

Contents

PCMI Lectures on the geometry of Outer space	
MLADEN BESTVINA	1
PCMI Lectures on the geometry of Outer space	3
Introduction	3
Lecture 1. Outer space and its topology	5
Lecture 2. Lipschitz metric, Train tracks	13
Lecture 3. Classification of automorphisms	21
Lecture 4. Hyperbolic features	27
Bibliography	31

PCMI Lectures on the geometry of
Outer space

Mladen Bestvina

PCMI Lectures on the geometry of Outer space

Mladen Bestvina

Introduction

Outer space is a contractible space on which $Out(\mathbb{F}_n)$ acts properly discontinuously. It was introduced by Marc Culler and Karen Vogtmann in [18]. Outer space is analogous to the symmetric space associated to an arithmetic lattice or the Teichmüller space associated to the mapping class group of a surface.

It is useful to keep in mind the comparison maps $Out(\mathbb{F}_n) \rightarrow GL_n(\mathbb{Z})$ and $Mod(S) \rightarrow Out(\pi_1(S))$. The first is obtained by sending an automorphism of \mathbb{F}_n to the induced automorphism of the abelianization \mathbb{Z}^n of \mathbb{F}_n . It is always surjective, and for $n = 2$ it is an isomorphism. The second comparison homomorphism is defined on the (extended) mapping class group of a punctured surface S with $\chi(S) < 0$ and it is always injective. When S is a punctured torus it is an isomorphism.

These notes are about the geometry of Outer space. The natural metric on it is not symmetric, but this is imposed by the features of $Out(\mathbb{F}_n)$; e.g. the growth rates of an automorphism and its inverse may be different. We give a metric classification of automorphisms, following Bers' proof of Thurston's classification of surface automorphisms. For the most part, the exposition is hands-on, with many exercises involving concrete examples. There are several harder exercises, indicated by asterisks. The last lecture gives a glimpse of the current developments.

I would like to thank Yael Algom-Kfir, Pritam Ghosh, Brian Mann and Catherine Pfaff for correcting some of the mistakes in an earlier version, and to the organizers of minicourses where I presented some variant of these notes at the following locations: University of Chicago, University of Utah, Goa, Technion, University of Buffalo, Yale University, Berlin and Vercors. Special thanks go to Yael Algom-Kfir and Catherine Pfaff for carefully reading the manuscript and suggesting several improvements.

Outer space and its topology

In this lecture we will talk about the topology of Outer space. For more information see the excellent survey [37].

A *graph* is a cell complex of dimension ≤ 1 . The *rose* R_n is the graph with 1 vertex and n edges.

1.1. Markings

A *marking* of a graph Γ is a homotopy equivalence $f : R_n \rightarrow \Gamma$. This is a convenient way of specifying an identification between $\pi_1(\Gamma)$ with the free group \mathbb{F}_n (thought of as being identified with $\pi_1(R_n)$ once and for all) with a (deliberate) ambiguity of composing with inner automorphisms (no basepoints!). Two marked graphs $f : R_n \rightarrow \Gamma$ and $f' : R_n \rightarrow \Gamma'$ are *equivalent* if there is a homeomorphism $\phi : \Gamma \rightarrow \Gamma'$ such that $\phi f \simeq f'$ (homotopic).

In practice one defines the *inverse* of a marking, i.e. a homotopy equivalence $\Gamma \rightarrow R_n$. If the edges of R_n are oriented and labeled by a basis a, b, \dots of \mathbb{F}_n (thus identifying $\pi_1(R_n) = \mathbb{F}_n$), the inverse marking can be defined by specifying a maximal tree T in Γ , orienting all edges in $\Gamma - T$, and labeling them with a (possibly different) basis of \mathbb{F}_n , expressed as words in a, b, \dots . Such a choice defines a map $\Gamma \rightarrow R_n$ by collapsing T to a point and sending each edge to the edge path specified by the label.

Exercise 1.1. Show that the two marked graphs in Figure 1.1 are equivalent. We follow the convention that capital letters represent inverses of lower case letters. Unlabeled edges form a maximal tree.




FIGURE 1.1. Equivalent marked graphs

1.2. Metric

A *metric* on a finite graph is an assignment ℓ of positive numbers $\ell(e)$, called *lengths*, to the edges e of Γ . The *volume* of a finite metric graph is the sum of the lengths of the edges.

A metric on a graph allows one to view the graph as a geodesic metric space, with each edge e having length $\ell(e)$. This point of view lets us assign lengths also to paths in the graph; in particular any closed immersed loop has finite length (an *immersion* is a locally injective map, in this case from the circle to the graph).

We will consider the triples (Γ, ℓ, f) where Γ is a finite graph with all vertices of valence ≥ 3 , ℓ is a metric on Γ with volume 1, and $f : R_n \rightarrow \Gamma$ is a marking. Two such triples (Γ, ℓ, f) and (Γ', ℓ', f') are *equivalent* if there is an *isometry* (i.e. a length-preserving homeomorphism) $\phi : \Gamma \rightarrow \Gamma'$ such that $\phi f \simeq f'$.

Definition 1.2. Outer space

$$\mathcal{X}_n = \{(\Gamma, \ell, f)\} / \sim$$

is the set of equivalence classes of finite marked metric graphs with vertices of valence ≥ 3 and of volume 1.

We will usually omit equivalence class, ℓ and f from the notation, and talk about points $\Gamma \in \mathcal{X}_n$ instead of $[(\Gamma, \ell, f)] \in \mathcal{X}_n$.

1.3. Lengths of loops

Once $\pi_1(R_n)$ is identified with \mathbb{F}_n we can view each nontrivial conjugacy class in \mathbb{F}_n as a loop in R_n , up to homotopy. The homotopy class has a unique immersed representative, up to parametrization. If α is a nontrivial conjugacy class and $(\Gamma, \ell, f) \in \mathcal{X}_n$, define the length $\ell_\Gamma(\alpha)$ of α in Γ as the length of the immersed loop $\alpha|_\Gamma$ in Γ homotopic to $f(\alpha)$.

1.4. \mathbb{F}_n -trees

If Γ is a marked metric graph, the universal cover $\tilde{\Gamma}$ is a (metric, simplicial) tree, and the marking (i.e. the identification $\pi_1(\Gamma) = \mathbb{F}_n$) induces an action of \mathbb{F}_n on $\tilde{\Gamma}$. The equivalence relation on marked metric graphs translates to saying that two metric simplicial \mathbb{F}_n -trees S, T are equivalent if there is an equivariant isometry $S \rightarrow T$. Thus \mathcal{X}_n can be alternatively defined as the space of minimal metric simplicial free \mathbb{F}_n -trees with covolume 1, up to equivariant isometry. The length of a conjugacy class becomes the translation length in the tree.

1.5. Topology and Action

\mathcal{X}_n can be naturally decomposed into open simplices. If Γ is a graph and $f : R_n \rightarrow \Gamma$ a marking, the set of possible metrics $\overset{\circ}{\Sigma}(\Gamma)$ on Γ is an open simplex

$$\{(\ell_1, \ell_2, \dots, \ell_E) \mid \ell_i > 0, \sum \ell_i = 1\}$$

of dimension $E - 1$ if E is the number of edges. If T is a forest (i.e. a disjoint union of trees) in Γ and $\Gamma' = \Gamma/T$ is obtained by collapsing all edges of T to points, then $\overset{\circ}{\Sigma}(\Gamma')$ can be identified with the open face of $\overset{\circ}{\Sigma}(\Gamma)$ in which the coordinates of edges in T are 0. Then Γ' is said to be obtained from Γ by *collapsing a forest*, and Γ is obtained from Γ' by *blowing up a forest*. The union $\Sigma(\Gamma)$ of $\overset{\circ}{\Sigma}(\Gamma)$ with all such open faces as T ranges over all forests in Γ is a simplex-with-missing-faces: it can be obtained from the closed simplex $\Sigma^*(\Gamma) = \{(\ell_1, \ell_2, \dots, \ell_E) \mid \ell_i \geq 0, \sum \ell_i = 1\}$ by deleting those open faces that assign 0 to a set of edges whose union contains a loop. For example, if Γ is the theta-graph with 2 vertices and 3 edges connecting them, $\Sigma(\Gamma)$ is the 2-simplex minus its vertices.

Exercise 1.3. The smallest dimension of a $\Sigma(\Gamma)$ is $n - 1$.

Exercise 1.4. The largest dimension of a $\Sigma(\Gamma)$ is $3n - 4$.

Top dimensional simplices correspond to 3-valent graphs, codimension 1 simplices to graphs with one valence 4 vertex and all others valence 3, etc.

In this way \mathcal{X}_n becomes a complex of simplices-with-missing-faces. We define the *simplicial topology* on \mathcal{X}_n just like on a simplicial complex: a subset $U \subset \mathcal{X}_n$ is open [closed] if and only if $U \cap \Sigma(\Gamma)$ is open [closed] in $\Sigma(\Gamma)$ for every $\Gamma \in \mathcal{X}_n$.

One can also put the missing faces in and get a simplicial complex \mathcal{X}_n^* , called the *simplicial completion* of \mathcal{X}_n . This complex is isomorphic to the *splitting complex* or equivalently to the *complex of spheres*, see Lecture 4. The fact that \mathcal{X}_n^* is a simplicial complex (e.g. simplices are determined by their vertices) is nontrivial.

Exercise 1.5. Show that every $\Sigma(\Gamma)$ is contained in only finitely many $\Sigma(\Gamma')$. Conclude that \mathcal{X}_n is locally compact and metrizable.

Another way to define a topology on \mathcal{X}_n is via length functions. Let \mathcal{C} be the set of all nontrivial conjugacy classes in \mathbb{F}_n . The *length function* is the function

$$\mathcal{L} : \mathcal{X}_n \rightarrow (0, \infty)^{\mathcal{C}}$$

to the space of functions $\mathcal{C} \rightarrow (0, \infty)$ with the product topology that to Γ assigns $(\alpha \mapsto \ell_\Gamma(\alpha))$. This function is injective (this is the Rigidity of the Length Spectrum, see Exercise 2.8), and if we identify \mathcal{X}_n with the image, the subspace topology induces a topology on \mathcal{X}_n , the *length function topology*. This topology is equivalent to the simplicial topology, i.e. \mathcal{L} is an embedding, see [22]. See also [17].

1.6. Thick part and spine

For a fixed small $\epsilon > 0$ define the *thick part* $\mathcal{X}_n(\epsilon)$ of \mathcal{X}_n as the set of $\Gamma \in \mathcal{X}_n$ such that $\ell_\Gamma(\alpha) \geq \epsilon$ for every nontrivial conjugacy class α . When $\epsilon > 0$ is sufficiently small the intersection of $\mathcal{X}_n(\epsilon)$ with every $\Sigma(\Gamma)$ is a nonempty convex set (e.g. taking $\epsilon \leq \frac{1}{3n-3}$ ensures that the barycenter of $\Sigma(\Gamma)$ is in $\mathcal{X}_n(\epsilon)$).

For each simplex-with-missing-faces $\Sigma(\Gamma)$ let $S(\Gamma)$ be the union of simplices in the barycentric subdivision of the closed simplex $\Sigma^*(\Gamma)$ that are contained in $\Sigma(\Gamma)$. Thus $S(\Gamma)$ is the dual of the missing faces. The *spine* $K_n \subset \mathcal{X}_n$ is the union of $S(\Gamma)$'s for all $\Gamma \in \mathcal{X}_n$.




FIGURE 1.2. Spine and thick part intersected with a simplex

1.7. Action of $Out(\mathbb{F}_n)$

There is a natural right action of $Out(\mathbb{F}_n)$ on \mathcal{X}_n by precomposing the marking. An element $\Phi \in Out(\mathbb{F}_n)$ can be thought of as a homotopy equivalence $\Phi : R_n \rightarrow R_n$ and then the action is:

$$[(\Gamma, \ell, f)] \cdot \Phi = [(\Gamma, \ell, f\Phi)]$$

Exercise 1.6. Show that the action is well-defined, i.e. does not depend on the representative in the conjugacy class.

Note that in terms of inverse markings, the action amounts to *postcomposing* with inverse automorphism. The action is simplicial and it is compatible with the action on conjugacy classes:

$$\ell_{\Gamma\Phi}(\alpha) = \ell_{\Gamma}(\Phi(\alpha))$$

It is sometimes convenient to write $\ell_{\Gamma}(\alpha)$ as a pairing $\langle \Gamma, \alpha \rangle$ and then the identity becomes

$$\langle \Gamma\Phi, \alpha \rangle = \langle \Gamma, \Phi(\alpha) \rangle$$

Exercise 1.7. Show that the point stabilizer $Stab(\Gamma, \ell, f)$ is isomorphic to the isometry group $Isom(\Gamma, \ell)$ of the underlying graph, with an isometry ϕ corresponding to the automorphism $f^{-1}\phi f$, where $f^{-1} : \Gamma \rightarrow R_n$ denotes the inverse marking.

Exercise 1.8. Show that there are only finitely many orbits of $\Sigma(\Gamma)$'s.

Exercise 1.9. Show that the action leaves the spine and the thick part invariant.

Proposition 1.10. *The action of $Out(\mathbb{F}_n)$ on \mathcal{X}_n is proper. The action on the thick part and on the spine is cocompact.*

Exercise 1.11. (Combinatorial description of the spine.) Show that the following simplicial complex \mathcal{P}_n , the *poset of marked graphs* is homeomorphic to the spine K_n . The vertices of \mathcal{P}_n are marked graphs (Γ, f) (with Γ having no vertices of valence ≤ 2) modulo equivalence $(\Gamma, f) \sim (\Gamma', f')$ if there is a homeomorphism $\phi : \Gamma \rightarrow \Gamma'$ with $\phi f \simeq f'$. A k -simplex in \mathcal{P}_n is induced by a sequence of nontrivial forest collapses $\Gamma_0 \rightarrow \Gamma_1 \rightarrow \dots \rightarrow \Gamma_k$.

Exercise 1.12. $\dim K_n = 2n - 3$.

Exercise 1.13. There are equivariant deformation retractions from Outer space \mathcal{X}_n to the thick part $\mathcal{X}_n(\epsilon)$ (for small $\epsilon > 0$) and from $\mathcal{X}_n(\epsilon)$ to the spine K_n .

Reduced Outer space \mathcal{R}_n is the subspace of \mathcal{X}_n consisting of those graphs that do not have a separating edge.

Exercise 1.14. Show that \mathcal{R}_n is an equivariant deformation retract of \mathcal{X}_n .

1.8. Rank 2 picture

Recall that the comparison map $Out(\mathbb{F}_n) \rightarrow GL_n(\mathbb{Z})$ is an isomorphism when $n = 2$ and that the (full) mapping class group of a punctured torus is $Out(\mathbb{F}_n)$. Since the symmetric space $SL_2(\mathbb{R})/SO_2$ and Teichmüller space of $(T^2, \{p\})$ is the hyperbolic plane \mathbb{H}^2 , it is not surprising that \mathcal{X}_2 is essentially also (a combinatorial version of) \mathbb{H}^2 . More precisely, the *reduced* Outer space in rank 2 is the *filled in Farey graph minus the vertices* pictured in Figure 1.3.




FIGURE 1.3. Reduced Outer space in rank 2. The circle and the vertices of the triangles are not part of the space.

Markings of three of the simplices are pictured in Figure 1.4. Observe that there are two ways to blow up a rose R_2 to a theta graph and this translates into the fact that reduced Outer space is a surface.




FIGURE 1.4. 3 2-simplices with their marked graphs

To obtain the whole Outer space, we also need to attach simplices corresponding to graphs with separating edges, see Figure 1.5. These simplices have missing vertices and are missing two of the sides. They are attached to the reduced Outer space along the third side.

Exercise 1.15. Find an automorphism of \mathbb{F}_2 that takes the bottom triangle in Figure 1.4 to the upper right triangle.




FIGURE 1.5. Simplices corresponding to graphs with separating edges

Exercise 1.16. What does the automorphism $a \mapsto a$, $b \mapsto ab$ do to the reduced Outer space? It fixes a missing vertex and ...

Exercise 1.17. Show that neither Outer space nor reduced Outer space is a manifold when $n \geq 3$. Hint: Find a graph in \mathcal{X}_n with one vertex of valence 4, the other vertices of valence 3, and all three blowups having no separating edges.

1.9. Contractibility

The central fact about Outer space is its contractibility, proved by Marc Culler and Karen Vogtmann.

Theorem 1.18. [18] *Outer space \mathcal{X}_n is contractible.*

Of course, this means that the thick part and the spine are also contractible.

Culler-Vogtmann use combinatorial Morse theory and argue that the spine is contractible. They carefully order the set of roses in \mathcal{X}_n : r_1, r_2, \dots and argue that for each i the union of stars of the first i roses is contractible. The difficult step is showing that the intersection of the star of the i th rose with the union of the previous stars is contractible.

An alternative proof, more in the spirit of these notes, was constructed by Skora [33], building on the ideas of Steiner. For each $\Gamma \in \mathcal{X}_n$ they construct a (folding) path to Γ from a point in the simplex containing the rose with identity marking and argue that the collection of these paths varies continuously in Γ (this is technically the hard step). These paths then determine a deformation retraction from \mathcal{X}_n to a simplex with missing faces. For more on folding paths see Lecture 2.

Neither Steiner's nor Skora's work was published; for details see [15] or [22].

1.10. Group theoretic consequences

Corollary 1.19. *$Out(\mathbb{F}_n)$ is finitely presented.*

PROOF. Recall that if a group acts freely and cocompactly on a simply connected simplicial complex, then it is finitely presented. More generally, it is finitely presented if it acts cocompactly on a simply connected complex with finitely presented vertex stabilizers and finitely generated edge stabilizers (see [14]). The action on the spine has finite stabilizers. \square

Proposition 1.20. *$Out(\mathbb{F}_n)$ is virtually torsion-free.*

In the proof we will use the fact that every finite subgroup of $Out(\mathbb{F}_n)$ fixes a point of \mathcal{X}_n . This is called the (Nielsen) realization theorem, see [39, 16, 29]. For a recent, more intrinsic, proof see [28].

PROOF. We claim that the kernel of $Out(\mathbb{F}_n) \rightarrow GL_n(\mathbb{Z}/3)$ is torsion-free. Let $1 \neq \Phi \in Out(\mathbb{F}_n)$ have finite order. By Nielsen realization, Φ is realized as a graph isomorphism $\phi : \Gamma \rightarrow \Gamma$. We may without loss collapse all separating edges of Γ , so every edge is contained in an embedded circle. If Φ is in the kernel, then ϕ maps any circle to itself preserving orientation. But for each circle C in Γ there is another circle C' (because $n \geq 2$ without loss) so that $C \cap C'$ is nonempty, connected, and $\neq C$. Thus ϕ is identity. \square

Corollary 1.21. *A torsion-free subgroup H of finite index has a compact classifying space $K(H, 1)$ of dimension $2n - 3$, and the virtual cohomological dimension of $Out(\mathbb{F}_n)$ is $2n - 3$.*

To see that $vcd(Out(\mathbb{F}_n)) \geq 2n - 3$ note that $Out(\mathbb{F}_n)$ contains an abelian subgroup of rank $2n - 3$. E.g. for $n = 3$ we can take the group of automorphisms of the form $a \mapsto a, b \mapsto a^p b, c \mapsto a^q c a^r$ for $p, q, r \in \mathbb{Z}$.

Exercise 1.22. Show that this group is indeed isomorphic to \mathbb{Z}^{2n-3} .

The next statement does not use contractibility of \mathcal{X}_n , only Nielsen realization plus the fact that the action on the spine is cocompact.

Corollary 1.23. *$Out(\mathbb{F}_n)$ has finitely many conjugacy classes of finite subgroups.*

Exercise* 1.24. Find a nontrivial element of finite order in the kernel of $Out(\mathbb{F}_n) \rightarrow GL_n(\mathbb{Z}/2)$. Show that every such element has order 2 and that therefore every finite subgroup of the kernel is abelian (in fact, a direct sum of $\mathbb{Z}/2$'s). Can you find the largest such subgroup?

Exercise* 1.25. Can you find estimates on the size of the largest finite subgroup of $Out(\mathbb{F}_n)$? For example, the stabilizer of a rose has order $2^n n!$. Can you find a larger finite group? What about $n = 2$ and 3? For the answer see [38].

LECTURE 2

Lipschitz metric, Train tracks

In this lecture we introduce the Lipschitz metric on Outer space. The definition, motivated by Thurston's metric on Teichmüller space [36], dates back to the 1990's when my former student Tad White proved the key Lemma 2.4. At the time we didn't have any applications for this metric. It recaptured my own interest when I realized that one can give a classification of automorphisms in the style of Bers using this metric. Bers [4] proved the Thurston classification theorem for mapping classes using the Teichmüller metric on Teichmüller space. This is the subject of Lecture 3.

Francaviglia and Martino were the first to study this metric systematically. Much of the material in this section is in their paper [21].

2.1. Definitions

Let $[(\Gamma, \ell, f)], [(\Gamma', \ell', f')] \in \mathcal{X}_n$ be two points in Outer space. A continuous map $\phi : \Gamma \rightarrow \Gamma'$ is a *difference of markings* map if $\phi f \simeq f'$. We will only consider Lipschitz maps and we denote by $\sigma(\phi)$ the Lipschitz constant of ϕ . When ϕ is homotoped rel vertices to a map ϕ' which has constant slope on each edge, then $\sigma(\phi') \leq \sigma(\phi)$. We define the distance:

$$d(\Gamma, \Gamma') = \inf_{\phi} \log \sigma(\phi)$$

as $\phi : \Gamma \rightarrow \Gamma'$ ranges over all difference of markings. Recall the Arzela-Ascoli theorem, which says that any sequence of L -Lipschitz maps between two compact metric spaces has a convergent subsequence. This theorem implies that the infimum above is realized. We will call a difference of markings $\phi : \Gamma \rightarrow \Gamma'$ *optimal* if it has constant slope on each edge and minimizes the Lipschitz constant (which is then the maximal slope).

2.2. Elementary facts

Proposition 2.1.

- $d(\Gamma_1, \Gamma_3) \leq d(\Gamma_1, \Gamma_2) + d(\Gamma_2, \Gamma_3)$.
- $d(\Gamma, \Gamma') \geq 0$ and equality implies $\Gamma = \Gamma'$.
- $d(\Gamma\Phi, \Gamma'\Phi) = d(\Gamma, \Gamma')$.

PROOF. The first claim follows from the general fact that $\sigma(\psi\phi) \leq \sigma(\psi)\sigma(\phi)$. For the second claim, let $\phi : \Gamma \rightarrow \Gamma'$ be an optimal map. If $d(\Gamma, \Gamma') < 0$ then all slopes of ϕ are < 1 . This implies that the volume of the image of ϕ is < 1 , so ϕ is not surjective. But a homotopy equivalence between finite graphs without vertices of valence 1 is always surjective.

If $d(\Gamma, \Gamma') = 0$ then all slopes of ϕ must be equal to 1 and the images of different edges can intersect only in finite sets. Thus ϕ is a quotient map that

identifies finitely many collections of finitely many points. The only way for such a map to be a homotopy equivalence (or even for Γ and Γ' to have the same rank) is for ϕ to be an isometry, so $\Gamma = \Gamma'$.

The third claim is an exercise. \square

2.3. Example

To illustrate the definition, let us compute the distance in the following example, see Figure 2.1.




FIGURE 2.1. A is the rose with edge lengths $\frac{1}{2}$ and B is the theta graph with edge lengths $\frac{1}{3}$, both in the same 2-simplex.

To compute $d(A, B)$ consider the difference of markings map ϕ that sends the vertex of A to the midpoint of the middle edge of B , the loop on the left homeomorphically to the circle formed by the middle and the left edges, and the loop on the right homeomorphically to the circle formed by the middle and the right edge. The slope of ϕ on both edges is $\frac{(\frac{2}{3})}{(\frac{1}{2})} = \frac{4}{3}$, so $d(A, B) \leq \log \frac{4}{3}$. We now claim that $d(A, B) = \log \frac{4}{3}$. To see this, observe that each of the two edges in A is a loop of length $\frac{1}{2}$ and any difference of markings map will map it to a loop homotopic to an immersed loop of length $\frac{2}{3}$. Thus the length of the image cannot be smaller than $\frac{2}{3}$, and so the slope of any difference of markings map on either edge cannot be less than $\frac{4}{3}$. More generally, we observe:

Lemma 2.2. *If α is any nontrivial conjugacy class then*

$$\log \frac{\ell_{\Gamma'}(\alpha)}{\ell_{\Gamma}(\alpha)} \leq d(\Gamma, \Gamma')$$

So for any α we obtain a lower bound on the distance. In our example, the lower bound agrees with the upper bound provided by the explicit difference of markings map. This determines the distance.

We will say that a conjugacy class α is a *witness* if equality holds in the statement of the Lemma.

In a similar way, one can compute that $d(B, A) = \log \frac{3}{2}$ by considering the map $B \rightarrow A$ that collapses the middle edge, and the witness loop formed by the other two edges.

Note in particular that $d(A, B) \neq d(B, A)$.

Exercise 2.3. Let A be as above and let C_ϵ be the graph in the same 1-simplex as A with lengths of edges ϵ and $1 - \epsilon$. Show that $d(C_\epsilon, A) \rightarrow \infty$ as $\epsilon \rightarrow 0$, but $d(A, C_\epsilon)$ stays bounded by $\log 2$.

Thus the distance function is not even quasi-symmetric, i.e. $\frac{d(X, Y)}{d(Y, X)}$ can be arbitrarily large. However, a theorem of Handel-Mosher [25] states that the restriction of d to any thick part $\mathcal{X}_n(\epsilon)$ is quasi-symmetric. See also [2] for a different proof.

2.4. Tension graph, train track structure

Here is the crucial fact. It is analogous to Teichmüller's theorem for Riemann surfaces. It states that witnesses always exist.

Lemma 2.4. *Suppose $d(\Gamma, \Gamma') = \log \lambda$. Then there is a conjugacy class $\alpha \in \mathcal{C}$ such that*

$$\frac{\ell_{\Gamma'}(\alpha)}{\ell_{\Gamma}(\alpha)} = \lambda$$

Note that for *any* α inequality \leq holds. So the lemma says that we can define the distance alternatively as

$$d(\Gamma, \Gamma') = \log \max_{\alpha} \frac{\ell_{\Gamma'}(\alpha)}{\ell_{\Gamma}(\alpha)}$$

The equality between the min and the max is an instance of the *max-flow min-cut* principle.

The proof of Lemma 2.4 introduces the key idea of train tracks.

PROOF. Fix a difference of markings map $\phi : \Gamma \rightarrow \Gamma'$ with $\sigma(\phi) = \lambda$. By $\Delta = \Delta_{\phi}$ denote the union of those edges of Γ on which the slope of f is λ . This subgraph of Γ is called the *tension graph* for ϕ , and some of its vertices may have valence 1 or 2. Now let v be a vertex of Δ . A *direction* at v in Δ is a germ of geodesic paths $[0, \epsilon] \rightarrow \Delta$ sending 0 to v . Alternatively, it is an oriented edge of Δ with initial vertex at v . Denote the set of these directions by $T_v(\Delta)$. Its cardinality is the valence of v in Δ and this set plays the role of the unit tangent space at v . Now ϕ induces a map (kind of a derivative)

$$\phi_* : T_v(\Delta) \rightarrow T_{\phi(v)}(\Gamma')$$

since for small ϵ it sends a geodesic $\gamma : [0, \epsilon] \rightarrow \Delta$ to a geodesic $\phi\gamma : [0, \epsilon] \rightarrow \Gamma'$ (parametrized with speed λ). Here $\phi(v)$ may not be a vertex, in which case $T_{\phi(v)}(\Gamma')$ naturally has two directions. Thus we have an equivalence relation on $T_v(\Delta)$:

$$d_1 \sim d_2 \iff \phi_*(d_1) = \phi_*(d_2)$$

A *train track structure* on a graph Δ is simply a collection of equivalence relations on the sets $T_v(\Delta)$ for all vertices v . Thus the tension graph is naturally equipped with a train track structure. The definition is motivated by Thurston's train tracks on surfaces.

It is customary to draw equivalent directions as tangent to each other. The equivalence classes are *gates*. An immersed path in Δ (thought of as a train route) is *legal* if whenever it passes through a vertex, the entering and the exiting gates are distinct. Otherwise, a path is *illegal*. Similarly, a turn (i.e. an unoriented pair of distinct directions) is illegal if the directions are equivalent; otherwise the turn is legal. More informally, legal paths do not make 180° turns.

Figure 2.2 shows the tension graphs with their train track structures from the examples in Section 2.3. The tension graph of $\phi : A \rightarrow B$ is all of A and the vertex has two gates. For the map $B \rightarrow A$ the tension graph is a circle formed by two edges and all turns are legal.

Now we make the following two observations:

- if the immersed loop $\alpha|_{\Gamma}$ representing a conjugacy class α in Γ is contained in Δ and is legal, then $\frac{\ell_{\Gamma'}(\alpha)}{\ell_{\Gamma}(\alpha)} = \lambda$, i.e. α is a witness,




FIGURE 2.2. Tension graphs with their train track structures from examples in 2.3.

- if every vertex of Δ has at least two gates, then Δ contains a legal loop; in fact this loop can be chosen to cross every oriented edge at most once.

The first of these claims is an exercise in definitions: f has slope λ on each edge of $\alpha|\Gamma$ and consecutive edges are mapped without backtracking by definition of legality. For the second claim, keep extending a legal path until the same oriented edge repeats.

Of course, in general Δ may have vertices with only one gate. To finish the proof we will show that ϕ may be perturbed so that every vertex has at least two gates.

Claim: Suppose v is a vertex of Δ_ϕ with only one gate. Then ϕ may be perturbed to $\phi' : \Gamma \rightarrow \Gamma'$ so that $\sigma(\phi') = \lambda$ and $\Delta_{\phi'} \subsetneq \Delta_\phi$.

Repeating this operation will eventually produce a perturbation of ϕ whose tension graph has at least two gates at every vertex (note that the set of edges where the slope is λ cannot become empty by the assumption that $d(\Gamma, \Gamma') = \log \lambda$).

Proof of Claim. The homotopy ϕ_t from ϕ to ϕ' will be stationary on all vertices except for v , and it will move $\phi(v)$ slightly in the direction $\phi_*(d)$, where $d \in T_v(\Delta)$. All maps ϕ_t are linear on edges. Thus the slope is unaffected on edges not incident to v , it decreases on edges in Δ incident to v , and it may increase on edges outside Δ incident to v . The perturbation is small so that even the increased slope on such edges is $< \lambda$. Thus $\Delta_{\phi'} \subset \Delta_\phi$ but $\Delta_{\phi'}$ does not contain v and edges incident to it. \square

Exercise 2.5. [21] Show that in any graph with a train track structure with at least two gates at every vertex, there is a legal loop that is either embedded, or it forms a “figure 8” crossing each edge once, or it forms a “dumbbell”, crossing edges in the two loops once and edges in the connecting arc twice. See Figure 2.3.




FIGURE 2.3. Possible forms of candidates. Train track structure is suggested by the pictures.

We say that an immersed loop in a graph Γ (without any train track structure) is a *candidate* if it has a form as in Exercise 2.5. Thus, given Γ, Γ' there is always a candidate in Γ which is a witness for $d(\Gamma, \Gamma')$. Thus there is a simple algorithm to compute distances $d(\Gamma, \Gamma')$ in Outer space. Simply look at the ratio of lengths in Γ' and in Γ of all candidate loops in Γ and take the log of the largest such ratio.

Exercise 2.6. Let R_3 be the rose in \mathcal{X}_3 with all edges of length $\frac{1}{3}$ and with inverse marking given by a, b, c , and let Γ be another such rose but with inverse marking

given by $abA, bacB, a$. Find all candidates in each that are witnesses for the distance to the other.

Exercise 2.7. Consider the automorphism Φ of $\mathbb{F}_4 = \langle a, b, c, d \rangle$ given by $a \rightarrow b \rightarrow c \rightarrow d \rightarrow ADCB$ (capital letters are inverses of the lowercase letters).

- Let R be the rose with the identity marking (so the edges correspond to a, b, c, d) and with all lengths $\frac{1}{4}$. Compute $d(R, R\Phi)$.
- Find the graph Γ in the same simplex as R (i.e. the same marking, but edge lengths can be arbitrary) so that $d(\Gamma, \Gamma\Phi)$ is minimal.
- Can you find a graph Γ' in a small neighborhood of Γ so that $d(\Gamma', \Gamma'\Phi) < d(\Gamma, \Gamma\Phi)$?

Exercise 2.8. If Γ, Γ' are distinct points in \mathcal{X}_n show that there are conjugacy classes α, β such that $\ell_\Gamma(\alpha) > \ell_{\Gamma'}(\alpha)$ and $\ell_\Gamma(\beta) < \ell_{\Gamma'}(\beta)$. Deduce that the length function $\mathcal{L} : \mathcal{X}_n \rightarrow (0, \infty)^{\mathcal{C}}$ and the projectivized length function $\mathcal{X}_n \rightarrow \mathbb{P}(0, \infty)^{\mathcal{C}}$ are injective. This is called the *length spectrum rigidity*.

Exercise 2.9. For a marked graph Γ let K_Γ be the finite set of candidates for Γ and for all marked graphs obtained from Γ by collapsing a forest. Show that lengths of elements of K_Γ determine each point of $\Sigma(\Gamma)$.

But, surprisingly, there is no finite collection of conjugacy classes whose lengths define an injection of \mathcal{X}_n into \mathbb{R}^N , see [34].

The following two properties of the Lipschitz metric point out similarities with the ℓ^∞ metric.

Exercise 2.10. Show that in each simplex straight lines are geodesics (not necessarily parametrized with unit speed).

Hint: Let $\Gamma_1, \Gamma_2, \Gamma_3$ be three points along a straight line with Γ_2 between the other two. Argue that any witness for $\Gamma_1 \rightarrow \Gamma_3$ is also a witness for $\Gamma_1 \rightarrow \Gamma_2$ and for $\Gamma_2 \rightarrow \Gamma_3$.

Exercise 2.11. Show that geodesics are not unique in general. Specifically, in rank 2, show that there are geodesics contained in a 2-simplex with endpoints on one edge, but with the geodesic intersecting the interior.

Exercise 2.12. Show that the distance function $d : \mathcal{X}_n \times \mathcal{X}_n \rightarrow [0, \infty)$ is continuous.

Hint: It suffices to prove continuity on $\Sigma(\Gamma) \times \mathcal{X}_n$ for every simplex $\Sigma(\Gamma)$. Now use a variant of Exercise 2.9.

2.5. Folding paths

A folding path is determined by an optimal map $\phi : \Gamma \rightarrow \Gamma'$ such that the tension graph Δ_ϕ is all of Γ and every vertex has at least two gates. It is a geodesic path Γ_t from Γ to Γ' and for each $t < t'$ it comes with an optimal map $\Gamma_t \rightarrow \Gamma_{t'}$ so that the tension graph is all Γ_t and these maps compose correctly for $t < t' < t''$. To define an initial segment of this path choose $\epsilon > 0$ smaller than half the length of the shortest edge, and for $t \in [0, \epsilon]$ define Γ_t by identifying segments of length t issuing from any vertex in equivalent directions. Then rescale to make volume equal to 1.

For example, for the map $\phi : A \rightarrow B$ considered in Section 2.3, the time t graph before rescaling would have one edge of length $2t$ and two edges of length $1 - 2t$.

There are naturally induced maps $\Gamma_t \rightarrow B$ so at $t = \epsilon$ one can repeat the procedure to continue the path.

It is not clear *a priori* that this defines a path globally. If ϕ is simplicial with respect to some subdivisions of Γ and Γ' and the lengths of all edges in each subdivision are equal, the procedure amounts to Stallings' folding, identifying two edges whenever they share a vertex and map to the same edge (but here we do it continuously resulting in a path in \mathcal{X}_n).

A very elegant definition of folding paths is due to Skora [33]. It is most conveniently described in terms of the universal cover $\tilde{\phi} : \tilde{\Gamma} \rightarrow \tilde{\Gamma}'$. Consider the graph of $\tilde{\phi}$:

$$Gr(\tilde{\phi}) = \{(u, v) \in \tilde{\Gamma} \times \tilde{\Gamma}' \mid \tilde{\phi}(u) = v\}$$

and define the *vertical t -neighborhood* of the graph $Gr(\tilde{\phi})$:

$$N_t = \{(u, v) \in \tilde{\Gamma} \times \tilde{\Gamma}' \mid d(\tilde{\phi}(u), v) \leq t\}$$

where d refers to the path metric on $\tilde{\Gamma}'$. Restrict the horizontal foliation of $\tilde{\Gamma} \times \tilde{\Gamma}'$ by $\tilde{\Gamma} \times \{v\}$, $v \in \tilde{\Gamma}'$ to N_t and define $\tilde{\Gamma}_t$ as the quotient space where all components of leaves are collapsed. Then $\tilde{\Gamma}_t$ is a tree and its quotient by the action of \mathbb{F}_n is the desired graph Γ_t (which needs to be rescaled). For $t = 0$ we have $\Gamma_t = \Gamma$ and for t large $\Gamma_t = \Gamma'$.

To get a feel for this definition, consider the ‘‘tent map’’ $\phi : [-1, 1] \rightarrow [0, \lambda]$ for $\lambda > 0$, which has slope λ on $[-1, 0]$ and slope $-\lambda$ on $[0, 1]$. The graph of this map is pictured in Figure 2.4 (with the target thought of as \mathbb{R}).




FIGURE 2.4. Construction of folding paths following Skora.

The metric on $\tilde{\Gamma}_t$ comes from projecting to the first coordinate and maps $\Gamma_t \rightarrow \Gamma_{t'}$ for $t < t'$ from inclusion $N_t \hookrightarrow N_{t'}$.

To see that a folding path is always a geodesic take any legal loop in Γ and observe that its image in Γ_t is legal for $\Gamma_t \rightarrow \Gamma_{t'}$ for any $t' > t$ and that it is a witness for that map (the slope of the map on each edge is the ratio of lengths of the loop at $\Gamma_{t'}$ and Γ_t).

Example 2.13. Let Γ be the rose in \mathcal{X}_2 with identity marking, and with $\ell(a) = \lambda^{-2}$ and $\ell(b) = \lambda^{-1}$ where $\lambda > 0$ satisfies $\lambda^{-1} + \lambda^{-2} = 1$ (see Example 3.6). Let $\phi : \Gamma \rightarrow \Gamma\Phi$ be the optimal map for Φ given by $a \mapsto b$, $b \mapsto ab$ suggested by Φ , so ϕ has slope λ on both edges and $\Delta = \Gamma$. The folding path from Γ to $\Gamma\Phi$ amounts to identifying the terminal portion of the edge b around the edge a in the direction of A (note that $\{A, B\}$ is the only illegal turn).

Proposition 2.14. [21] *d is a geodesic metric.*

PROOF. Choose an optimal map $\phi : \Gamma \rightarrow \Gamma'$. If $\Delta_\phi = \Gamma$ (and all vertices have ≥ 2 gates) the folding path is a geodesic from Γ to Γ' . If $\Delta_\phi \neq \Gamma$ start by scaling Δ_ϕ up and the edges in the complement up until the tension graph Δ increases. If there are any vertices with one gate, adjust ϕ . Continue until $\Delta = \Gamma$ and then follow with a folding path. See also [21] and [7] for further discussion. \square

Exercise 2.15. What is the geodesic constructed in this proof in the case $\Gamma = B$ and $\Gamma' = A$ in the example in Section 2.3?

Exercise 2.16. Find geodesics from R to Γ and from Γ to R in Exercise 2.6.

Exercise 2.17.** Can a folding path intersect some $\Sigma(\Gamma)$ in a disconnected set?

Classification of automorphisms

Recall the classification of isometries of hyperbolic space.

Definition 3.1. Let $\Phi \in \text{Out}(\mathbb{F}_n)$. The *displacement function* of Φ is the function

$$D = D_\Phi : \mathcal{X}_n \rightarrow [0, \infty)$$

given by $D(\Gamma) = d(\Gamma, \Gamma\Phi)$.

We denote $\tau(\Phi) = \inf D_\Phi$, the *translation length* of Φ .

Definition 3.2. Φ is

- *hyperbolic* if $\inf D > 0$ and minimum is realized,
- *elliptic* if $\inf D = 0$ and minimum is realized (equivalently, Φ fixes a point of \mathcal{X}_n),
- *parabolic* if the minimum of D is not realized.

We now describe the quality of each of these classes, following Bers' approach to Thurston's classification of mapping classes [4].

3.1. Elliptic automorphisms

Example 3.3. Let $\Phi : \mathbb{F}_2 \rightarrow \mathbb{F}_2$ be given by $a \mapsto b, b \mapsto a$. Then Φ is elliptic as it fixes the rose with identity marking and edge lengths $\frac{1}{2}$.

The following is an immediate consequence of the fact that point stabilizers are finite.

Proposition 3.4. *Every elliptic automorphism has finite order.*

The converse also holds, namely every automorphism of finite order is elliptic, by Nielsen Realization.

Exercise 3.5. Show that the automorphism Φ defined by $a \mapsto b, b \mapsto Ab$ is elliptic and find a fixed point. Hint: First find the fixed points of Φ^3 .

3.2. Hyperbolic automorphisms

Example 3.6. Let Φ be given by $a \mapsto b, b \mapsto ab$. Let $\phi : \Gamma \rightarrow \Gamma$ be the map suggested by Φ on the rose Γ . Now assign lengths so that ϕ has the same slope on both edges, say λ . Temporarily assigning 1 to a we see from $\phi(a) = b$ that b must have length λ . Then from $\phi(b) = ab$ we get the equation $\lambda^2 = 1 + \lambda$, whose only positive root is the golden ratio $\lambda = \frac{1+\sqrt{5}}{2}$. Now we must rescale to get volume 1, i.e. we must set $\ell(a) = \frac{1}{1+\lambda} = \lambda^{-2}$ and $\ell(b) = \frac{\lambda}{1+\lambda} = \lambda^{-1}$.

Now consider the train track structure on Γ induced by ϕ . There are 3 gates: $\{a\}$, $\{b\}$ and $\{A, B\}$. Observe that ϕ sends legal paths to legal paths. To prove this observation, one only needs to check:

- the image of each edge is a legal path, and

- legal turns are mapped to legal turns (equivalently, f induces an injective map on the set of gates at every vertex).

Both of these are easy to check: ab is a legal path and the map on gates is $\{a\} \mapsto \{b\}$, $\{b\} \mapsto \{a\}$, $\{A, B\} \mapsto \{A, B\}$.

In particular, for any $m = 1, 2, \dots$ the map $\phi^m : \Gamma \rightarrow \Gamma$ is optimal with the same train track structure. We now see that $d(\Gamma, \Gamma\Phi^m) = m \log \lambda$. This implies $\tau(\Phi) = \log \lambda$ and so Φ is hyperbolic.

Exercise 3.7. Prove the assertion in the last sentence.

Hint: The triangle inequality. If $d(\Gamma', \Gamma'\Phi) < \log \lambda$ one can get from Γ to $\Gamma\Phi^m$ via $\Gamma', \Gamma'\Phi, \dots, \Gamma'\Phi^m, \Gamma\Phi^m$. For large m this is a contradiction.

Exercise 3.8. Suppose $\phi : \Gamma \rightarrow \Gamma$ is any map between graphs that sends vertices to vertices and edges to nontrivial immersed paths. Let e_1, \dots, e_k be the list of all 1-cells in Γ (ignore orientations). Form the *transition matrix* M : it's a $k \times k$ matrix whose ij -entry is the number of times $f(e_j)$ crosses e_i with either orientation. For example, the transition matrix both for Example 3.6 and for Example 3.5 is $M = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$. Assume that some positive power of M has all entries positive. The classical theorem of Perron-Frobenius says that the largest in norm eigenvalue λ of M is > 1 , its eigenspace is 1-dimensional and spanned by a vector with all coordinates positive. Show that there is a metric on Γ such that ϕ has slope λ on every edge.

Definition 3.9. Let $\phi : \Gamma \rightarrow \Gamma$ be an optimal map. We say that ϕ is a *train track map* if $\Delta_\phi = \Gamma$, every vertex has at least two gates, and ϕ maps legal paths to legal paths. The slope λ on each edge is the *dilatation* of ϕ .

The arguments of Example 3.6 prove:

Proposition 3.10. *If Φ admits a train track representative $\phi : \Gamma \rightarrow \Gamma$ then $d(\Gamma, \Gamma\Phi) = \tau(\Phi)$. Thus Φ is hyperbolic unless the dilatation is 1, and then Φ is elliptic.*

More generally, if $\phi : \Gamma \rightarrow \Gamma$ is an optimal map representing Φ such that $\phi(\Delta_\phi) \subset \Delta_\phi$ and so that $\phi : \Delta_\phi \rightarrow \Delta_\phi$ is a train track map, then Φ is hyperbolic.

Example 3.11. Let $\Phi \in \text{Out}(\mathbb{F}_3)$ be given by $a \mapsto b, b \mapsto ab, c \mapsto ca$. Then Φ is hyperbolic with $\tau(\Phi) = \log \lambda$ with λ the golden ratio just like in Example 3.6. For Γ take the rose with the metric on a, b a scaled down version of the metric in Example 3.6 and let the length of c be close to 1. Then the map $\Gamma \rightarrow \Gamma$ suggested by Φ is a train track map.

Theorem 3.12. *Every hyperbolic automorphism can be represented by an optimal map $\phi : \Gamma \rightarrow \Gamma$ so that Δ_ϕ is an invariant subgraph and $\phi : \Delta_\phi \rightarrow \Delta_\phi$ is a train track map.*

Note that Δ_ϕ may be a proper subgraph of Γ . For a proof see [6]. One can also arrange that ϕ sends vertices to vertices, but then the train track structure on Δ_ϕ has to be modified by declaring $d_1 \sim d_2$ provided there exists $k > 0$ such that $\phi_*^k(d_1) = \phi_*^k(d_2)$.

Example 3.13. The following example was obtained by entering a random (surface) automorphism into Peter Brinkmann's program *XTrain* available on the web at

<http://math.sci.cuny.cuny.edu/pages?name=XTrain>. Capital letters denote inverses of lower case letters. With notation as in Figure 3.1, the map is $a \mapsto Da$, $b \mapsto CAdaCCAdaCCAda$, $c \mapsto cADacADaccADaccADac$, $d \mapsto ddabccA$. The program also tells us that the dilatation λ is a root of the polynomial $x^4 - 10x^3 + 10x^2 - 10x + 1$ and is approximately 9.012144.




FIGURE 3.1. A more typical train track map.

Train track structure can be computed by looking at the “derivative” map. By a we denote the direction where a begins, and by A where a ends: $a \mapsto D \mapsto a$, $b \mapsto C \mapsto C$, $c \mapsto c$, $d \mapsto d$, $B \mapsto A \mapsto A$. Two directions at a vertex form an illegal turn if they eventually map to the same direction, so $A \sim B$ and $b \sim C$. This is indicated in the diagram. As an exercise, compute the lengths of edges.

Example 3.14. Here is another example, this time it does not come from a surface automorphism. The map is given by $a \mapsto c$, $b \mapsto bcAdEaCb$, $c \mapsto cAeDa$, $d \mapsto BcAeAd$, $e \mapsto BcAe$.




FIGURE 3.2. Another train track map. The only nontrivial gate is $\{d, e\}$.

When Φ is hyperbolic and Γ achieves the minimum of D_Φ , choose a geodesic path from Γ to $\Gamma\Phi$ and take the union of all Φ^m -translates of the path, $m \in \mathbb{Z}$. This is a geodesic line and Φ acts on it by translation by $\tau(\Phi)$. Such a line is an *axis* of Φ .

Exercise 3.15. Let $\phi : \Gamma \rightarrow \Gamma$ be a train track map with dilatation λ . Show that for every conjugacy class α the sequence $\ell_\Gamma(\Phi^k(\alpha))/\lambda^k$, $k = 1, 2, \dots$ is non-increasing, and it is constant if $\alpha|\Gamma$ is legal.

The limiting values in the exercise are translation lengths of an \mathbb{F}_n -action on an \mathbb{R} -tree, called the *stable tree* of Φ .

3.3. Parabolic automorphisms

Example 3.16. Let Φ be given by $a \mapsto a$, $b \mapsto ab$. Then $\tau(\Phi) = 0$ as can be seen by taking $\ell(a) = \epsilon$, $\ell(b) = 1 - \epsilon$ with $\epsilon \rightarrow 0$. But Φ has infinite order, so it must be parabolic.

Definition 3.17. An automorphism $\Phi \in \text{Out}(\mathbb{F}_n)$ is *reducible* if it can be represented as $\phi : \Gamma \rightarrow \Gamma$ so that for some subgraph $\Gamma' \subsetneq \Gamma$ we have $\phi(\Gamma') \subseteq \Gamma'$, and Γ' is not a forest (i.e. a disjoint union of trees).

Otherwise we say that Φ is *irreducible*. An automorphism is *fully irreducible* if every nonzero power is irreducible.

Examples 3.3, 3.11, 3.16 are all reducible (for 3.3 consider the dumbbell graph).

Example 3.6 is irreducible and so is Example 3.13, but this is more difficult to prove.

Theorem 3.18. *Every parabolic automorphism is reducible.*

PROOF. Fix a sequence $\Gamma_i \in \mathcal{X}_n$ such that $D_\Phi(\Gamma_i) \rightarrow \tau(\Phi)$. The key claim is:
The sequence Γ_i leaves every thick part $\mathcal{X}_n(\epsilon)$.

The claim implies the theorem by the following argument that goes back to Bers. For large i the graph Γ_i will contain very small loops. Of course there can be several scales of smallness (e.g. ϵ and ϵ^2) but we are guaranteed to have an arbitrarily large ratio between two consecutive scales. More precisely, for $\epsilon > 0$ define $\Gamma_i(\epsilon)$ to be the union of all essential loops (not necessarily immersed) of length $< \epsilon$. By construction Φ is represented as $\phi : \Gamma_i \rightarrow \Gamma_i$ with Lipschitz constant uniformly bounded by some K . Thus $f(\Gamma_i(\epsilon)) \subset \Gamma_i(K\epsilon)$, and “large ratio between two consecutive scales” means that $\Gamma_i(K\epsilon)$ deformation retracts to $\Gamma_i(\epsilon)$. Thus ϕ can be homotoped so that the core subgraph of $\Gamma_i(\epsilon)$ is invariant. Formally, one finds such a large ratio between consecutive scales by considering a long finite sequence of subgraphs such as $\Gamma_i(1) \supset \Gamma_i(1/K) \supset \Gamma_i(1/K^2) \supset \dots$. They have no contractible components and are nonempty for large i , and in a sufficiently long chain of nonempty subgraphs two consecutive ones will have the same core (i.e. minimal deformation retract).

It remains to prove the claim. The idea is that if the claim fails, we could find a graph in \mathcal{X}_n where D_Φ achieves the minimum. Conceptually, the simplest argument is to translate the claim to a statement about the quotient $\mathcal{X}_n/\text{Out}(\mathbb{F}_n)$. For clarity, let’s pretend that \mathcal{X}_n is a complete Riemannian manifold with $\text{Out}(\mathbb{F}_n)$ acting as a deck group by isometries, so that $\mathcal{X}_n/\text{Out}(\mathbb{F}_n)$ is also a complete Riemannian manifold. The statement that Φ is parabolic amounts to saying that a loop representing Φ in $\pi_1(\mathcal{X}_n/\text{Out}(\mathbb{F}_n))$ cannot be homotoped to a loop of length $\tau(\Phi)$. Projecting a geodesic path from Γ_i to $\Gamma_i\Phi$ gives a loop in $\mathcal{X}_n/\text{Out}(\mathbb{F}_n)$ based at the image $[\Gamma_i]$ of Γ_i . If all Γ_i stay in some thick part $\mathcal{X}_n(\epsilon)$ then after passing to a subsequence we may assume that $[\Gamma_i] \rightarrow [\Gamma]$. We also note that the projected loops α_i stay in some larger compact set, since the distance from a graph with a tiny loop back to the thick part is very large. Thus by Arzela-Ascoli after a further subsequence we have a limiting loop α at $[\Gamma]$ of length $\tau(\Phi)$. Now for large i the loops α and α_i are homotopic, so they all represent the conjugacy class of Φ . The length of α is $\tau(\Phi)$, contradiction.

There are some technical issues coming from the non-symmetry of the metric on \mathcal{X}_n and from the non-freeness of the action. For a less conceptual, but more elementary proof, following Bers, see [6]. \square

If $\tau(\Phi) > 0$ and Φ is parabolic, there is an “axis at infinity”.

Putting the above discussion together, we obtain the following theorem, originally proved by different methods.

Theorem 3.19 ([10]). *Every irreducible automorphism is represented by a train track map.*

3.4. Reducible automorphisms

Every automorphism of \mathbb{F}_n has a representative $\phi : \Gamma \rightarrow \Gamma$ called a *relative train track map*. It is built from train track maps like an upper triangular block matrix. More precisely,

- there is a filtration

$$\Gamma_0 \subset \Gamma_1 \subset \cdots \subset \Gamma_k = \Gamma$$

into invariant subgraphs (i.e. $\phi(\Gamma_i) \subset \Gamma_i$), and

- for every i and every edge e in $\Gamma_i - \Gamma_{i-1}$ the paths $\phi^m(e)$ for $m = 1, 2, \dots$ can backtrack only within Γ_{i-1} .

Here we take $\Gamma_{-1} = \emptyset$. Therefore $\phi : \Gamma_0 \rightarrow \Gamma_0$ is a train track map. We also assume that the filtration is maximal. The transition matrix is upper triangular, and there are two kinds of strata $\Gamma_i - \Gamma_{i-1}$:

- polynomially growing (PG): ϕ induces a cyclic permutation of the edges in $\Gamma_i - \Gamma_{i-1}$, and
- exponentially growing (EG): the transition matrix restricted to $\Gamma_i - \Gamma_{i-1}$ grows exponentially and is irreducible (for every entry some power is nonzero in that entry).

Example 3.11 has two strata, the lower is (EG) and the upper is (PG).

Relative train track maps are useful because cancellation under iterations is controlled. There are various improvements that control cancellation even further; they go under the generic name of “improved relative train tracks”, see [10, 8, 12, 20]. Some of the improvement require first passing to a suitable positive power of the automorphism, but this is usually harmless in applications. The following is a partial list of results proved using the improved relative train track technology:

- [10] For every automorphism $f : \mathbb{F}_n \rightarrow \mathbb{F}_n$ the rank of the fixed subgroup $Fix(f) = \{x \in \mathbb{F}_n \mid f(x) = x\}$ has rank $\leq n$.
- [5, 8, 9, 3] Tits Alternative: If H is any subgroup of $Out(\mathbb{F}_n)$ then either H contains \mathbb{F}_2 or H is virtually abelian.
- [19] The abstract commensurator of $Out(\mathbb{F}_n)$ is $Out(\mathbb{F}_n)$.
- [12] The mapping torus of any automorphism of \mathbb{F}_n satisfies quadratic isoperimetric inequality.
- [24] If H is any finitely generated subgroup of $Out(\mathbb{F}_n)$ then either H contains a fully irreducible automorphism, or a finite index subgroup of H fixes a nontrivial free factor of \mathbb{F}_n .

3.5. Growth

If $\Phi \in Out(\mathbb{F}_n)$ and γ is a nontrivial conjugacy class, we can define the *growth rate* of γ with respect to Φ :

$$\tau(\Phi, \gamma) = \limsup_{m \rightarrow \infty} \log \frac{\ell_\Gamma(\Phi^m(\gamma))}{m}$$

for a fixed $\Gamma \in \mathcal{X}_n$. If we decide on some other Γ' instead, the ratio of lengths with respect the two is uniformly bounded, so $\tau(\Phi, \gamma)$ is independent of the choice of Γ . In fact, as a consequence of the theory, one may replace \limsup by \lim in the definition.

If Φ is hyperbolic and we choose $\phi : \Gamma \rightarrow \Gamma$ to be as in Theorem 3.12 and if γ is a legal loop in Δ_ϕ , then $\ell_\Gamma(\Phi^m(\gamma)) = \ell_\Gamma(\gamma)\lambda^m$ for $m > 0$, so $\tau(\Phi, \gamma) = \log \lambda = \tau(\Phi)$.

Exercise 3.20. Take the automorphism Φ with $a \mapsto b$, $b \mapsto ab$, $c \mapsto d$, $d \mapsto cad$. Show that $\tau(\Phi) = \log \lambda$ for $\lambda = \frac{1+\sqrt{5}}{2}$ and that $\ell_\Gamma(\Phi^m(c)) \sim m\lambda^m$. Deduce that Φ is parabolic.

A (weak) Perron number is a positive real number which is an algebraic integer and it is greater (or equal) than the norm of any of its Galois conjugates.

Theorem 3.21. *For any $\Phi \in \text{Out}(\mathbb{F}_n)$ there are finitely many weak Perron numbers $\lambda_1, \dots, \lambda_k > 1$ so that for any γ*

$$\tau(\Phi, \gamma) = \log \lambda_i$$

for some i or $\tau(\Phi, \gamma) = 0$. Moreover, $\tau(\Phi) = \max \log \lambda_i$ (or 0 if the collection of λ_i is empty).

The numbers λ_i are the growth rates of the block of the transition matrix corresponding to the strata $\Gamma_i - \Gamma_{i-1}$.

It was shown by Thurston [35] that for every weak Perron number $\lambda > 1$ there is an automorphism Φ represented by a train track map $\phi : \Gamma \rightarrow \Gamma$ with dilatation λ .

For a more detailed information about growth, see [30].

3.6. Pathologies

For those familiar with mapping class groups, the following facts will seem like pathologies.

- $\tau(\Phi)$ may be different from $\tau(\Phi^{-1})$.
- Φ may be hyperbolic and Φ^{-1} parabolic.
- If $\tau(\Phi) = 0$ then Φ grows polynomially, but not necessarily linearly.
- If $\tau(\Phi) = \log \lambda$, λ may not be an algebraic unit (but can be any weak Perron number).

All of these facts are obstructions to an automorphism being realizable as a homeomorphism of a surface.

Hyperbolic features

To what extent is \mathcal{X}_n negatively curved? It is not completely clear what we mean by this, since the metric is not symmetric. But even if we could make sense of the question, \mathcal{X}_n could not be negatively curved, since it contains “flats” for $n \geq 4$. Consider the commuting automorphisms of F_4 , $\Phi: a \mapsto b, b \mapsto ab, c \mapsto c, d \mapsto d$, $\Psi: a \mapsto a, b \mapsto b, c \mapsto d, d \mapsto cd$. Then there are constants $C_1, C_2 > 0$ so that for the rose R with identity marking and edge lengths $1/4$ and for $k, l \in \mathbb{Z}$ we have

$$C_1(|k| + |l|) \leq d(R, R\Phi^k\Psi^l) \leq C_2(|k| + |l|)$$

i.e. we have a quasi-isometric embedding of a flat.

Exercise 4.1. Prove the inequalities.

The first negatively curved phenomenon was observed by Yael Algom-Kfir, by analogy with Minsky’s theorem [32]. First note that in hyperbolic space \mathbb{H}^n every geodesic is *strongly contracting*: there is a universal constant C so that if B is any metric ball in \mathbb{H}^n disjoint from a geodesic line ℓ then the image of B under the nearest point projection to ℓ has diameter $\leq C$. Also note that this property fails in Euclidean space, so we can view it as an indicator of negative curvature.

Theorem 4.2 ([1]). *Let Φ be a fully irreducible automorphism. Then any orbit $\{\Gamma\Phi^k\}$ of Φ is strongly contracting: there is a constant $C = C(\Phi, \Gamma)$ so that if*

$$B_{\rightarrow}(\Delta, R) = \{\Omega \in \mathcal{X}_n \mid d(\Delta, \Omega) \leq R\}$$

is a ball disjoint from the orbit then the nearest point projection of the ball to the orbit has diameter $\leq C$.

For $\Delta \in \mathcal{X}_n$ the function $k \mapsto d(\Delta, \Gamma\Phi^k)$ is proper and the minset is the nearest point projection of Δ .

To motivate what happens next, let’s look at mapping class groups. If S is a compact surface, define the *curve complex* of S to be the simplicial complex $\mathcal{C}(S)$ whose vertices are isotopy classes of essential simple closed curves (a curve is essential if it is not homotopic into the boundary or to a point), and a collection of vertices spans a simplex if the corresponding isotopy classes can be represented by pairwise disjoint curves. If S is a hyperbolic surface with totally geodesic boundary, we can work with simple closed geodesics – they are automatically disjoint if they can be isotoped to be disjoint.

In the case of the torus, or the torus with one boundary component, the curve complex is a discrete set since non-isotopic essential simple closed curves always intersect. In these cases one modifies the definition of the curve complex and puts an edge between two vertices if they can be isotoped so that they intersect in one point. The resulting graph is the classical *Farey graph*, pictured below (cf picture of reduced Outer space in Lecture 1).




FIGURE 4.1. Farey graph.




FIGURE 4.2. Farey graph in upper half plane.

Theorem 4.3 ([31]). *The curve complex is hyperbolic.*

The statement means that the 1-skeleton is δ -hyperbolic for some δ , with respect to the geodesic metric where every edge has length 1.

Teichmüller space, like Outer space, is not hyperbolic. It has a coarse map to the curve complex, and this map “kills flats”, so that the curve complex captures hyperbolic aspects of Teichmüller space.

Going back to Outer space, there are two analogs of the curve complex.

4.1. Complex of free factors \mathcal{F}_n

This complex was defined by Hatcher and Vogtmann in [27]. The definition is analogous to the Bruhat-Tits building for $SL_n(\mathbb{Z})$, which is a simplicial complex whose vertices are proper vector subspaces of \mathbb{Q}^n and simplices are chains of subspaces.

Recall that a subgroup $A < \mathbb{F}_n$ is a *free factor* if there is a subgroup $B < \mathbb{F}_n$ such that $\mathbb{F}_n = A * B$. The vertices of \mathcal{F}_n are conjugacy classes of proper (i.e. not 1 nor \mathbb{F}_n) free factors, and a simplex is induced by a chain $A_0 < A_1 < \dots < A_k$. Again in rank 2 the complex is a discrete set, so we modify the definition and put an edge between the conjugacy classes of $\langle a \rangle$ and $\langle b \rangle$ provided a and b form a basis of \mathbb{F}_2 . Since the conjugacy class of a rank 1 free factor in \mathbb{F}_2 corresponds exactly to a simple closed curve in a punctured torus, the graph \mathcal{F}_n is also the Farey graph.

Theorem 4.4. [7] \mathcal{F}_n is hyperbolic.

4.2. The complex \mathcal{S}_n of free factorizations

This complex was introduced by Hatcher in [26] in the form of a complex of spheres. The easiest way to define it is as the simplicial completion \mathcal{X}_n^* of \mathcal{X}_n (see Lecture 1). Thus a vertex of \mathcal{S}_n is a *1-edge free splitting* of \mathbb{F}_n , i.e. a minimal simplicial \mathbb{F}_n -tree with trivial edge stabilizers, 1 orbit of edges, and no global fixed points. It is convenient to draw the quotient space of such a tree and label vertices by their stabilizers, as in Bass-Serre theory. Thus a vertex looks like a picture below.




FIGURE 4.3. 1-edge splittings $A * B$ and $A *_1$ of \mathbb{F}_n .

Such graphs arise when a marked graph is equipped with a degenerate metric. Similarly, an edge of \mathcal{S}_n is a 2-edge splitting, and the endpoints of this edge are 1-edge splittings obtained by collapsing one of two orbits of edges.

Theorem 4.5 ([23]). \mathcal{S}_n is hyperbolic.

4.3. Coarse projections

Recall that when we discuss distance in \mathcal{F}_n or \mathcal{S}_n we consider points in the 1-skeleton only. We will view the distance function on all of \mathcal{F}_n or \mathcal{S}_n as being defined only “coarsely”, i.e. with a bounded ambiguity. If $x, y \in \mathcal{F}_n$ we set $d(x, y)$ to be the diameter of the union of the 1-skeleta of the simplices containing x and y respectively, and similarly for \mathcal{S}_n .

There are coarse Lipschitz projection maps $\mathcal{X}_n \rightarrow \mathcal{S}_n \rightarrow \mathcal{F}_n$. The word “coarse” means that the image of a point is a uniformly bounded set, and “Lipschitz” means that there are constants $A, B > 0$ so that if two points are at distance $\leq d$ then the union of their images has diameter $\leq Ad + B$.

To define $\pi : \mathcal{X}_n \rightarrow \mathcal{S}_n^{(0)}$, for $\Gamma \in \mathcal{X}_n$ let $\pi(\Gamma)$ be the set of vertices of the smallest simplex containing Γ . Similarly, $\rho : \mathcal{S}_n^{(0)} \rightarrow \mathcal{F}_n^{(0)}$ is defined to be the set of stabilizers of the vertices of the given splitting, viewed as an \mathbb{F}_n -tree.

Here are some facts.

- Both $\pi : \mathcal{X}_n \rightarrow \mathcal{S}_n^{(0)}$ and $\rho : \mathcal{S}_n^{(1)} \rightarrow \mathcal{F}_n$ are coarse Lipschitz maps and both are coarsely onto (i.e. there is $R > 0$ such that every R -ball in the target intersects the image).
- Both π and ρ have unbounded point inverses, even when π is restricted to the spine.

- A fully irreducible automorphism $\Phi \in \text{Out}(\mathbb{F}_n)$ acts as a hyperbolic isometry on both \mathcal{F}_n and \mathcal{S}_n .
- An automorphism which is not fully irreducible acts with bounded orbits on \mathcal{F}_n , but may be hyperbolic on \mathcal{S}_n .

4.4. Idea of the proof of hyperbolicity

Let Y be a connected graph with edge lengths 1. How does one go about proving that Y is hyperbolic. The definition requires checking that geodesic triangles are uniformly thin, but in practice it is hard to decide if a given path is a geodesic. The following outline applies to all three complexes mentioned above: the curve complex, the complex of free factors, and the free factorization complex.

The strategy is to work with “natural” paths. Suppose one is given a collection of paths \mathcal{P} in Y , each joining a pair of vertices, and assume the following:

- The collection \mathcal{P} is *transitive*, meaning that every pair of vertices in Y is connected by a path in \mathcal{P} ,
- Each path in \mathcal{P} is a *reparametrized quasigeodesic*, i.e. there are constants L, A so that for every path $\alpha : [a, b] \rightarrow Y$ in \mathcal{P} there is a homeomorphism $\tau : [a', b'] \rightarrow [a, b]$ such that $\alpha\tau : [a', b'] \rightarrow Y$ is an (L, A) -quasi-geodesic.
- There is some $\delta \geq 0$ so that every triangle formed by three paths in \mathcal{P} is δ -thin.

The following is a variant of hyperbolicity criteria in [31] and [11].

Proposition 4.6. *Suppose Y admits a collection of paths \mathcal{P} as above. Then Y is hyperbolic.*

PROOF. It suffices to prove that any loop in Y of length ℓ bounds a disk of area $\leq C\ell \log \ell$ (see [13]). Of course, Y is a graph and has no 2-cells, but one imagines attaching disks to all loops in Y of length bounded by some fixed constant and then the area of a loop is the least number of these attached disks a null-homotopy of the loop crosses, counted with multiplicity. For simplicity, assume that $\ell = 3 \cdot 2^n$ for some $n > 0$ so that we may think of the loop as a polygon with $3 \cdot 2^n$ sides. Such a polygon can be triangulated so that combinatorially it is the n -neighborhood of a fixed triangle in the Farey graph (the 1-neighborhood of a triangle consists of 4 triangles). For each diagonal in this triangulation choose a path in \mathcal{P} connecting the same pair of points. Thus we have a map from the Farey-triangulated polygon to Y that takes outer edges to edges in Y and all other edges to paths in \mathcal{P} . The area of a thin triangle is bounded by a linear function of its diameter. The diameter of the central triangle and of the triangles adjacent to it is bounded by 2^n , but subsequent layers have diameters bounded by $2^{n-1}, 2^{n-2}, \dots$ and each subsequent layer has twice as many triangles as the previous layer. Adding these numbers we get $\sim n \cdot 2^n$ which is about $\ell \cdot \log \ell$. \square

It remains to define the collection \mathcal{P} and check the above properties. For both \mathcal{F}_n and \mathcal{S}_n the collection of paths is defined by projecting folding paths in \mathcal{X}_n using the coarse projection map. The first bullet is easy to verify. In practice, the second bullet is verified by constructing a coarse Lipschitz retraction from Y to the image of a given path in \mathcal{P} . Both the second and the third bullets require hard work.

Bibliography

1. Yael Algom-Kfir, *Strongly contracting geodesics in outer space*, *Geom. Topol.* **15** (2011), no. 4, 2181–2233. MR 2862155
2. Yael Algom-Kfir and Mladen Bestvina, *Asymmetry of outer space*, *Geom. Dedicata* **156** (2012), 81–92. MR 2863547
3. Emina Alibegović, *Translation lengths in $\text{Out}(F_n)$* , *Geom. Dedicata* **92** (2002), 87–93, Dedicated to John Stallings on the occasion of his 65th birthday. MR 1934012 (2003m:20045)
4. Lipman Bers, *An extremal problem for quasiconformal mappings and a theorem by Thurston*, *Acta Math.* **141** (1978), no. 1-2, 73–98. MR 0477161 (57 #16704)
5. M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, *Geom. Funct. Anal.* **7** (1997), no. 2, 215–244. MR 1445386 (98c:20045)
6. Mladen Bestvina, *A Bers-like proof of the existence of train tracks for free group automorphisms*, *Fund. Math.* **214** (2011), no. 1, 1–12. MR 2845630
7. Mladen Bestvina and Mark Feighn, *Hyperbolicity of the complex of free factors*, arXiv:1107.3308.
8. Mladen Bestvina, Mark Feighn, and Michael Handel, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms*, *Ann. of Math. (2)* **151** (2000), no. 2, 517–623. MR 1765705 (2002a:20034)
9. ———, *The Tits alternative for $\text{Out}(F_n)$. II. A Kolchin type theorem*, *Ann. of Math. (2)* **161** (2005), no. 1, 1–59. MR 2150382 (2006f:20030)
10. Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, *Ann. of Math. (2)* **135** (1992), no. 1, 1–51. MR 1147956 (92m:20017)
11. Brian H. Bowditch, *Intersection numbers and the hyperbolicity of the curve complex*, *J. Reine Angew. Math.* **598** (2006), 105–129. MR MR2270568
12. Martin R. Bridson and Daniel Groves, *The quadratic isoperimetric inequality for mapping tori of free group automorphisms*, *Mem. Amer. Math. Soc.* **203** (2010), no. 955, xii+152. MR 2590896 (2011g:20058)
13. Martin R. Bridson and André Haefliger, *Metric spaces of non-positive curvature*, *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*, vol. 319, Springer-Verlag, Berlin, 1999. MR 1744486 (2000k:53038)
14. Kenneth S. Brown, *Cohomology of groups*, *Graduate Texts in Mathematics*, vol. 87, Springer-Verlag, New York, 1994, Corrected reprint of the 1982 original. MR 1324339 (96a:20072)
15. Matt Clay, *Contractibility of deformation spaces of G -trees*, *Algebr. Geom. Topol.* **5** (2005), 1481–1503 (electronic). MR 2186106 (2006i:20026)
16. Marc Culler, *Finite groups of outer automorphisms of a free group*, *Contributions to group theory*, *Contemp. Math.*, vol. 33, Amer. Math. Soc., Providence, RI, 1984, pp. 197–207. MR 767107 (86g:20027)
17. Marc Culler and John W. Morgan, *Group actions on \mathbf{R} -trees*, *Proc. London Math. Soc. (3)* **55** (1987), no. 3, 571–604. MR MR907233 (88f:20055)
18. Marc Culler and Karen Vogtmann, *Moduli of graphs and automorphisms of free groups*, *Invent. Math.* **84** (1986), no. 1, 91–119. MR MR830040 (87f:20048)
19. Benson Farb and Michael Handel, *Commensurations of $\text{Out}(F_n)$* , *Publ. Math. Inst. Hautes Études Sci.* (2007), no. 105, 1–48. MR 2354204 (2008j:20102)
20. Mark Feighn and Michael Handel, *Abelian subgroups of $\text{Out}(F_n)$* , *Geom. Topol.* **13** (2009), no. 3, 1657–1727. MR 2496054 (2010h:20068)
21. Stefano Francaviglia and Armando Martino, *Metric properties of Outer space*, *Pub. Mat.* **55** (2011), 433–473.

22. Vincent Guirardel and Gilbert Levitt, *Deformation spaces of trees*, Groups Geom. Dyn. **1** (2007), no. 2, 135–181. MR 2319455 (2009a:20041)
23. Michael Handel and Lee Mosher, *The free splitting complex of a free group I: Hyperbolicity*, arXiv:1111.1994.
24. ———, *Subgroup classification in $Out(F_n)$* , arXiv:0908.1255.
25. ———, *Parageometric outer automorphisms of free groups*, Trans. Amer. Math. Soc. **359** (2007), no. 7, 3153–3183 (electronic). MR MR2299450 (2008c:20045)
26. Allen Hatcher, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv. **70** (1995), no. 1, 39–62. MR 1314940 (95k:20030)
27. Allen Hatcher and Karen Vogtmann, *The complex of free factors of a free group*, Quart. J. Math. Oxford Ser. (2) **49** (1998), no. 196, 459–468. MR 1660045 (99i:20038)
28. Sebastian Hensel, Damian Osajda, and Piotr Przytycki, *Realisation and dismantlability*, arXiv:1205.0513.
29. D. G. Khramtsov, *Finite groups of automorphisms of free groups*, Mat. Zametki **38** (1985), no. 3, 386–392, 476. MR 811572 (87c:20071)
30. Gilbert Levitt, *Counting growth types of automorphisms of free groups*, Geom. Funct. Anal. **19** (2009), no. 4, 1119–1146. MR 2570318 (2011f:20068)
31. Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149. MR MR1714338 (2000i:57027)
32. Yair N. Minsky, *Quasi-projections in Teichmüller space*, J. Reine Angew. Math. **473** (1996), 121–136. MR 1390685 (97b:32020)
33. Richard K. Skora, *Deformations of length functions in groups*, preprint, 1989.
34. John Smillie and Karen Vogtmann, *Length functions and outer space*, Michigan Math. J. **39** (1992), no. 3, 485–493. MR 1182503 (93j:20054)
35. William P. Thurston, *Entropy in dimension one*, preprint 2012.
36. ———, *Minimal stretch maps between hyperbolic surfaces*, arxiv:math/9801039.
37. Karen Vogtmann, *Automorphisms of free groups and outer space*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000), vol. 94, 2002, pp. 1–31. MR 1950871 (2004b:20060)
38. Shi Cheng Wang and Bruno Zimmermann, *The maximum order of finite groups of outer automorphisms of free groups*, Math. Z. **216** (1994), no. 1, 83–87. MR 1273467 (95a:20039)
39. Bruno Zimmermann, *Über Homöomorphismen n -dimensionaler Henkelkörper und endliche Erweiterungen von Schottky-Gruppen*, Comment. Math. Helv. **56** (1981), no. 3, 474–486. MR 639363 (83f:57025)