each closed (closed in Y_n) k-simplex $S \subseteq Y_n$ can be imbedded in a compact k-simplex $\overline{S}$. Now consider the barycentric subdivision of $\overline{S}$. Each simplex A of the subdivision has as vertices a sequence of marked graphs $\rho_1, \rho_2, \ldots, \rho_m$ (for some m) such that ρ_{i+1} is obtained by collapsing an edge from ρ_i. Thus there is some l such that $\rho_i \in Y_n$ iff $i \leq l$, and the face of A spanned by $\rho_1, \ldots, \rho_l$ is the maximal closed face of A contained in Y_n. We will then define the retraction in $A \cap Y_n$ as linearly retracting it to the face spanned by $\rho_1, \ldots, \rho_l$. It is not hard to see that this retraction is well defined in all of X_n, continuous (using the weak topology) and equivariant.

The image of this retraction, K_n, is a simplicial complex. Its vertices are barycenters of simplices of Y_n, so they are marked graphs with volume 1, without separating edges and with all edges of the same length. What is the dimension of K_n? Well, as we have seen, vertices of a given simplex of K_n

are sequences of marked graphs $\rho_1, \ldots, \rho_l$, each one obtained by collapsing an edge of the previous one. The proof of 3.4 implies that we can obtain ρ_1 with $3n-3$ edges and ρ_l with n edges, for a maximum $l = 2n-2$. So,

Proposition 3.10. *K_n is a cell complex of dimension $2n-3$.* □

As we will see in 4, the cell complex structure of K_n, together with its action by Out($\mathbb{F}_n$), will give us information about the group Out($\mathbb{F}_n$). We will need another result about K_n:

Proposition 3.11. *K_n has finite cell stabilizers.*

Proof. It is enough to prove that it has finite vertex [3] stabilizers. Given a vertex of K_n, $\rho = (g, G)$, and $f \in$Out($\mathbb{F}_n$), $f\rho = \rho$ means that there is a diagram commuting up to homotopy

$$\begin{array}{ccc} R_n & \xrightarrow{g} & G \\ \downarrow{\scriptstyle \overline{f}} & & \downarrow{\scriptstyle \phi} \\ R_n & \xrightarrow{g} & G \end{array}$$

, where ϕ is an isometry and $\overline{f}$ induces the automorphism f. However there are only finitely many isometries ϕ from G to itself, thus there are only finitely many possible automorphisms f which can make the diagram commute. □

3.4 The case $n = 2$

We can give a nice, explicit description of the complex Y_2 (and of X_2, but it is not so nice).

By 3.4, Y_2 will be an 'incomplete' CW complex of dimension 2. 2-cells contain marked graphs with 3 edges and 2 vertices, and as they have no separating edges or vertices of valence ≤ 2, these graphs are simply 2 nodes with 3 edges joining them. 1-cells of Y_2 contain just wedges of 2 circles, and there are no 0-cells, as graphs with 1 edge cannot have fundamental group $\equiv \mathbb{F}_2$.

This section will describe a homeomorphism from Y_2 to the hyperbolic plane $\mathbb{H}^2$ in its upper half plane model. The 1-cells will correspond to geodesics between certain points $(q, 0) \in \mathcal{H}^2$, with the 2-cells being the triangles left between these geodesics.

But first we need some facts about the free group $\mathbb{F}_2$ first.

The abelianization map ab:$\mathbb{F}_2 \to \mathbb{Z}^2$ induces a surjective map Out($\mathbb{F}_2$) $\to$Aut($\mathbb{Z}^2$) =GL($2, \mathbb{Z}$). This map is injective, as proved by Nielsen in [Ni17]. This is equivalent to the following statement:

Proposition 3.12. *For each pair of points $p = (a, b), q = (c, d)$ in $\mathbb{Z}^2$ with $ad - bc = \pm 1$ there is a unique base (w_1, w_2) of $\mathbb{F}_2$ up to conjugacy with $ab(w_1) = p$ and $ab(w_2) = q$.* □

We also need some notation about the upper half plane: For any two points q_1, q_2 of $\mathbb{R} \cup \infty \equiv \partial\mathbb{H}^2$, we will let γ_{q_1,q_2} be the geodesic joining them.

[3]We have been calling 0-cells of G-trees 'nodes', but we will call 0-cells of deformation spaces 'vertices' to distinguish between them.

Let's see now how we can identify outer space with the upper half plane. We will begin with 1-simplices: let $\rho = (r, R)$ be a rose with edges e_1, e_2 representing a basis (w_1, w_2) of $\mathbb{F}_2$, and let $(a, b) = ab(w_1)$, $(c, d) = ab(w_2)$. Then our homeomorphism will send the open 1-simplex centered in ρ to $\gamma_{\frac{a}{b},\frac{c}{d}}$ (we will see how exactly in 3.18).

So 1-cells correspond to pairs $(a, b), (c, d)$ such that $ad - bc = \pm 1$.

Definition 3.13. Let $\frac{a}{c}$ and $\frac{b}{d}$ be distinct points of $\mathbb{Q} \cup \infty$ given by irreducible fractions (∞ is given by $\frac{1}{0}$). We say they are **related** if $ad - bc \in \{1, -1\}$.

Let $\mathcal{G}$ be the set of geodesics in $\mathcal{H}^2$ between $\left(\frac{a}{b}, 0\right)$ and $\left(\frac{c}{d}, 0\right)$, where $\frac{a}{b}$ and $\frac{c}{d}$ are related.

For given a, b with $\gcd(a, b) = 1$, the set of points related to $\frac{a}{b}$ is given by $\frac{c_0+ma}{d_0+mb}$, for $m \in \mathbb{Z}$ and some initial solution $\frac{c_0}{d_0}$. Using this it is easy to deduce the following statement:

Lemma 3.14. *Given related points $\frac{a}{b}, \frac{c}{d} \in \mathbb{Z}^2$, the only points related to both of them are $\frac{a+c}{b+d}$ and $\frac{a-c}{b-d}$.* □

As any two 1-cells have to be disjoint from each other, we need the geodesics of $\mathcal{G}$ to be pairwise disjoint. This is a consequence of the following technical lemma:

Lemma 3.15. *If $\frac{a_1}{b_1}$ is related to $\frac{a_2}{b_2}$ and $\frac{c_1}{d_1}$ is related to $\frac{c_2}{d_2}$, then $\gamma_{\frac{a_1}{b_1},\frac{a_2}{b_2}}$ and $\gamma_{\frac{c_1}{d_1},\frac{c_2}{d_2}}$ either are equal or do not intersect.*

Proof. We will prove the case where $\frac{a_1}{b_1}, \frac{c_1}{d_1}, \frac{a_2}{b_2}, \frac{c_2}{d_2}$ are rational and all distinct, the other cases are easier. If the geodesics from the statement intersect, then $\frac{a_1}{b_1} < \frac{c_1}{d_1} < \frac{a_2}{b_2} < \frac{c_2}{d_2}$ (up to relabelling and sign).

So $\frac{c_1}{d_1} - \frac{a_1}{b_1} < \frac{a_2}{b_2} - \frac{a_1}{b_1} = \frac{\pm 1}{b_1 b_2}$. Comparing the denominators, we have that $d_1 b_1 > b_1 b_2$, so $d_1 > b_2$. Similarly we use $\frac{c_2}{d_2} - \frac{a_2}{b_2} < \frac{c_2}{d_2} - \frac{c_1}{d_1} = \frac{\pm 1}{d_1 d_2}$ to deduce that $b_2 > d_1$, a contradiction. □

So here is the picture: we have a set $\mathcal{G}$ of disjoint geodesics in the upper half plane $\mathcal{H}^2$ between related points of $\mathbb{Q} \cup \{\infty\}$. Moreover, each geodesic γ is contained in two geodesic triangles T_1, T_2 with vertices in $\mathbb{Q} \cup \{\infty\}$. As the edges of both triangles cannot intersect outside of γ, one of them has to be inside the half disk between γ and $\mathbb{R}$, and the other one is outside of it. Also, no geodesic of $\mathcal{G}$ passes inside the interior of these geodesic triangles, because they would intersect their edges in that case.

Proposition 3.16. *The set $\mathcal{G}$ is locally finite in $\mathcal{H}^2$ and the geodesics of $\mathcal{G}$ and the triangles they form cover all the upper half plane.*

Proof. In the first part, it is clear that the set of geodesics joining ∞ with rational numbers is locally finite (these geodesics are given by $x = n$, for $n \in \mathbb{Z}$). For geodesics γ joining two rational numbers, notice that if a geodesic γ joins $\frac{a}{b}$ with $\frac{c}{d}$, then $ad - bc = 1$ so $\frac{a}{b} - \frac{c}{d} = \frac{1}{bd}$, so it is a semicircumference of radius $\frac{1}{2bd}$ centered in the real line. Meaning that, given a point $p = (x, y) \in \mathbb{R}^2$, for γ to pass near p

we need $\frac{1}{2bd} > \frac{y}{2}$ (which implies $b, d < \frac{1}{y}$) and, as the radius of the geodesic is ≤ 1, the extreme points $\frac{a}{b}$ and $\frac{c}{d}$ of the geodesic have to be at distance ≤ 2 of p. This leaves finitely many possibilities for a, b, c, d, which proves the first part.

For the second part, take some point $p \in \mathbb{H}^2$. If $p \in \mathbb{Z}$, then p is contained in a geodesic of $\mathcal{G}$. If not, if we let $x(p)$ be the x coordinate of p, then there is a geodesic of G going from $\lfloor x(p) \rfloor$ to $\lfloor x(p) \rfloor + 1$. If p is above this geodesic, then it is contained in the triangle between $\lfloor x(p) \rfloor, \lfloor x(p) \rfloor + 1$ and ∞. If not, by the first part of this proposition there has to be some geodesic γ directly above p (that is, the vertical segment from p to γ doesn't intersect other geodesics). This implies that p is contained in the triangle below γ. $\square$

As [CV86] mentions, the structure we have been describing of $\mathbb{H}^2$ as a CW complex is the **Farey complex** of the modular group. The figure to the right contains a few 1-cells of the complex in $[0, 1] \times \mathbb{R}^+ \subseteq \mathbb{H}^2$.

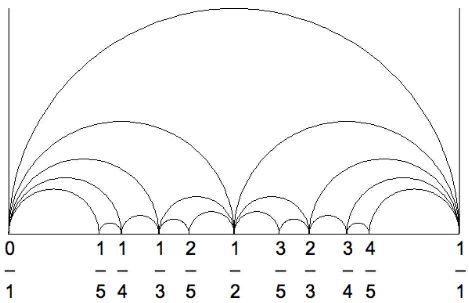

To complete the homeomorphism from Y_2 to $\mathbb{H}^2$ we just have to define the map in 2-cells.

Let a marked graph $\rho = (g, G)$ be the barycenter of a 2-simplex, so it has three edges of length $\frac{1}{3}$. Then, up to homotopy, there are only 6 loops of length $\frac{2}{3}$ in G. Take $u, v \in \mathbb{F}_n$ of length $\frac{2}{3}$ and such that their directions agree in their common edge. The next figure represents a marking of ρ with the loops representing u, v and uv^{-1}:

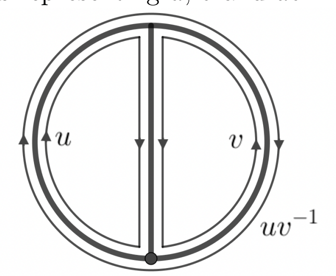

And the six elements of length $\frac{2}{3}$ are $u, v, u^{-1}, v^{-1}, uv^{-1}, vu^{-1}$. Under the abelianization $\mathbb{F}_2 \xrightarrow{\text{ab}} \mathbb{Z}^2$, these six elements are represented by three points: $u, u^{-1} \mapsto (a, b)$, $v, v^{-1} \mapsto (c, d)$ and $uv^{-1}, vu^{-1} \mapsto (a - c, b - d)$. Note that u, v form a basis of the fundamental group, so the points $\frac{a}{b}$ and $\frac{c}{d}$ are related, thus by 3.14, the points $\frac{a}{b}$, $\frac{c}{d}$ and $\frac{a-c}{b-d}$ form a geodesic triangle of $\mathcal{G}$ in the upper half plane. We will identify this triangle with the 2-simplex centered at ρ.

Proposition 3.17. *The correspondence described above gives a bijection*

$$\{\textit{Open 2-simplices of } Y_2\} \longrightarrow \{\textit{Geodesic triangles between the geodesics of } \mathcal{G} \textit{ in } \mathbb{H}^2\}$$

Proof. The map is well defined (it doesn't depend on the choice of u, v) because for each 2-simplex s of Y_2, the points described above, $\frac{a}{b}$, $\frac{c}{d}$ and $\frac{a-c}{b-d}$, are the abelianizations of the only loops with length $\frac{2}{3}$ of the marked graph in the barycenter of s.

To conclude that the map is surjective, first notice that every geodesic triangle is of the form $\frac{a}{b}, \frac{c}{d}, \frac{a-c}{b-d}$ for some $a, b, c, d \in \mathbb{Z}$. So we just have to prove that if $\frac{a}{b}, \frac{c}{d} \in \mathbb{R} \cup \infty$ are related, then there is some basis u, v of $\mathbb{F}_2$ with $\mathrm{ab}(u) = (a, b)$, $\mathrm{ab}(v) = (c, d)$. If the vectors (a, b) and (c, d) have norm 1, this is obvious. Moreover, if this is true for a pair of vectors $(a, b), (c, d)$, it is true for $(a + c, b + d), (c, d)$ and $(a - c, b - d), (c, d)$. Also notice that if $|(a, b)| > |(c, d)|$, then either $(a - c, b - d)$ or $(a + c, b + d)$ has norm $< |(a, b)|$: this is an easy consequence of the fact that, as $|(a, b)| \geq \sqrt{2}$ and $ad - bc = 1$, the vectors (a, b) and (c, d) form an angle $\leq \frac{\pi}{4}$. So we can conclude by induction on $\max(a^2 + b^2, c^2 + d^2)$.

To prove injectivity, suppose two simplices s_1, s_2 with barycenters $\rho_1 = (g_1, G_1)$, $\rho_2 = (g_2, G_2)$ have the same geodesic triangle as an image, with vertices $\frac{a}{b}, \frac{c}{d}$ and $\frac{a-c}{b-d}$ for some $a, b, c, d \in \mathbb{Z}$. For $i = 1, 2$, let u_i, v_i be the elements of $\mathbb{F}_2$ with $\mathrm{ab}(u_i) = (a, b)$, $\mathrm{ab}(v_i) = (c, d)$ and such that u_i, v_i have length $\frac{2}{3}$ in ρ_i. Then the directions of u_i, v_i agree in their common edge (if not, $\frac{a+c}{b+d}$ would be the vertex of the triangle, instead of $\frac{a-c}{b-d}$). Thus there is an isometry $\varphi : G_1 \to G_2$ sending the loops representing u_1, v_1 in ρ_1 to the loops representing u_2, v_2 in ρ_2. Moreover, due to 3.12, the base (u_1, v_1) is the same as the base (u_2, v_2) up to conjugacy. So φ makes the markings ρ_1 and ρ_2 equivalent, proving injectivity. $\square$

Theorem 3.18. *Y_2 is homeomorphic to the upper half plane.* $\square$

Proof. We will first see how to define the map in 2-cells. Let ρ be the barycenter of a closed 2-simplex $s \subseteq Y_2$, and let u, v, a, b, c, d as in the discussion before theorem 3.17. We will let $\rho_{x,y,z}$, with $x+y+z = 1$, be the marked graph of the simplex of ρ with $l(u) = 2x, l(v) = 2y, l(uv^{-1}) = 2z$. We will also identify the geodesic triangle with vertices $\frac{a}{b}, \frac{c}{d}$ and $\frac{a-c}{b-d}$ with an euclidean triangle, like the one shown below with vertices $p_{\frac{a}{b}}, p_{\frac{c}{d}}$ and $p_{\frac{a-c}{b-d}}$.

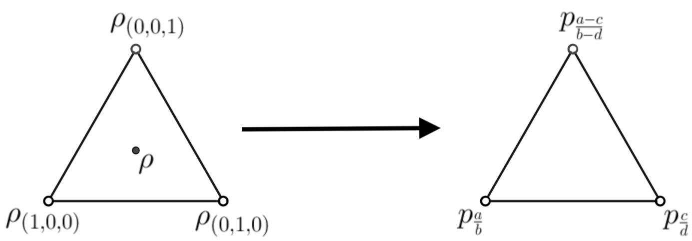

Then the homeomorphism will send the marked graph $\rho_{x,y,z}$ to the point $xp_{\frac{a}{b}} + yp_{\frac{c}{d}} + z\frac{a-c}{b-d}$. Every 1-simplex is sent to the same geodesic by the maps of the two adjacent 2-simplices, and in the same orientation.

So we can define a map $F : Y_2 \to \mathbb{H}^2$ in the following way:

First define the restriction of F to 1-simplices, giving them the correct orientation (so that it will agree with the maps we have described of 2-simplices). Then, for each 2-simplex s, the map $F|_{\partial s}$ is already determined, and it sends ∂s to the ∂T, where $T \subseteq \mathbb{H}^2$ is the geodesic triangle corresponding to s. So we just have to extend the map $F|_{\partial s}$ to a homeomorphism $F_s : s \to T$, and we are done. $\square$

Now we will study the action of Out($\mathbb{F}_2$) on Y_2.

As mentioned in the beginning of the section, Out($\mathbb{F}_2$) $\equiv$GL$_2(\mathbb{Z})$. And GL$_2(\mathbb{Z})$ acts on $\mathbb{H}^2 \equiv \frac{\mathbb{C}\setminus\mathbb{R}}{x\sim -x}$ by Möbius transformations, with each matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ sending z to $\frac{az+b}{cz+d}$. Notice that $M\frac{x}{y} = \frac{ax+by}{cx+dy}$, so if we identify rationals $\frac{x}{y}$ with vectors $\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right)$, product still works the same way.

Now, two rationals $\frac{x_1}{y_1}$ and $\frac{x_2}{y_2}$ are related iff $\det\left(\begin{smallmatrix} x_1 & x_2 \\ y_1 & y_2 \end{smallmatrix}\right) \in \{1, -1\}$, so as matrices of GL$_2(\mathbb{Z})$ have determinant ± 1, $\frac{x_1}{y_1}$ and $\frac{x_2}{y_2}$ are related iff $M\frac{x_1}{y_1}$ and $M\frac{x_2}{y_2}$ are related. Thus the action of GL$_2(\mathbb{Z})$ on $\mathcal{H}^2$ preserves the simplicial structure of the Farey diagram. In fact the action of GL$_2(\mathbb{Z})$ on the set of cells of the Farey complex is the same as the action of Out($\mathbb{F}_2$) (under the identification Out($\mathbb{F}_2$) $\equiv$GL$_2(\mathbb{Z})$).

In the proof of theorem 3.18, if we identify the we define the maps from the edges and simplices of Y_2 to $\mathbb{H}^2$ equivariantly, we get that:

Theorem 3.19. *There is an homeomorphism $F : Y_2 \to \mathbb{H}^2$ such that the action of Out($\mathbb{F}_2$) on Y_2 corresponds via F to the action of GL$_2(\mathbb{Z})$ on $\mathcal{H}^2$.*

4 The proof of contractibility

This chapter is mostly based in section 6 of [GL06] and section 3 of [GL07], and it gives sufficient conditions for deformation spaces to be contractible.

Concretely, we will prove the following results:

Theorem 4.1. *Any deformation space is contractible in the weak topology.*

Theorem 4.2. *If a deformation space $\mathcal{D}$ contains a tree with finitely generated vertex stabilizers, then $\mathcal{D}$ is contractible in the Gromov topology.*

Notice that, due to proposition 1.11, it is equivalent to prove the contractibility of $\mathcal{D}$ or $P\mathcal{D}$, so:

Corollary 4.3. *Outer Space is contractible.* $\square$

4.1 Outline of the proof

Let $\mathcal{D}$ be a deformation space. To prove the contractibility of $\mathcal{D}$, we will define a homotopy $H : \mathcal{D} \times [0,1] \to \mathcal{D}; H(T,t) = H_t(T)$, with $H_0 =$Id$_{\mathcal{D}}$ and $H_1(\mathcal{D})$ contained in $\mathcal{C}_0$, the closed[4] cone in $\mathcal{D}$ of some tree $T_0 \in \mathcal{D}$. The tree T_0 will be chosen to have finitely generated vertex stabilizers in the case of theorem 4.2. This proves the theorems above because cones are contractible in the weak topology, and also in the Gromov topology (in cones both topologies coincide, see 2.10).

[4] In [GL06], the tree T_0 is assumed to be reduced. In that case, $\mathcal{C}_0$ is in fact an open cone.

This chapter will contain a lot of $\varepsilon - \delta$ proofs, as will use almost exclusively the Gromov topology. This may seem to contradict the fact that theorem 4.1 makes a statement about the weak topology, however the only thing we will use about the weak topology apart from its definition is that it coincides with the Gromov topology in finite unions of cones, that is, proposition (2.10). During this chapter, when we make a statement about a deformation space $\mathcal{D}$, we will be referring to $\mathcal{D}$ with the Gromov topology.

As in [GL07], we define a topology similar to the Gromov topology but for maps between G-trees:

Definition 4.4. Let $\mathcal{M}$ be the set of equivariant maps (continuous functions) between G-trees, where we identify two maps $f : S \to T$ and $g : S' \to T'$ if there are equivariant isometries $\varphi : S \to S'$ and $\psi : T \to T'$ such that $g \circ \varphi = \psi \circ f$.

For any given map $f : T \to T_\infty$ between G-trees, a neighborhood of f is denoted by $W_f(X, A, \varepsilon)$, where $X \subseteq T$ is finite, $A \subseteq G$ is finite and $\varepsilon > 0$. $W_f(X, A, \varepsilon)$ is the set of maps $f' : T' \to T'_\infty$ such that there is a map $\psi : X \to T'; x \to x'$ with $|d(x, ay) - d(x', ay')| < \varepsilon$ and $|d(f(x), f(ay)) - d(f'(x'), f'(ay'))| < \varepsilon$.

We also define $\mathcal{D}'$ as the set of not necessarily minimal G-trees whose minimal subtree is in $\mathcal{D}$, with the Gromov topology[5].

Defining the homotopy is not going to be hard: it will be given by the following composition of several maps that we will outline now and explain in detail in the next sections. Defining these maps is not hard; the complicated part is showing that they are continuous.

$$\begin{array}{ccccccc} \mathcal{D} \times [0,\infty] & \to & \mathcal{M} \times [0,\infty] & \mapsto & \mathcal{D}' & \to & \mathcal{D} \\ (T,t) & \mapsto & (f_T,t) & \mapsto & T_t & \mapsto & T_t^m \end{array}$$

- The map $(T,t) \mapsto (f_T,t)$ comes from associating to each tree $T \in \mathcal{D}$ a map $f_T : T_0(T) \to T$, which will be a function continuously dependent on T and with $T_T \in \mathcal{C}_0$.
- The map $\mathcal{M} \times [0,\infty] \mapsto \mathcal{D}'$ will be defined in the section about Skora's deformation: we will take the function $f_T : T_0(T) \to T$ and construct from it a path of trees T_t, with $t \in [0,\infty]$, going from $T_0(T)$ to T, and depending continuously on T and t.
- The last map just takes the tree T_t to its minimal subtree T_t^m.

We can show directly that the last map is continuous:

Proposition 4.5. *The map $\mathcal{D}' \to \mathcal{D}$ associating to each tree its minimal subtree is continuous.*

Proof. Let T be a tree, T^m its minimal subtree and $V_{T^m}(X^m, A, \varepsilon)$ a neighborhood of T^m. We want to find a Gromov neighborhood U of T such that if $T' \in U$ via a function $X \to T'; x \mapsto x'$, then $(T')^m \in$

[5] As defined, $\mathcal{D}'$ is not a set (neither is $\mathcal{M}$). This can be fixed by bounding the cardinal of the trees.

$V_{T^m}(X^m, A, \varepsilon)$, that is, there is a function $f^m : X^m \to (T')^m$ such that $|d(x, ay) - d(f^m(x), af^m(y))| < \varepsilon$ $\forall x \in X^m, a \in A$.

Specifically, we will take $f^m = \pi' \circ f$, where $\pi' : T' \to (T')^m$ is the projection. So, we want $|d(x, ay) - d(\pi'(x'), a\pi'(y'))| < \varepsilon$ to hold $\forall x \in X^m, a \in A$. This is implied by the following conditions:

1. $|d(x, ay) - d(x', ay')| < \frac{\varepsilon}{3}$ $\forall x \in X^m, a \in A$.

2. $|d(ax', \pi'(ax'))| < \frac{\varepsilon}{3}$ $\forall x \in X^m, a \in A$.

Condition 1 is satisfied for any tree $T' \in V_T(X^m, A, \frac{\varepsilon}{3})$. We can also find a Gromov neighborhood of T in which condition 2 holds: as X^m and A are finite, it is enough to do it for a fixed $x \in AX^m$. As $x \in T^m$, there is some T-hyperbolic element g with $x \in A_g$, the axis of g. So, $2d(g, x) - d(g^2x, x) = 0$. There is a Gromov neighborhood V of T such that for every $T' \in V$, g is hyperbolic in T' (by 1.7, length functions are continuous in the Gromov topology) and $2d(g, x') - d(g^2x', x') < \frac{2\varepsilon}{3}$, thus $d(x', \pi'(x')) = d(x', (T')^m) \leq d(x', A_g) = \frac{2d(g,x') - d(g^2x', x')}{2} < \frac{\varepsilon}{3}$ and we are done. □

4.2 Skora's Deformation

Definition 4.6. A **morphism** between $\mathbb{R}$-trees T, T' is an equivariant map $f : T \to T'$ such that any segment $[a, b]$ of T can be subdivided into finitely many segments in such a way that the restriction of f to those segments is an isometry.

Let $f : T_0 \to T_\infty$ be a morphism between $\mathbb{R}$-trees. As in [GL07], we want to use f to construct a continuous path of trees T_t, for $t \in [0, \infty]$, such that when $t = 0, \infty$, T_t coincides with our given trees T_0 and T_∞ and we have morphisms $T_0 \xrightarrow{\varphi_t} T_t \xrightarrow{\psi_t} T_\infty$ with $\psi \circ \varphi = f$. This implies that if T_0 and T_∞ are in the same deformation space $\mathcal{D}$, then $T_t \in \mathcal{D}$ $\forall t$.

Definition 4.7. Let $\sim_t$ be the equivalence relation in T_0 given by $a \sim_t b$ if $f(a) = f(b)$ and $f([a, b]) \subseteq B(f(a), t)$[6]. Then we define $T_t = T_0/\sim_t$. The map f descends to the quotient, inducing $\psi_t : T_t \to T_\infty$.

Notice that in the cases $t = 0, \infty$, the definition of T_t coincides with the original trees T_0, T_∞. We want to define a metric d_t on T_t that satisfies the following properties:

- T_t is a tree
- $\varphi_t : T_0 \to T_t$ is a morphism of trees

These two conditions uniquely determine d_t. As T_t is defined as a quotient of T_0, we can imagine paths in T_t to be like paths on T_0 but we will be able to 'jump' between two points a, b if they are identified. This inspires the following definition:

[6]In any metric space, we let $B(x, r)$ be the ball of center x and radius r.

Definition 4.8. A t-path γ between two points $a, b \in T_0$ is a sequence of segments $[x_0, y_0], \ldots, [x_n, y_n]$ such that $x_0 = a, y_n = b$ and $y_i \sim_t x_{i+1}$ for all i. The *length* of γ, $l(\gamma)$, is defined to be $\sum_i d(x_i, y_i)$

We define a pseudometric δ_t in T_0 by $\delta_t(a,b) = \inf\{l(\gamma); \gamma\ t\text{-path from } a \text{ to } b\}$. Clearly if $a \sim_t a'$ and $b \sim_t b'$, then $\delta_t(a,b) = \delta_t(a',b')$, so δ_t descends to a pseudometric d_t in T_t.

A priori we don't even know if the pseudometric d_t could be 0, but in fact the space (T_t, d_t) turns out to be very nice:

Proposition 4.9. *Given a morphism $f : T \to T'$ between $\mathbb{R}$-trees, T_t is a $\mathbb{R}$-tree with the metric d_t and φ_t, ψ_t are morphisms $\forall t$. Moreover, the maps $\Psi : \mathcal{M} \times [0,\infty] \to \mathcal{M}; (f,t) \mapsto \psi_t$ and $\Phi : \mathcal{M} \times [0,\infty] \to \mathcal{M}; (f,t) \mapsto \varphi_t$ are continuous.* □

Of course, as the map $\psi_t \mapsto T_t$ is continuous (by definition of the topologies), this proposition implies the continuity of the map $(f_T, t) \mapsto T_t$ from section 4.1. The proof of proposition 4.9 is long, it can be found divided in several parts in section 3 of [GL07].

4.3 Choice of basepoint P_T

Before defining the map $T \mapsto f_T$ from section 4.1, we need to define a basepoint $P_T \in T$ 'continuously on T'. The definition of P_T will be similar (but not exactly the same) to the one in section 6 of [GL06], and we will prove continuity in a slightly different way.

As $\mathcal{D}$ is either genuine abelian or irreducible, we can pick two elements $g, h \in G$ hyperbolic in $\mathcal{D}$ with different axes (g, h having different axes is independent of the tree inside $\mathcal{D}$). Let A_g, A_h be the axes of g, h in some tree T. h translates A_h towards an end, which we call A_h^+. We can suppose A_h^+ is not an end of A_g in any tree of $\mathcal{D}$[7], if not change h by h^{-1}. Then we let P_T be the point of A_g 'closest' to A_h^+, that is, the ray from any other point of A_g towards A_h^+ passes through P_T.

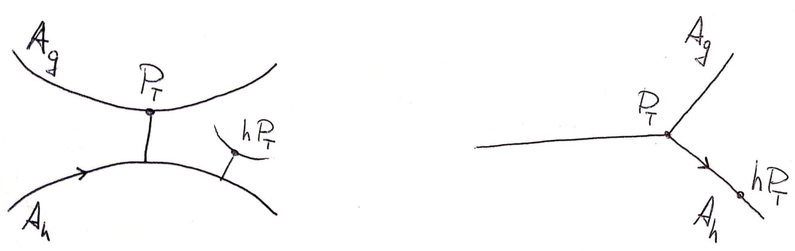

Figure 1: The two possible situations near P_T. The intersection between A_g and A_h may be compact.

But what do we mean exactly with 'P_T depends continuously on T'?

[7] A_h^+ being an end of A_g is an invariant of the deformation space: we can check that with the equivariant quasi-isometry definition of deformation space, because A_h^+ is an end of A_g iff for any point $x \in T$, the sequence $d_n = d(h^n x, gh^n x)$ does not tend to infinity.

Proposition 4.10. *Let $T \in \mathcal{D}$ and $\varepsilon > 0$. Then there is some nhood U of T such that, if $T' \in U$ via a map $X \to T'; x \mapsto x'$, with $P_T \in X$, then $d(P'_T, P_{T'}) < \varepsilon$.*

Proof. By definition, the point P_T satisfies the following equations:

- $d(P_T, gP_T) = l(g)$, because $P_T \in A_g$.
- $d(P_T, hP_T) = l(h) + 2d(A_g, A_h)$, because P_T is the closest point to A_h in A_g.
- $d(hP_T, ghP_T) = l(g) + 2l(h) + 4d(A_g, A_h)$, because hP_T is at distance $l(h) + 2d(A_g, A_h)$ of A_g.

Moreover, P_T is the only point which satisfies the three equations. In fact, suppose we have another point Q such that $|d(Q, gQ) - l(g)|, |d(Q, hQ) - l(h) - 2d(A_g, A_h)|$ and $|d(hQ, ghQ) - l(g) - 2l(h) - 4d(A_g, A_h)|$ are all $< \delta$. Then $d(Q, A_g) < \delta$, $d(Q, A_h) < d(A_g, A_h) + \frac{\delta}{2}$ and $d(hQ, A_g) > l(h) + 2d(A_g, A_h) - \frac{\delta}{2}$. Dividing in cases in the figures above, it is not hard to see that this implies $d(P_T, Q) < 3\delta$.

So, take a nhood U of T such that if $T' \in U$ via $X \to T'; x \mapsto x'$, $l(g), l(h), d(A_g, A_h)$ change by less than $\frac{\varepsilon}{7}$ from T to T', and $|d(P'_T, hP'_T) - d(P_T, hP_T)|, |d(P'_T, gP'_T) - d(P_T, gP_T)|, |d(hP'_T, ghP'_T) - d(hP_T, ghP_T)| < \frac{\varepsilon}{7}$. Then in T', $|d(P'_T, gP'_T) - l(g)|, |d(P'_T, hP'_T) - l(h) - 2d(A_g, A_h)|$ and $|d(hP'_T, ghP'_T) - l(h) - 2l(g) - 4d(A_g, A_h)|$ are all $< \varepsilon$, so P'_T is at distance $< 3\varepsilon$ of $P_{T'}$. □

Remark 4.11. The proof of proposition 4.10 gives us an example of what Guirardel and Levitt call in [GL06] a **stable** system of equations: a system of equations satisfied by a single point P and such that $\forall\varepsilon\exists\delta$ such that if some point Q satisfies those equations up to δ, then $d(P, Q) < \varepsilon$.

Remark 4.12. I wanted to express lemma 4.10 in a way that clearly conveys that P_T depends continuously on T, but in [GL06] Guirardel and Levitt do it differently. Their lemma 6.7, analogous to 4.10, reads as follows:

For any $a \in G$, the map $T \mapsto d(P_T, aP_T)$, from $\mathcal{D}$ to $\mathbb{R}$, is continuous.

Why does this imply that P_T depends continuously on T? The idea is that the system of equations $\mathrm{eq}_a = \{d(Q, aQ) = d(P_T, aP_T)\}$, for a in some appropriate finite subset $A \subseteq G$, is stable. Thus, if $d(P_T, aP_T)$ depends continuously on T and in lemma 4.10 we pick $X = \{P_T\}$ and the set A such that the system of equations is stable, then when ε tends to 0, $d(P'_T, P_{T'})$ will converge to 0.

4.4 Choice of f_T. Contractibility in the Gromov topology

This section defines the map $\mathcal{D} \to \mathcal{M}; T \mapsto f_T$, where f_T is a morphism from a tree $T_0(T) \in \mathcal{C}_0$ to T, from section 4.1. We define it as in [GL06] and prove its continuity, thus proving theorem 4.2.

Fix a tree $T_0 \in \mathcal{D}$ with finitely generated vertex stabilizers and let $\mathcal{C}_0$ be its closed cone in $\mathcal{D}$.

We define $f_T : T_0 \to T$ (we will later change the domain to $T_0(T)$) in the following way: first choose a set $\{v_i\}_i \in I$ of representatives of the orbits of vertices of T_0. Now for each v_i define $f_T(v_i)$ to be the projection of the point P_T in the fixed point set of G_{v_i}, where G_{v_i} is the stabilizer subgroup of v_i.

As defined, f_T may not be a morphism, but this can be solved changing the lengths of the edges of T_0. This is why we have to change the domain of f from T_0 to other tree $T_0(T)$ of its cone, so that the length of an edge between two vertices v, w in $T_0(T)$ is exactly $d(f_T(v), f_T(w))$. As the combinatorial structures of T_0 and $T_0(T)$ are the same, we will sometimes call their points/edges in the same way. Note that $T_0(T)$ is contained in the closed cone of T_0, and $T_0(T) \in \mathcal{D}$, because T_0 dominates $T_0(T)$ and $T_0(T)$ dominates T.

The last step to prove contractibility in the Gromov topology is to show that the function $T \mapsto f_T$ is continuous. As in [GL06], to do it we will use this lemma:

Lemma 4.13. *If for every pair of vertices v, w, the function $T \mapsto d(f_T(v), f_T(w))$ is continuous, then the function $T \mapsto f_T$ is continuous.*

Proof. First we will prove that if in the lemma we allow any two points p, q instead of just vertices, then the lemma is true: consider any nhood $W_{f_T}(X, A, \varepsilon)$ of f_T. Using the continuity of the functions $T \mapsto d(f_T(p), f_T(q))$ we can find a nhood U of T such that $\forall T' \in U$, $|d(f_T(x), f_T(ay)) - d(f_{T'}(x), f_{T'}(ay))| < \varepsilon$ for all $x, y \in X$, $a \in A$. As the length of an edge (v, w) in $T_0(T)$ is just $d(f_T(v), f_T(w))$, which depends continuously on T, we can shrink the neighborhood U so that if $T' \in U$, $|d(x, ay) - d(x', ay')| < \varepsilon \ \forall x, y \in X, a \in A$, thus $f_{T'} \in W_{f_T}(X, A, \varepsilon) \forall T' \in U$.

Now suppose that $T \mapsto d(f_T(v), f_T(w))$ is continuous for every pair of vertices v, w, and let's prove that $T \mapsto d(f_T(p), f_T(q))$ is continuous $\forall p, q \in T_0$. Let u be a vertex of T_0 and let $x \in T_0$ be in an edge between vertices v, w. We first prove that $d(f_T(x), f_T(u))$ depends continuously on T. To do it, consider the point $p = p_T \in T$ obtained from projecting $f(u)$ onto the segment $[f(v), f(w)] \subseteq T$. The situation is as shown in the figure:

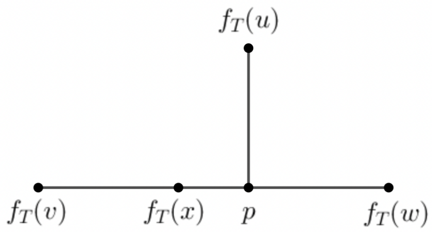

So, $d(f_T(x), f_T(u)) = d(p, f_T(u)) + d(f_T(x), p) = d(p, f_T(u)) + |d(f_T(v), p) - d(f_T(v), f_T(x))|$. These three functions depend continuously on T, because:

- $d(f_T(u), p) = d(f_T(u), f_T(w)) + d(f_T(u), f_T(v)) - d(f_T(v), f_T(w))$.
- $d(f_T(v), p) = d(f_T(v), f_T(u)) - d(f_T(u), p)$.

- $d(f_T(v), f_T(x)) = d(f_T(v), f_T(w)) \cdot \frac{d(v,x)}{d(v,w)}$, due to f_T being defined linearly in edges.

So the functions $d(f_T(x), f_T(u))$ are continuous for any $x \in T_0$ and any vertex u of T_0.

Now we can use the same strategy again: given points $p, q \in T_0$, we let v, w be the vertices of the edge of q, and we can show that $d(f_T(p), f_T(q))$ is continuous using that the functions $d(f_T(v), f_T(w))$, $d(f_T(p), f_T(v))$ and $d(f_T(p), f_T(w))$ are continuous. $\square$

So to prove contractibility in the Gromov topology, it is enough to prove that the functions $T \mapsto d(f_T(v), f_T(w))$ are continuous for any vertices v, w. That is a consequence of the following proposition:

Proposition 4.14. *Let v be a vertex of T_0 such that the stabilizer of v, G_v, is finitely generated. Then for every $T \in \mathcal{D}$ and $\varepsilon > 0$ there is a neighborhood U of T such that if $T' \in U$ via a map $X \to T'$; $x \mapsto x'$, with $f_T(v) \in X$, then $d(f_T(v)', f_{T'}(v)) < \varepsilon$.*

We will denote the fixed point set of a subset $H < G$ in a G-tree T by $\mathrm{Fix}_T(H)$.

Proof. Let $g_1, \dots, g_n$ be a set of generators of G_v and $\delta > 0$ small. Also let $D = d(P_T, \mathrm{Fix}_T(G_v)) = d(P_T, f_T(v))$ and let Q be the projection of $f_T(v)'$ onto $\mathrm{Fix}_{T'}(G_v)$.

Take a neighborhood U of T such that if $T' \in U$ via a map $X \to T'$; $x \mapsto x'$, then:

1. $d(f_T(v)', g_i f_T(v)') < \delta$ $\forall i$. This condition implies that $d(f_T(v)', Q) = d(f_T(v)', \mathrm{Fix}_{T'}(G_v)) < \frac{\delta}{2}$.
2. $|d(P'_T, f_T(v)') - d(P_T, f_T(v))| < \delta$.
3. $|d(P_T, g_i P_T) - d(P'_T, g_i P'_T)| < \delta$. This implies that $d(P'_T, \mathrm{Fix}_{T'}(G_v)) > D - \frac{\delta}{2}$
4. $d(P'_T, P_{T'}) < \delta$ (this last part is possible thanks to 4.10).

Then $d(P_{T'}, Q) < d(P_{T'}, P'_T) + d(P'_T, f_T(v)') + d(f_T(v)', Q) < \delta + (d(P_T, f_T(v)) + \delta) + \frac{\delta}{2} = D + \frac{5\delta}{2}$. However Q is in $\mathrm{Fix}_{T'}(G_v)$, so $f_{T'}(v)$ is inside the segment $[P_{T'}, Q]$. Moreover, $d(P_{T'}, f_{T'}(v)) = d(P_{T'}, \mathrm{Fix}_{T'}(G_v)) > d(P'_T, \mathrm{Fix}_{T'}(G_v)) - d(P'_T, P_{T'}) > D - \frac{3\delta}{2}$.

So $d(f_{T'}(v), Q) = d(P_{T'}, Q) - d(P_{T'}, f_{T'}(v)) < D + \frac{5\delta}{2} - (D - \frac{3\delta}{2}) = 4\delta$, thus $d(f_{T'}(v) - f_T(v)') \leq d(f_{T'}(v), Q) + d(Q, f_T(v)') < 4\delta + \frac{\delta}{2}$. Setting $\delta = \frac{\varepsilon}{4}$, we are done.

$\square$

4.5 Contractibility in the weak topology

This section proves theorem 4.1, following section 6 of [GL06]. The homotopy $H : \mathcal{D} \times [0,1] \to \mathcal{D}$ is defined the same way as in previous sections, but now the tree T_0 need not have finitely generated vertex stabilizers.

To prove continuity of H, we just need to prove continuity of $H|_{\mathcal{C} \times [0,\infty]}$, the restriction of H to any closed cone $\mathcal{C}$ of $\mathcal{D}$. We will use the following lemma:

Lemma 4.15. *For any closed cone $\mathcal{C} \subseteq \mathcal{D}$, the image of $H|_{\mathcal{C}\times[0,\infty]}$ is contained in a finite union of cones of $\mathcal{D}$.*

As with 4.9, this is a technical lemma about Skora's deformation with a tedious proof. Its proof can be found in lemmas 6.5 and 6.4 of [GL06].

Proof of theorem 4.2. Due to lemma 4.15 and 2.10, it is enough to prove that $H|_{\mathcal{C}\times[0,\infty]}$ is continuous in the Gromov topology for any closed cone $\mathcal{C}$. When we proved that H is continuous in the Gromov topology, the only placed where we used that T_0 has finitely generated vertex stabilizers is in order to prove that for every two vertices v, w of T_0, the function $T \mapsto d(f_T(v), f_T(w))$ is continuous (the hypothesis of lemma 4.13).

However for any tree T in a given cone $\mathcal{C}$, $d(f_T(v), f_T(w))$ depends linearly on the lengths of the edges of T, so $T \mapsto d(f_T(v), f_T(w))$ is continuous. □

Proof of theorem 2.9. It is enough to check that every tree of $\mathcal{D}$ is related to T_0 by a sequence of contractions and retractions. Let $T \in \mathcal{D}$ and let $\mathcal{S}$ be the finite set of cones which intersect $H(\{T\} \times [0,\infty])$. Then as H is continuous in the weak topology, $\mathcal{S}$ is a finite union of cones which is connected in the weak topology. Thus there is a sequence of closed cones $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_n$ in $\mathcal{D}$ from $\mathcal{C}_0$ to the cone of T such that $\mathcal{C}_{i+1}$ either contains or is contained in $\mathcal{C}_i$. Thus T and T_0 are related by a sequence of elementary contractions and expansions. □

5 Consequences for Out($\mathbb{F}_n$)

Having proved at last that the Outer Space is contractible, let's see what that implies about the group Out($\mathbb{F}_n$). Some reasonings from this section can be applied much more generally for groups acting on cell complexes, so it is a good example of how to deduce information about a group from its actions. This is the main theorem of the section:

Theorem 5.1. *Out($\mathbb{F}_n$) is of type WFL and has virtual cohomological dimension $2n-3$.*

To understand what this means and prove it we will need to introduce some concepts from chapter 8 of [Br82]:

Definition 5.2. Let G be a group. The *cohomological dimension* of G, cd(G), is the projective dimension of $\mathbb{Z}$ as a module over $\mathbb{Z}G$.

This means that cd(G) is the minimum i such that you have a projective resolution of length i for $\mathbb{Z}$ over $\mathbb{Z}G$,

$$0 \to P_i \to \cdots \to P_0 \to \mathbb{Z} \to 0,$$

or ∞ if such i doesn't exist.

We also introduce a condition stronger than finite cohomological dimension:

Definition 5.3. We say G is of type FL if $\mathbb{Z}$ admits a finite free resolution over $\mathbb{Z}G$.

As explained in [Br82], this concepts can be easily related to the cell cohomology of G**-complexes**, that is, CW complexes with an action of G which permutes their cells. Indeed, the cellular chain complex of a G-complex X of dimension k is a sequence

$$0 \to C_k^{\text{cell}}(X) \cdots \to C_1^{\text{cell}}(X) \to C_0^{\text{cell}}(X) \to \mathbb{Z} \to 0.$$

If G acts freely on the set of cells, the groups $C_i^{\text{cell}}(X)$ are free $\mathbb{Z}G$-modules. If X is contractible, then the sequence is exact, so it is a free resolution of $\mathbb{Z}$, meaning that G is of type FL and $\text{cd}(G) \leq k$.

When we try to apply this to the action of $\text{Out}(\mathbb{F}_n)$ on the spine K_n of the Outer Space, we run into a problem: the action of $\text{Out}(\mathbb{F}_n)$ on K_n is not free. In fact, $\text{Out}(\mathbb{F}_n)$ does not have finite cohomological dimension nor is it of type FL. This failure is caused by the fact that $\text{Out}(\mathbb{F}_n)$ has non trivial torsion:

Proposition 5.4. *If a group G has non trivial torsion, then cd$(G) = \infty$, thus G is not of type FL.*

Proof. First of all, a projective resolution over $\mathbb{Z}G$ is also a projective resolution over $\mathbb{Z}H$, for any subgroup H of G. So we only need to prove this for $G = \mathbb{Z}_n$. In that case, let t be a generator of G and consider the projective resolution

$$\cdots \mathbb{Z}G \overset{t-1}{\to} \mathbb{Z}G \overset{N}{\to} \mathbb{Z}G \overset{t-1}{\to} \mathbb{Z}G \to \mathbb{Z} \to 0,$$

where $N = \sum_{g \in G} g$, for which infinitely many cohomology groups are non trivial. $\square$

But it is 'almost' free in the sense that it has finite vertex stabilizers. This will allow us to deduce that $\text{Out}(\mathbb{F}_n)$ has two slightly weaker versions of the properties above:

Definition 5.5. We say G has *virtual cohomological dimension* n, $\text{vcd}(G) = n$, if there is some finite index subgroup $H < G$ with $\text{cd}(H) = n$. We say G is of type WFL if it is virtually torsion free and all its torsion free finite index subgroups are of type FL.

Virtual cohomological dimension is well defined because if any finite index subgroup of G has virtual cohomological dimension n, then all torsion free finite index subgroups of G have that same dimension. This can be directly deduced from Theorem VIII.3.1 of [Br82].

To show that $\mathrm{vcd}(\mathrm{Out}(\mathbb{F}_n)) \leq 2n-3$ and $\mathrm{Out}(\mathbb{F}_n)$ is of type WFL, it will be enough to prove that $\mathrm{Out}(\mathbb{F}_n)$ has finite index torsion free subgroups. Indeed, if H is such a subgroup, then its action on cells of K_n will be free, because K_n has finite cell stabilizers (3.11). So the chain complex of K_n (which has dimension $2n-3$ by 3.10) would give a free resolution of $\mathbb{Z}$ over $\mathbb{Z}H$ of length $2n-3$, which proves H is FL and $\mathrm{cd}(H) \leq 2n-3$.

Vogtmann proves the existence of torsion free finite index subgroups of $\mathrm{Out}(\mathbb{F}_n)$ in [Vo02] as follows: consider the natural map from $\mathrm{Aut}(\mathbb{F}_n)$ to $\mathrm{Aut}(\mathbb{Z}^n) \cong \mathrm{GL}_n(\mathbb{Z})$, which descends to a map from $\mathrm{Out}(\mathbb{F}_n)$ to $\mathrm{GL}_n(\mathbb{Z})$. Baumslag and Taylor proved in [BT68] that the kernel of this map is torsion free, so the inverse image of any torsion free finite index subgroup of $\mathrm{GL}_n(\mathbb{Z})$ by this map will be torsion free and of finite index in $\mathrm{Out}(\mathbb{F}_n)$. To prove that there are torsion free finite index subgroups of $\mathrm{GL}_n(\mathbb{Z})$ we can apply the same strategy again, proving that the kernel of the natural map $\mathrm{GL}_n(\mathbb{Z}) \to \mathrm{GL}_n(\mathbb{Z}_3)$ is torsion free, which is an easy exercise in number theory.

So the only thing left to prove is $\mathrm{vcd}(\mathrm{Out}(\mathbb{F}_n)) \geq 2n-3$. We will need three facts:

- If H is a subgroup of G, then $cd(H) \leq cd(G)$.

  This follows from the fact that any projective resolution over $\mathbb{Z}G$ is also a projective resolution over $\mathbb{Z}H$.

- $\mathrm{Out}(\mathbb{F}_n)$ contains subgroups isomorphic to $\mathbb{Z}^{2n-3}$.

  One such subgroup is the one generated by the automorphisms a_i and b_i, $i = 2, \dots n$, defined as follows: let $x_1, \dots, x_n$ be a base of $\mathbb{F}_n$, and let $a_i(x_j) = b_i(x_j) = x_j$ if $j \neq i$, with $a_i(x_i) = x_i x_1$ and $b_i(x_i) = x_1 x_i$.

- $\mathrm{cd}(\mathbb{Z}^k) = k$ for all k.

  This can be deduced from the following theorem (VIII.7.5 from [Br82]) applied to the additive action of $\mathbb{Z}^k$ on $\mathbb{R}^k$:

  Theorem 5.6. *If X is a contractible, free G-complex with compact quotient X/G, then there is an isomorphism*
  $$H^*(G, \mathbb{Z}G) \cong H^*_c(X, \mathbb{Z})$$

With that in mind, consider a torsion free finite index subgroup H of $\mathrm{Out}(\mathbb{F}_n)$. Then H contains a subgroup isomorphic to $\mathbb{Z}^{2n-3}$ (namely, its intersection with a subgroup of $\mathrm{Out}(\mathbb{F}_n)$ isomorphic to $\mathbb{Z}^{2n-3}$), so $\mathrm{cd}(H) \geq 2n-3$ and we are done.

References

[BT68] G. Baumslag and T. Taylor. *The centre of groups with one defining relator.* Math Ann. 175 (1968), 315-319.

[Br82] Kenneth S. Brown. *Cohomology of Groups.* Springer, 1982.

[CM87] M. Culler, J.W. Morgan. *Group actions on* $\mathbb{R}$*-trees.* Proc. London Math. Soc. (3), Volume 55 (1987) no. 3, pp. 571-604.

[CV86] Marc Culler and Karen Vogtmann. *Moduli of graphs and automorphisms of free groups.* Invent. Math., 84(1):91–119, 1986.

[GL06] V. Guirardel, G. Levitt. *Deformation Spaces of Trees.* Groups, Geometry, and Dynamics 1 (2006): 135-181.

[GL07] V. Guirardel, G. Levitt. *The outer space of a free product.* Proc. London Math. Soc. Volume 94, Issue 3, May 2007, pp. 695–714.

[Ha00] Allen Hatcher. *Algebraic topology.* Cambridge Univ. Press, Cambridge, 2000.

[Ni17] J. Nielsen. *Die Isomorphismen der allgemeinen, unendlicher Gruppe mit zwei Erzeugenden.* Math. Ann. 78 (1917), 385-397.

[Pa89] F. Paulin. *The Gromov Topology on* $\mathbb{R}$*-trees.* Topology Appl. 32 (1989) 197-221.

[Vo02] Karen Vogtmann. *Automorphisms of Free Groups and Outer Space.* Geometriae Dedicata vol. 94, pg. 1–31 (2002)