COURSE NOTES FOR MATH 636: OUTER AUTOMORPHISM GROUPS OF FREE GROUPS

Instructor: Alex Wright

Main authors:

Alex Bauman, Ram Ekstrom, Jordan Grant, Simran Khunger, Kaya Lakein, Yucong Lei, Nathan Li, Sadie Lipman, Yuqi Meng, Urshita Pal, Lukas Scheiwiller, Katia Shchetka, Henry Talbott, Yumeng Wang, Shuhong Yang, Sicheng Zhao

With contributions by: Zach Himes, Max Lahn, Qianwen Sun

And a guest lecture by: Mladen Bestvina

Contents

Introduction: These are notes on a graduate topics course at the University of Michigan in Fall 2023. Corrections are welcome and may be sent via email to Alex Wright.

Audience and scope: These notes might be useful for students who would like preparation before tackling the literature on $\mathrm{Out}(F_n)$. We will begin without assuming any familiarity with geometric group theory, and focus on geometric aspects of the theory. A good deal of the course follows expository material of Bestvina; the reader may wish to consult, for example, the course notes of Vogtmann [Vog] to see another point of view.

Authorship: For each lecture, one course participant was designated as the author, and another as the editor. The notes for each lecture are labelled with the initials of the author followed by the initials of the editor. Additionally Alex Wright edited the notes.

Citations: Only a small number of citations are provided. Often these are not original sources, but rather the sources that the lectures were based off of.

Acknowledgements: Alex Wright would like to especially thank Mladen Bestvina and Ric Wade, for generously explaining a great many things related to this course. Alex Wright would additionally like to thank Matt Clay, Armando Martino, Kasra Rafi, and Richard Webb for helpful conversations. Finally, Alex Wright is grateful to Jean Pierre Mutanguha, Caglar Uyanik, and Sanghoon Kwak for both for giving seminar talks in Fall 2023 which very nicely complemented the course and for interesting discussions during their visits; and again to Sanghoon Kwak for carefully reading these notes and pointing out a list of typos and corrections.

1. Introduction (08/28, HT, KS)

1.1. **Introduction and Cayley Graphs.** This course broadly lies within the topic of geometric group theory. As such, we'll be interested in studying groups by understanding their actions on geometric spaces. The geometric spaces will often be derived from the groups themselves. The prototypical example of a geometric space derived from a group is a *Cayley graph.*

Definition 1.1. Fix a group G and a generating set $S \subset G$ that is symmetric (i.e., $s \in S$ if and only if $s^{-1} \in S$). The Cayley graph $\mathrm{Cay}(G, S)$ is the graph produced by the following construction:

- The vertex set of $\mathrm{Cay}(G, S)$ is the set of elements of G.
- There is an edge between two vertices $g, h \in G$ if and only if there is an $s \in S$ such that $h = gs$.

A Cayley graph has a few nice elementary properties stemming from the group axioms:

- As is required for an undirected graph, there is an edge from g to h if and only if there is an edge from h to g. Formally, $h = gs$ if and only if $g = hs^{-1}$, and $s \in S \iff s^{-1} \in S$.
- A Cayley graph is always regular (every vertex has the same degree), and its degree is the cardinality of S. For any vertex $g \in G$, the set of adjacent vertices is by definition $\{gs : s \in S\}$, and $gs = gs'$ for $s, s' \in S$ if and only if $s = s'$ by left multiplication by g^{-1}.

Notice that this definition still makes sense if S is infinite, although we'll rarely encounter examples where this is the case.

A key property of $\mathrm{Cay}(G, S)$ is that G acts on it. Since the Cayley graph itself was defined via right multiplication, the action of G on $\mathrm{Cay}(G, S)$ will act by left multiplication. Specifically, a given $h \in G$ sends each vertex g of the Cayley graph to hg, and sends each edge between g and some gs to an edge between hg and hgs.

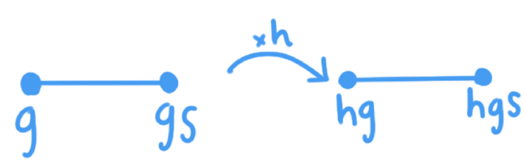

Figure 1

Since we're interested in geometric rather than purely topological spaces, we'll make $\mathrm{Cay}(G, S)$ a *metric graph* by thinking of each edge as a copy of the interval $[0, 1]$ with the standard Euclidean metric (the full metric space is then constructed by identifying appropriate endpoints). Between vertices, this metric is equivalent to word distance:

Definition 1.2. Fix a group G and a generating set $S \subset G$. For any $g \neq e$, the word length of G is the minimal positive integer l such that there exist $s_1, \ldots, s_l \in S$ with $s_1 \cdots s_l = g$. The word length of the identity is 0, and the word distance between two $g, h \in G$ is the word length of $g^{-1}h$.

1.2. **Quasi-Isometry.** When studying geometric spaces, we'll need a notion of spaces being equivalent that's a little looser than isometry. Specifically, we want both some 'local wiggle room' and some overall scaling distortion.

Definition 1.3. A map $f : (M, d_1) \to (N, d_2)$ between metric spaces is a quasi-isometry if there exist real numbers $K \geqslant 1$ and $C \geqslant 0$ such that:

(1) For all $x, y \in M$, we have
$$\frac{1}{K} d_1(x, y) - C \leqslant d_2(f(x), f(y)) \leqslant K d_1(x, y) + C.$$

(2) For all $z \in N$, there exists an $x \in M$ such that $d_2(f(x), z) \leqslant C$.

The second property on its own is called *coarsely surjective*. Here's a picture of what a quasi-isometry might 'look like' (although note that all finite-diameter spaces are trivially quasi-isometric, so this picture is just meant to give a rough idea).

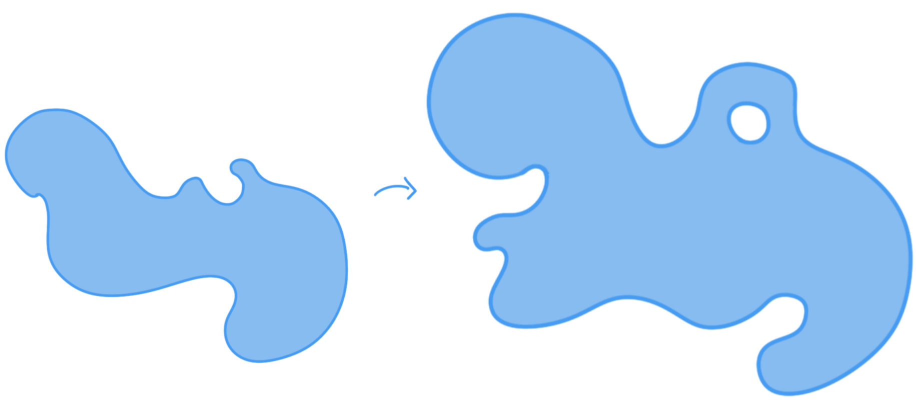

FIGURE 2

Exercise 1. *Let G be a group and let $S_1, S_2 \subset G$ be two finite generating sets. Then* $\mathrm{Cay}(G, S_1)$ *and* $\mathrm{Cay}(G, S_2)$ *are quasi-isomorphic.*

1.3. **Important Classes of Groups.** In this course, we'll restrict our attention to three important classes of groups.

The first and most well-understood class consists of lattices in Lie groups. Recall that a lattice Γ lying inside a Lie group G is a discrete subgroup such that the quotient G/Γ has finite measure. One basic example is $\mathbb{Z}^n \subset \mathbb{R}^n$, with the quotient being the unit-volume flat torus. A more geometrically interesting example is $SL_n(\mathbb{Z}) \subset GL_n(\mathbb{Z})$; the $n = 2$ case yields the unit tangent bundle over the modular curve as a quotient.

The second class consists of mapping class groups of finite-genus surfaces (with or without an orientability assumption).

Definition 1.4. Fix a finite-genus topological surface Σ. Its mapping class group $MCG(\Sigma)$ is the group of all homeomorphisms $\phi : S \to S$ quotiented by the subgroup of all homeomorphisms homotopic to the identity map.

While mapping class groups are generally difficult to understand concretely, one can think of $MCG(\Sigma)$ as the group of 'symmetries' of the surface that don't come simply from nudging points around via homotopy (and thus represent interesting topological information).

The last class consists of outer automorphisms of free groups.

Definition 1.5. Fix a positive integer n, and let F_n be the free group on n generators. Let $\mathrm{Aut}(F_n)$ be its automorphism group, and let $\mathrm{Inn}(F_n)$ be the subgroup consisting of all automorphisms that can be written in the form $\phi(g) = aga^{-1}$ for some $a \in F_n$. Then the outer automorphism group $\mathrm{Out}(F_n)$ is defined as $\mathrm{Aut}(F_n)/\mathrm{Inn}(F_n)$.

Remark 1.6. If F_n is defined to be the free group on generators $a_1, \ldots, a_n$, any homomorphism ϕ from F_n to another group G is uniquely determined by images of the a_i. Moreover, for any choice of n elements $b_1, \ldots, b_n \in G$, there is a unique homomorphism $\phi : F_n \to G$ such that $\phi(a_i) = b_i$.

For any free group $F_n = \langle a_1, \ldots, a_n \rangle$, recall that there is an abelianization map $\alpha : F_n \to \mathbb{Z}^n$ defined by choosing a coordinate basis $e_1, \ldots, e_n$ for $\mathbb{Z}^n$ and declaring $\mathrm{Ab}(a_i) = e_i$. Moreover, for any homomorphism $\phi : F_n \to F_n$, we can easily check that ϕ sends the commutator subgroup of F_n into itself, and hence this commutator subgroup is contained in the kernel of $\alpha \circ \phi$. By standard group theory, this implies the existence of a unique map $\mathrm{Ab}(\phi) : \mathbb{Z}^n \to \mathbb{Z}^n$ such that the following diagram commutes:

$$\begin{tikzcd} F_n \arrow[r, "\phi"] \arrow[d, "\alpha"'] & F_n \arrow[d, "\alpha"] \\ \mathbb{Z}^n \arrow[r, "\mathrm{Ab}(\phi)"'] & \mathbb{Z}^n \end{tikzcd}$$

If ϕ is in fact an automorphism, then it descends to an automorphism on the commutator subgroup, and $\mathrm{Ab}(\phi)$ is also an automorphism. The automorphism group of $\mathbb{Z}^n$ is $GL_n(\mathbb{Z})$, so we obtain a homomorphism $\mathrm{Ab} : \mathrm{Aut}(F_n) \to GL_n(\mathbb{Z})$. Moreover, since $\alpha \circ \phi$ is the identity whenever $\phi \in \mathrm{Inn}(F_n)$, Ab descends to a homomorphism

$$\overline{\mathrm{Ab}} : \mathrm{Out}(F_n) \to GL_n(\mathbb{Z}).$$

Next time, we'll prove the following:

Lemma 1.7. *The map* $\overline{\mathrm{Ab}}$ *defined above is surjective.*

2. Analogies and big picture (08/30, KL, ML)

To prove Lemma 1.7, we will use the following fact:

Fact 2.1. *Let* E_{ij} *denote the* $n \times n$ *square matrix with a* 1 *in the* ij*th component and* 0*s elsewhere. Then* $\mathrm{GL}(n, \mathbb{Z})$ *is generated as a group by all matrices of the form* $\mathrm{Id} + E_{ij}$ *with* $1 \leqslant i \neq j \leqslant n$*, together with all signed permutation matrices.*

(This fact is proved using Smith normal form, a relative of row reduction for integer matrices.)

Example 2.2. In $\mathrm{GL}(3,\mathbb{Z})$, an example of a matrix of the form $\mathrm{Id}+E_{ij}$ and of a signed permutation matrix are given by

$$\mathrm{Id}+E_{13} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 0 & 0 & -1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

respectively.

We will moreover make use of the following definition:

Definition 2.3. Given a fixed choice of free generating set for F_n, an automorphism of F_n is called a signed permutation automorphism if it sends each generator to another generator or its inverse.

Example 2.4. A signed permutation automorphism of $F_3 = \langle a, b, c\rangle$ is given by $a_1 \mapsto a_2$, $a_2 \mapsto a_3$, $a_3 \mapsto a_1^{-1}$.

Proof of Lemma 1.7. Given Fact 2.1, it suffices to show that all signed permutation matrices and all matrices of the form $\mathrm{Id}+E_{ij}$ with $1 \leqslant i \neq j \leqslant n$ are in the image of $\mathrm{Ab} : \mathrm{Aut}(F_n) \to \mathrm{GL}(n,\mathbb{Z})$. Now, signed permutation automorphisms of F_n map to signed permutation matrices under Ab (since $\mathrm{Ab}(\rho)$ outputs the matrix representation of the abelianization of ρ). Moreover, for fixed $1 \leqslant i \neq j \leqslant n$, the automorphism of $F_n = \langle a_1, \dots, a_n\rangle$ defined by

$$a_k \mapsto \begin{cases} a_k & \text{if } k \neq j \\ a_j a_i & \text{if } k = j \end{cases}$$

maps to $\mathrm{Id}+E_{ij}$ under Ab. (To see that this is indeed an automorphism of F_n, we can write down its inverse.) $\square$

Lemma 1.7 describes a relationship between $\mathrm{Out}(F_n)$ and the lattice $\mathrm{GL}(n,\mathbb{Z})$. Next, we seek to describe a relationship to mapping class groups. We will be guided by the following observation:

Remark 2.5. The fundamental group of a surface with at least one puncture is free. For instance, if $\Sigma_{g,n}$ is a surface of genus g with $n \geqslant 1$ punctures, then $\pi_1(\Sigma_{g,n}) \cong F_{2g+n-1}$.

We begin with a brief review of some covering space theory.

- Given a homeomorphism f of a surface Σ, there is an induced group homomorphism of fundamental groups $f_* : \pi_1(\Sigma, x_0) \to \pi_1(\Sigma, f(x_0))$. This is an isomorphism.
- If a is an arc from x_0 to $f(x_0)$ in Σ, then there is an isomorphism

$$I_a : \pi_1(\Sigma, f(x_0)) \xrightarrow{\sim} \pi_1(\Sigma, x_0)$$

defined by $[\gamma] \mapsto [a^{-1}\gamma a]$. Moreover, if a' is another arc, then $I_a \circ I_{a'}^{-1} : \pi_1(\Sigma, x_0) \to \pi_1(\Sigma, x_0)$ is an inner automorphism of $\pi_1(\Sigma, x_0)$. See Figure 3.

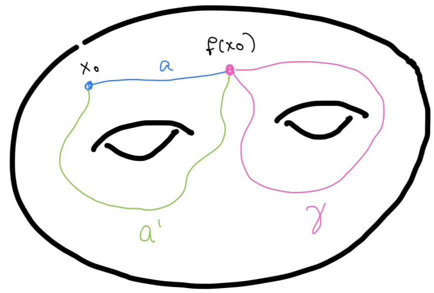

FIGURE 3

In particular, this second remark means that we can make the following definition:

Definition 2.6. Let $\Psi : \mathrm{Homeo}(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma, x_0))$ denote the map defined by $\Psi(f) = I_a \circ f_*$.

Indeed, for any two arcs a and a', we have that $I_a \circ f_*$ and $I_{a'} \circ f_*$ differ by $I_a \circ I_{a'}^{-1}$, which is an inner automorphism of $\pi_1(\Sigma, x_0)$. Thus, $I_a \circ f_*$ is a well-defined element of $\mathrm{Out}(\pi_1(\Sigma, x_0)) = \mathrm{Aut}(\pi_1(\Sigma, x_0))/\,\mathrm{Inn}(\pi_1(\Sigma, x_0))$.

Exercise 2. *If $f, g \in \mathrm{Homeo}(\Sigma)$ are homotopic, then $\Psi(f) = \Psi(g)$.*

This exercise allows us to make the following definition:

Definition 2.7. Let $\overline{\Psi} : MCG(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma))$ be given by $\overline{\Psi}([f]) = \Psi(f)$.

One can check that $\overline{\Psi}$ is a group homomorphism. We now claim the following:

Lemma 2.8. *The map $\overline{\Psi}$ is injective.*

Proof Sketch. Use a result from Chapter 1 of Hatcher, which states that two maps on $K(\pi, 1)$'s whose induced maps on fundamental groups agree must be homotopic [Hat02, Proposition 1B.9]. (Later in the course, in Lemma 11.4, we'll do a different proof of this proposition in the case of graphs, and that proof can also be adapted to surfaces. But in any case the proof in Hatcher is not hard.) □

Thus, what we have established so far in Lemmas 2.8 and 1.7 is that we can think of mapping class groups as subgroups of outer automorphisms, and of $\mathrm{GL}(n, \mathbb{Z})$ as a quotient of $\mathrm{Out}(F_n)$.

Remark 2.9. If a surface Σ is closed, meaning that it has no punctures and no boundary, and Σ is not a sphere or the projective plane, then $\overline{\Psi}$ is bijective. Otherwise, it is usually very far from being bijective! See [FM12a, Chapter 8] for more details.

These relationships, while precise, are not as helpful as we may hope, as the two maps $\overline{\mathrm{Ab}}$ and $\overline{\Psi}$ typically do not capture most of $\mathrm{Out}(F_n)$. We thus describe a deeper analogy between our three types of groups.

Things to note about the spaces appearing in this analogy are that all the red ones are all contractible, and all the blue ones have rich negatively curved behaviour (though

Lattices in Lie groups $\Gamma \subset G$	$MCG(\Sigma)$	$\mathrm{Out}(F_n)$
acts on symmetric space $K\backslash G$	acts on Teichmüller space $T(\Sigma) = \{$marked hyp. surfaces$\}$	acts on outer space $\{$marked metric graphs$\}$
quotient: $K\backslash G/\Gamma$ locally symmetric space	quotient: $T(\Sigma)/MCG(\Sigma)$ moduli space of hyp. surfaces	quotient: moduli space of metric graphs
(no good analogy in line with flavor of course)	acts on the curve graph	acts on free factor graph, free splitting graph
Jordan normal form	Nielsen-Thurston classification	(partial) Nielsen-Thurston classification

they are not always true negatively curved spaces), and all the brown ones are Gromov hyperbolic (which means that they have some tree-like behavior).

Finally, we end class by mentioning several different perspectives we might take/things we might ask about in studying $\mathrm{Out}(F_n)$.

- Dynamics of individual automorphisms (e.g., we might seek to describe the word length of an automorphism raised to the nth power applied to a word)
- Classical group theory (e.g. with regards to the study of presentations)
- Algorithmic (e.g. given an element of a free group, we might seek to determine whether it is part of a free basis)
- Metric graphs (their fundamental groups are free)
- Groups acting on trees (the universal cover of a graph is a tree, which means that group actions on trees are ubiquitous in the study of F_n; this leads us to Bass-Serre theory)
- Algebraic topology (e.g. we might ask about the group cohomology of $\mathrm{Out}(F_n)$)
- Metric (e.g. we might ask what $\mathrm{Out}(F_n)$ looks like up to quasi-isometry)
- Probabilistic (e.g. we might ask what a typical element of $\mathrm{Out}(F_n)$ looks like as sampled by a random walk)
- Three manifolds (e.g. handlebodies and doubled handlebodies have free fundamental group, and one especially thinks about discs or spheres embedded in these three manifolds)

3. Graphs and graph morphisms (09/01, ML, KS)

We will follow expository lectures of Bestvina for at least the next week [Besb, Besa].

Definition 3.1 (Graph)**.** A *graph* is a cell complex of dimension 1. Such a complex can also be viewed combinatorially via its vertices and edges.

Recall that if Γ is a graph, a choice of (a) a maximal tree $T \subseteq \Gamma$; and (b) orientations and labels to all edges not in T gives an identification of $\pi_1(\Gamma, v)$ with a free group.

Example 3.2. Let Γ denote the graph in Figure 5. Drawn in red is a maximal tree T, and edges not in T are oriented and labeled. This gives an identification of $\pi_1(\Gamma, v)$ with the free group $\langle a, b, c, d, e\rangle$ as follows: for a loop in Γ, read the labels (with orientation) as the loop traverses the edges of Γ, ignoring edges that are in T.

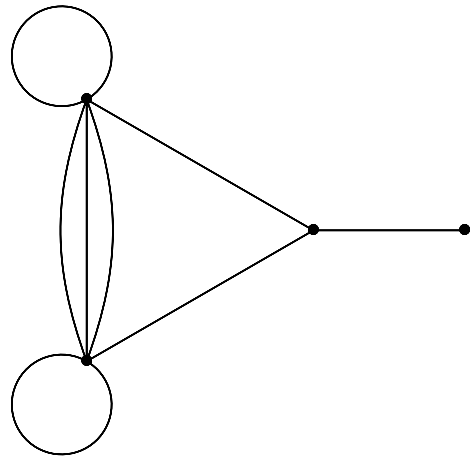

FIGURE 4. An example of a graph

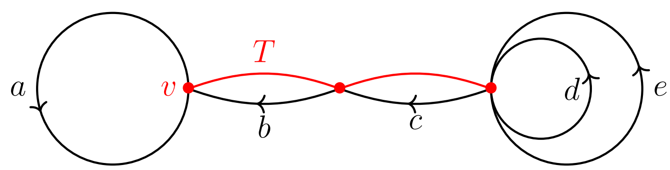

FIGURE 5

For a more sophisticated approach, note that the map from Γ to the wedge of five circles collapsing T to a point is a homotopy equivalence and therefore induces an isomorphism of fundamental groups. See [Hat02, Proposition 1A.2] for a review.

Definition 3.3 (Graph morphism)**.** A *morphism* of graphs is a continuous map $f : \Gamma \to \Gamma'$ which sends vertices to vertices and edges to edges. We typically implicitly assume parameterizations of the edges of Γ and Γ' are given, and implicitly require that f restrict to linear maps on edges.

We should note that morphisms of graphs don't collapse edges; in particular, the continuous map which collapses the red edge of graph in Figure 6 (obtaining a graph with one vertex and two edges, which are loops) is not a graph morphism. Similarly,

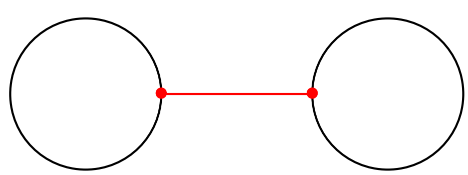

FIGURE 6

morphisms of graphs don't expand edges to paths of more than one edge; in particular, the continuous map which expands the above edge into two consecutive edges (obtaining the graph in Figure 7) is not a graph morphism either.

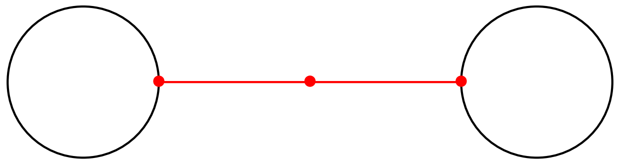

FIGURE 7

Example 3.4. Consider the graph morphism illustrated in Figure 8, obtained by identifying the left two vertices and the two edges from the other vertex to these vertices: This graph morphism is locally injective everywhere except at the right-most vertex.

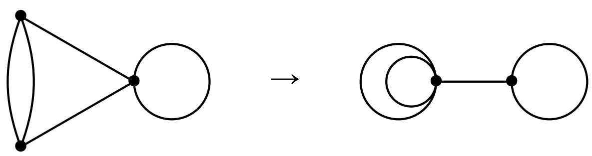

FIGURE 8

Definition 3.5 (Immersion)**.** A morphism of graphs is called an *immersion* if it is locally injective at every vertex. Note that morphisms of graphs are always locally injective at non-vertex points.

Note that injective graph morphisms are immersions, and so are covering maps. Also, compositions of immersions are immersions.

Lemma 3.6. *If $f\colon \Gamma \to \Gamma'$ is an immersion of finite graphs, then it can be factored as $f = c \circ i$ for some finite graph Γ'', injective graph morphism $i\colon \Gamma \to \Gamma''$, and covering map $c\colon \Gamma'' \to \Gamma'$.*

Moreover, if all vertices in Γ' have the same number of preimages in Γ, then we can pick i to be a bijection on vertices.

Proof Sketch. Let d be the maximal number of preimages of a vertex of of Γ'. We will first construct a space Γ''_0 (not technically a graph) in which Γ embeds and which maps to Γ'. The space Γ''_0 consists of Γ along with enough isolated vertices so that every vertex of Γ' has exactly d preimages and half-edges around each vertex so that the map to Γ' is a local homeomorphism at every vertex.

Glue corresponding half-edges to get Γ''. Why does this work? If e is an edge of Γ', then each of the half-edges comprising e has d preimages in Γ''_0. There are $d - |f^{-1}(e)|$ copies of each half-edge in $\Gamma''_0 \backslash \Gamma$, and so there is a pairing which leads to a valid gluing. $\square$

Example 3.7. Consider the graph morphism $f\colon \Gamma \to \Gamma'$ drawn in Figure 9. In particular, f identifies the two left-most vertices, the two right-most vertices, and the two horizontal edges. Then Γ_0'' and Γ'' are drawn in Figure 10.

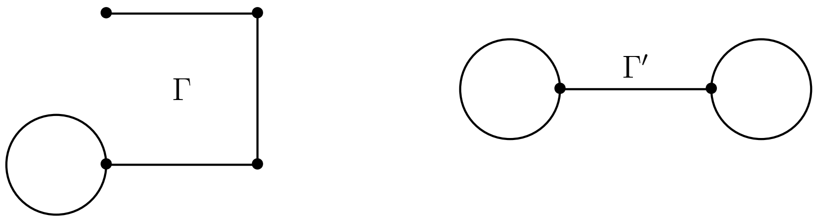

FIGURE 9

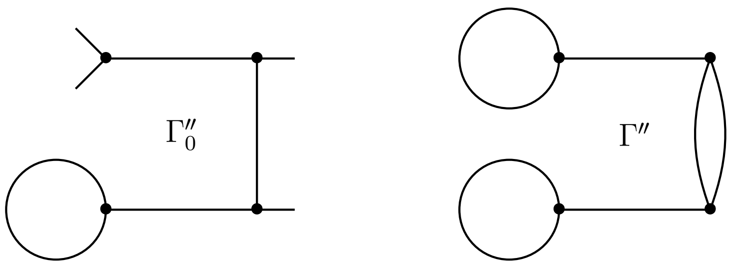

FIGURE 10

We end with one additional observation on immersions. Note that there is a unique path without backtracking between any two points of a tree. Since the universal cover of a connected graph Γ is a tree, every element of $\pi_1(\Gamma, v)$ is represented by a unique path with no backtracking. In particular, every such path is non-trivial in $\pi_1(\Gamma, v)$.

Lemma 3.8. *If $f\colon \Gamma \to \Gamma'$ is an immersion of graphs, then the induced homomorphism $f_*\colon \pi_1(\Gamma, v) \to \pi_1(\Gamma', f(v))$ is injective.*

Proof. Immersions map non-backtracking paths to non-backtracking paths. □

4. FOLDING (09/06, UP, HT)

For the lectures on folding, in addition to the notes of Bestvina we're following, one can consult the original short paper of Stallings, which still reads very well 40 years after it was written [Sta83]

Definition 4.1. A **fold** is a graph morphism $f : \Gamma \to \Gamma'$ obtained as follows: Let e_1, e_2 be two oriented edges with the same initial point. Let $\Gamma' = \Gamma/(e_1 \sim e_2)$. Lastly, let f be the corresponding quotient map.

We classify folds into four types, depending on which of e_1 and e_2 are loops and, if both are not loops, whether they have the same endpoint.

Specifically:

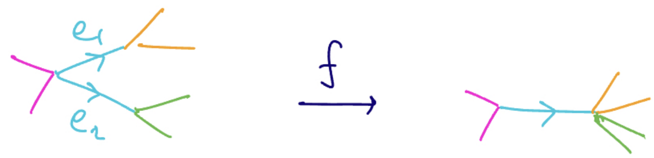

FIGURE 11. Type 1

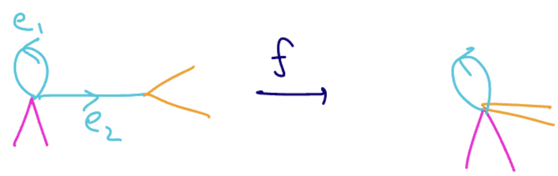

FIGURE 12. Type 2

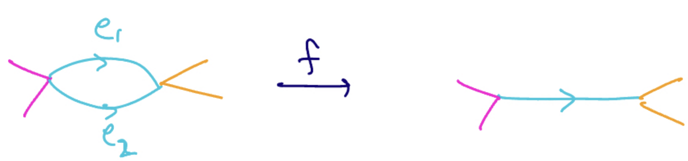

FIGURE 13. Type 3

- If both e_1, e_2 are loops, then it's type 4
- If exactly one of e_1, e_2 is a loop, then it's type 2
- If neither is a loop, and:
  - The other end point of e_1, e_2 is distinct, then type 1
  - The other end point is the same, then type 3

Remark 4.2. A fold is somewhat analogous to a row operation on a matrix.

Lemma 4.3. *If $f : \Gamma \to \Gamma'$ is type 1 or 2, then*

$$f_* : \pi_1(\Gamma) \to \pi_1(\Gamma')$$

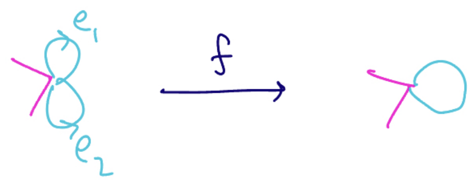

FIGURE 14. Type 4

is an isomorphism. If f is type 3 or 4, then f_* is surjective and

$$\mathrm{Rank}(\pi_1(\Gamma')) = \mathrm{Rank}(\pi_1(\Gamma)) - 1.$$

Proof. Strategy: Pick maximal trees and get explicit identifications with free groups so f_* becomes very simple.

Type 1: Pick T a maximal tree in Γ containing e_1 and e_2. Set $T' = f(T)$. Note that T' is a maximal tree: T' has all vertices of the quotient graph, and that all loops in T' lift to loops in T, so T' has no loops.

FIGURE 15

Pick an orientation of the edges of $\Gamma - T$. This gives an orientation on the edges of $\Gamma' - T'$. Now, note that $\pi_1(\Gamma)$ is isomorphic to the free group on edges in $\Gamma - T$, and $\pi_1(\Gamma')$ is isomorphic to the free group on edges in $\Gamma' - T'$. Since f induces a bijection from $\Gamma - T$ to $\Gamma' - T'$, f_* sends one basis to the other. See Figure 15.

Type 2: Pick T containing e_2 (assuming that e_1 is a loop and e_2 isn't). Pick $T' = f(T - e_2)$.

Then f induces a bijection from $\Gamma - T$ to $\Gamma' - T'$ as above, and f_* is an isomorphism. For instance, in the example, f_* sends $a \mapsto \alpha$, $b \mapsto \beta$, and $c \mapsto \gamma$. See Figure 16.

Type 3: Pick T containing e_1 but not e_2, and set $T' = f(T)$.

Then f_* 'deletes' the generator of the fundamental group corresponding to e_2, but makes no other changes. In the example, f_* sends $a \mapsto \alpha, b \mapsto \beta, c \mapsto \mathrm{id}$. See Figure 17.

Type 4: Pick T arbitrarily, let $T' = f(T)$. See Figure 18. Generators corresponding to

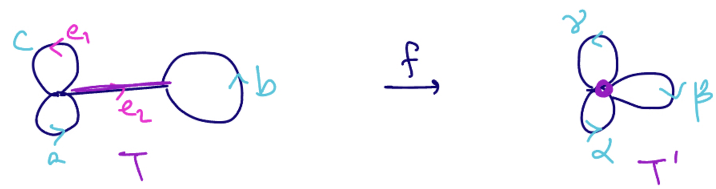

FIGURE 16

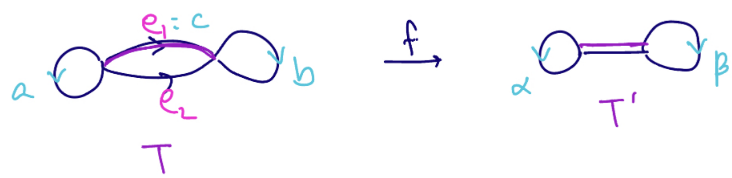

FIGURE 17. This figure has a typo: it should be e_2 that is labelled c, not e_1.

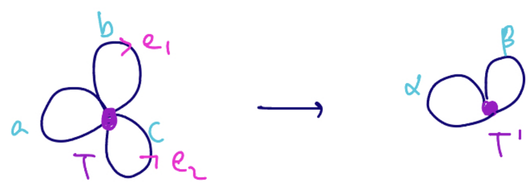

FIGURE 18

e_1 and e_2 get identified, and otherwise f_* is a bijection on generators. In our example, f_* is described by $a \mapsto \alpha, b \mapsto \beta$ and $c \mapsto \beta$. $\square$

Remark 4.4. We could have shortened our analysis noting that (if you add vertices), a type 2 fold is a sequence of type 1 folds. See Figure 19.

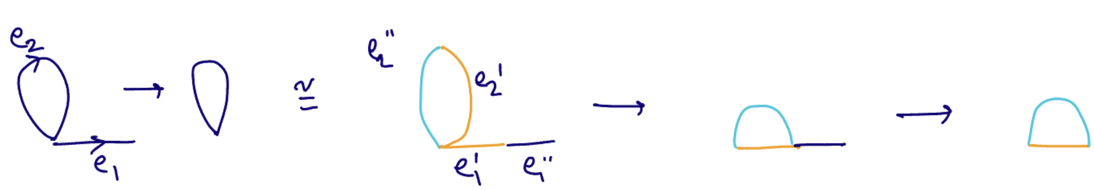

FIGURE 19

Theorem 4.5 (Stallings, 1983)**.** *Every morphism $\Gamma \to \Gamma'$ between finite graphs can be factored as*

$$\Gamma = \Gamma_0 \to \Gamma_1 \to \cdots \to \Gamma_k \to \Gamma'$$

where the last map is an immersion and the rest are folds.

Proof. By induction on the number of edges in Γ, if $\Gamma \to \Gamma'$ is not an immersion, it is not locally injective near some vertex v, so a pair of edges e_1, e_2 leaving v get identified. Fold this pair to get

$$\Gamma \to \Gamma/(e_1 \sim e_2) \to \Gamma'$$

and apply induction to $\Gamma/(e_1 \sim e_2) \to \Gamma'$. $\square$

Remark 4.6. It is important to note that

$$\mathrm{Im}(\pi_1(\Gamma) \to \pi_1(\Gamma')) = \mathrm{Im}(\pi_1(\Gamma_k) \to \pi_1(\Gamma')),$$

since folds are surjective on π_1.

Remark 4.7. Given a finitely generated subgroup $H = \langle w_1, w_2, \ldots, w_k \rangle$ of the free group $F_n = \langle a_1, a_2, \ldots, a_n \rangle$, we naturally obtain a morphism of graphs as follows. The codomain will be the rose with n (oriented) petals labelled $a_1, \ldots, a_n$. The domain is a rose with k (oriented) petals, with the i-th petal subdivided into a number of segments equal to the word length of w_i, with those segments labelled by generators of F_n so the petal spells out w_i. This is illustrated by the following example.

Example 4.8. $F_2 = \langle a, b \rangle$. $H = \langle a^3 b, \overline{a}bab, a^2\overline{b}a \rangle$. (Convention: $\overline{a} = a^{-1}$). See Figure 20. Note that each generator w_i of H becomes a subdivided loop in a bouquet graph,

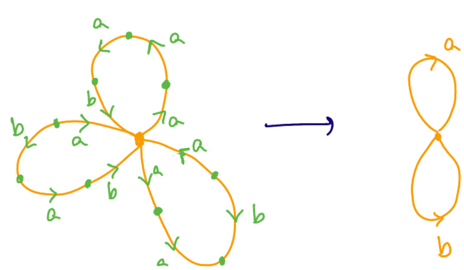

FIGURE 20

and the edges of this subdivided loop are labeled according to the letters in the word w_i.

5. First applications (09/08, SL, SK)

Corollary 5.1. *For every finitely generated subgroup $H \leqslant F_n$, there is an immersion g to a rose with $im(g_*) = H$.*

Proof. We begin by making a graph morphism whose image is H. Then by the factorization we get from Theorem 4.5 and applying Remark 4.6, we get such a g. $\square$

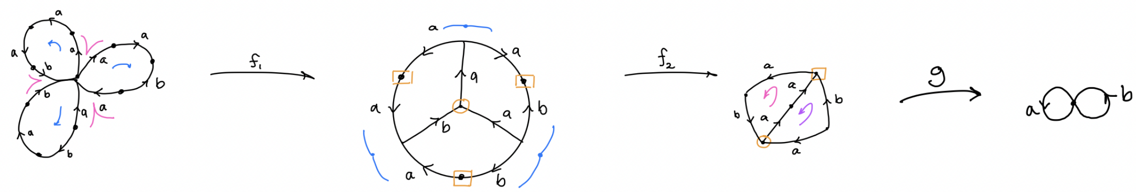

FIGURE 21

We now return to Example 4.8 to demonstrate the explicit folds. See Figure 21. The map f_1 is given by folding each of the adjacent pink edges together, and f_2 is given by folding the adjacent blue edges. The map g is an immersion.

Lemma 5.2. *If $g : \Gamma \to \Gamma'$ is an immersion, then $[\gamma'] \in \pi_1(\Gamma', v')$ is in the image of g_* if and only if γ' has a lift, γ, to Γ starting and ending at v, where $v \in g^{-1}(v')$.*

Remark 5.3. The lift is unique if it exists since g is locally injective.

Proof. If the lift exists, certainly it maps to γ. Conversely, suppose $[\gamma'] = g_*[\gamma]$. Without loss of generality, assume γ is non-backtracking. Since g is an immersion, $g\gamma$ is non-backtracking. Thus, by uniqueness of non-backtracking representative, we have $g\gamma = \gamma'$. $\square$

We again return to Example 4.8 to demonstrate how we can use Lemma 5.2 to conclude properties of our subgroup. See Figure 22. Firstly, we can conclude which group elements are in the subgroup. For example, if we consider the element $a^2 \in F_2$, this corresponds to a loop $[\gamma']$, with base point v_1. We then lift this loop to a path in Γ_k which starts at v, but does not end at v. Thus, $a^2 \notin H$.

We can also conclude, since g_* is an isomorphism to H, that a^3b and $\overline{a}b\overline{a}^2$ form a basis for H and the pink and purple loops are a basis of $\pi_1(\Gamma_k)$. In particular, the rank of H is 2.

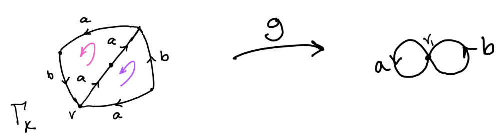

FIGURE 22

Definition 5.4. A subgroup A of a free group F_r is a **free factor** if there exists a subgroup $W < F_r$ such that $F_r = A * W$. Equivalently, A is a free factor if there exists a basis of F_r containing a basis of A.

Exercise 3. *$A = \langle aba^{-1}b^{-1}\rangle$ and $A = \langle a^2\rangle$ are not free factors of $F_2 = \langle a, b\rangle$.*

Example 5.5. $A = \langle a\rangle$ is a free factor of F_2, with $W = \langle b\rangle$. Note that W is not unique. For example, we could choose $W = \langle ab\rangle$.

Theorem 5.6 (Marshall Hall). *If H is a finitely generated subgroup of F_n, then there exists a finite index subgroup $H' < F_n$ such that H is a subgroup of H' and H is a free factor of H'.*

Before we prove this theorem, we first introduce a lemma.

Lemma 5.7. *If Γ' is a subgraph of Γ, $\pi_1(\Gamma')$ is a free factor of $\pi_1(\Gamma)$.*

Proof. Pick a maximal tree T' of Γ' and extend it to a maximal tree T of Γ. $\pi_1(\Gamma)$ has a basis corresponding to the edges $\Gamma - T$, and $\pi_1(\Gamma')$ has a basis corresponding to the edges $\Gamma' - T'$. □

We now prove Theorem 5.6.

Proof of Theorem 5.6. Pick an immersion $g : \Gamma' \to R_n$ where R_n is the rose with n petals and $im(g_*) = H$. Factor g as $g = c \circ i$, where $i : \Gamma' \to \Gamma$ is an injection and $c : \Gamma \to R_n$ is a covering map. Then, $\pi_1(\Gamma)$ is finite index in $\pi_1(R_n)$ since c is a finite cover, and $\pi_1(\Gamma')$ is a free factor in $\pi_1(\Gamma)$ by Lemma 5.7. □

Definition 5.8. A group is **residually finite** if the intersection of all the finite index subgroups is the identity. In other words, for any non trivial element, you can find a finite index subgroup not containing it.

As a warm up, we prove the following lemma:

Lemma 5.9. *F_n is residually finite.*

Proof. Take a non-backtracking loop γ in R_n. One can naturally associate to γ an immersion g from a subdivided interval to R_n, so the image of g is γ. The domain of g can be viewed as a line segment that's split up into multiple labelled edges so that it reads off the word associated to γ.

Factor g as an inclusion, i, from the subdivided interval to some Γ' and a cover $c : \Gamma' \to R_n$. Then, $\pi_1(\Gamma')$ is finite index and $[\gamma] \notin \pi_1(\Gamma')$. □

6. Cores (09/11, RE, JG)

Theorem 6.1. *If $H < F_n$ is finitely generated and $g \in F_n \backslash H$, then there is a finite index subgroup $H' \subset F_n$ containing H with $g \notin H'$.*

Proof Sketch. Choose an immersion $f : \Gamma \to R_n$ of a graph Γ to the rose R_n such that $\mathrm{Im}(f_*) = H$. It follows that the word g may not be read off in Γ, starting and ending at the basepoint. If necessary we may enlarge Γ to a graph Γ_0 so that we can read off g in Γ_0 starting at the basepoint and ending at a point that isn't the basepoint. The enlarged graph looks like the original graph with a new subdivided interval sticking off the point where we originally got stuck reading g; the segments of this interval are labelled so we can finish reading g. Assume that g is reduced so that the map $\Gamma_0 \to R_n$ is still an immersion.

We may now construct a finite cover $\tilde{\Gamma} \to R_n$ and an inclusion $\Gamma_0 \hookrightarrow \tilde{\Gamma}$ so that one obtains a factorization:

$$\begin{tikzcd} \Gamma_0 \arrow[r, hook] \arrow[dr] & \tilde{\Gamma} \arrow[d] \\ & R_n \end{tikzcd}$$

Put $H' = \pi_1(\tilde{\Gamma})$ and observe that $H \subset H'$. Then H' is a finite index subgroup of F_n because $\tilde{\Gamma}$ is a finite cover, and $g \notin H'$. $\square$

Remark 6.2. If H is as in the above theorem, then the conclusion of the theorem is equivalent to the statement that H is the intersection of all finite index subgroups F_n containing H.

Theorem 6.3. *F_n is Hopfian, meaning that every surjective group homomorphism $f\colon F_n \to F_n$ is injective.*

Proof. Suppose that F_n is freely generated by the elements $a_1, \ldots, a_n$. Consider the following map of graphs below where the petals of the graph on the left are subdivided and labeled in such a way to spell out the word $f(a_i)$ as one traverses around the petal.

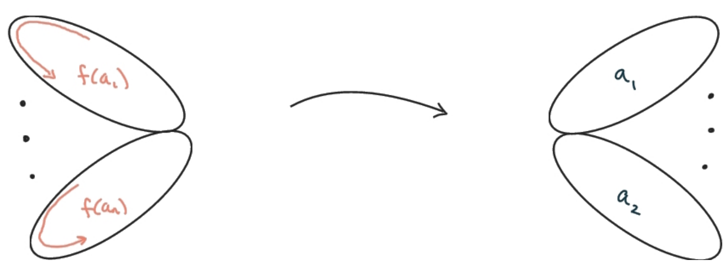

FIGURE 23

Folding the map above we obtain a sequence of maps

$$\Gamma_0 \to \Gamma_1 \to \cdots \to \Gamma_h \to R_n$$

where $\Gamma_h \to R_n$ is an immersion. Since the composition is surjective on π_1, it follows that the map $\Gamma_h \to R_n$ is a graph homeomorphism for otherwise there is some word in F_n that can't be read off in Γ_h and this word is not in the image. One may also observe that the folds above are necessarily of type 1 or 2 (since the rank may not decrease). As these folds induce isomorphisms on the fundamental group we deduce that f is an isomorphism. $\square$

Exercise 4. *Give an algorithm to decide if a homomorphism $F_n \to F_n$ is:*

(1) Surjective.
(2) Injective.

Exercise 5. *Suppose w is a reduced word. Show that $a_1, \ldots, a_{n-1}, w$ is a basis if and only if w is of the form $w = w' a_n^{\pm 1} w''$, with $w', w'' \in \langle a_1, \ldots, a_{n-1} \rangle$. (See Hint A.1.)*

Corollary 6.4. *A set of n elements of F_n is a basis if and only if it generates F_n.*

Proof. Say $\{w_1, \ldots, w_n\}$ generate F_n. The map defined by $a_i \mapsto w_i$ extends to a surjective group homomorphism $F_n \to F_n$. This must be an isomorphism, implying that $\{w_1, \ldots, w_n\}$ form a basis of F_n. □

We now turn our attention to the main topic of this section which are "cores" of graphs.

Definition 6.5 (Unpointed Cores)**.** Let Γ be a possibly infinite graph. The (unpointed) core $\mathrm{Core}(\Gamma)$ is the union of all immersed loops, i.e. non-backtracking loops of Γ.

We emphasize that our loops do not have a distinguished basepoint. An immersed loop can be see as an immersion from a subdivided circle to the graph.

Definition 6.6 (Pointed Cores)**.** Let (Γ, v) be a graph with a base point v. The pointed core $\mathrm{Core}(\Gamma, v)$ is the union of all immersed paths that start and end at v.

As an example, consider the graph below with labeled unpointed and pointed cores.

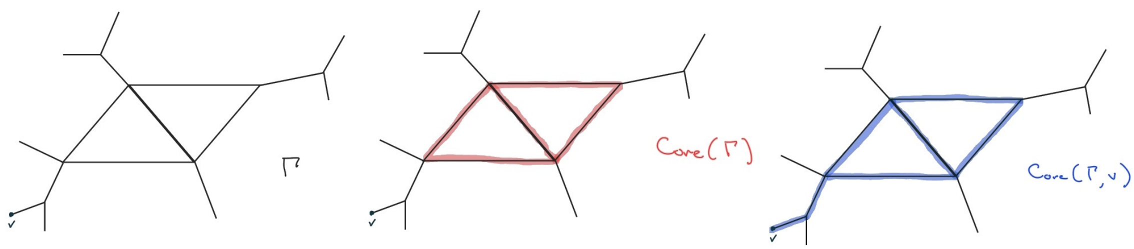

FIGURE 24

Definition 6.7. A hanging tree of Γ is a subgraph T that is a tree and only meets the closure of its complement at a single point.

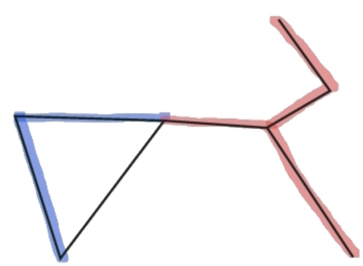

FIGURE 25. The red subgraph is a hanging tree while the blue subgraph is not a hanging tree for it meets its complement at two vertices.

Lemma 6.8. *Assume that Γ is a connected graph that is not a tree.*

(1) $\mathrm{Core}(\Gamma)$ *does not share edges with any hanging tree.*
(2) $\mathrm{Core}(\Gamma)$ *is connected.*
(3) Every component of $\Gamma \backslash \mathrm{Core}(\Gamma)$ *is a hanging tree.*

Proof. Addressing (1), let T be a hanging tree that meets its complement at v. Evidently by its definition, any loop entering T must backtrack along some edge in T. Consequently no immersed loop of Γ shares an edge with T. This establishes (1).

For (2) say that w_1 and w_2 are vertices of the core. It suffices to show that they lie in the same connected component of $\mathrm{Core}(\Gamma)$. Select two immersed loops Γ_1 and Γ_2 so that w_i is a vertex lying on γ_i (for $i = 1, 2$). If γ_1 and γ_2 intersect it is immediate that w_1 and w_2 lie in the same connected component. Otherwise, since Γ is connected, there exists a minimal length path α joining γ_1 to γ_2. The graph determined by the union of the paths $\gamma_1 + \gamma_2 + \alpha$ can be realized as the image of an immersed loop (see the figure below).

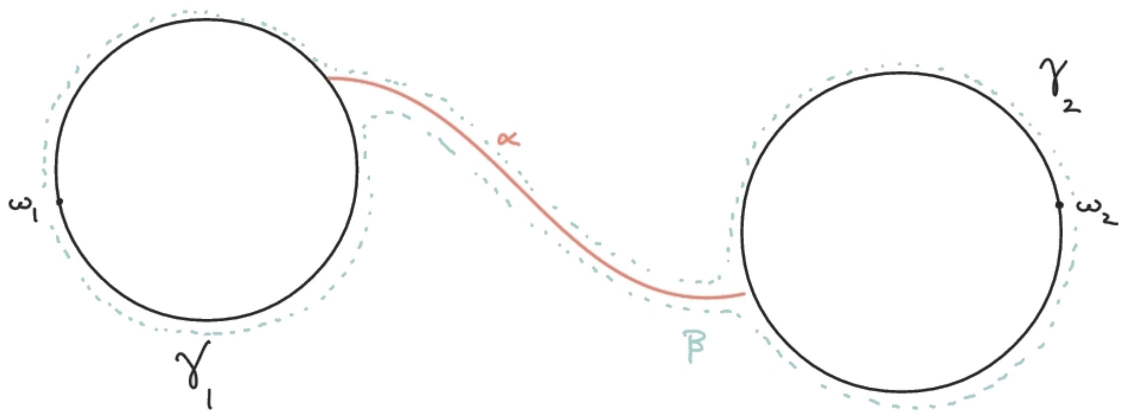

Figure 26. The "dumbbell" immersed loop

Whence w_1 and w_2 lie in the same connected component of $\mathrm{Core}(\Gamma)$. Thus $\mathrm{Core}(\Gamma)$ is connected.

Finally, for (3) let T be a component of $\Gamma \backslash \mathrm{Core}(\Gamma)$. Then T is necessarily a tree. It suffices to show that T cannot meet $\mathrm{Core}(\Gamma)$ at a pair of points p_1 and p_2. If it did, then the union a minimal length path from p_1 to p_2 in $\mathrm{Core}(\Gamma)$ and the union of a minimal length path from p_1 to p_2 in T (such paths exist by connectedness) determines an immersed loop of Γ not lying entirely within $\mathrm{Core}(\Gamma)$—contrary to the definition of the core. □

The following corollary is now immediate.

Corollary 6.9. *The complement* $\Gamma \backslash \mathrm{Core}(\Gamma)$ *is the disjoint union of the maximal hanging trees of* Γ.

7. Cores (09/13, SZ, KS)

Lemma 7.1. *If* $v \in \mathrm{Core}(\Gamma)$*, then* $\mathrm{Core}(\Gamma, v) = \mathrm{Core}(\Gamma)$*. If* $v \notin \mathrm{Core}(\Gamma)$ *and* α *is the unique immersed path from* v *to* $\mathrm{Core}(\Gamma)$*, then* $\mathrm{Core}(\Gamma, v) = \mathrm{Core}(\Gamma) \cup \alpha$.

Remark 7.2. If $v \notin \mathrm{Core}(\Gamma)$, it is in a hanging tree T that attaches to $\mathrm{Core}(\Gamma)$ at p. α is the unique immersed path in T from p to v.

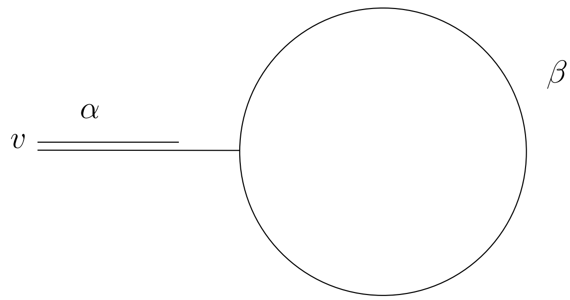

FIGURE 27. Picture for $v \notin \mathrm{Core}(\Gamma)$

Proof. First we show $\mathrm{Core}(\Gamma) \subset \mathrm{Core}(\Gamma, v)$. Indeed if β is an immersed loop, then a path like shows $\alpha, \beta \subset \mathrm{Core}(\Gamma, v)$. Therefore $\mathrm{Core}(\Gamma) \cup \alpha \subset \mathrm{Core}(\Gamma, v)$.

For the other inclusion, by Lemma 6.8, $\Gamma \backslash (\mathrm{Core}(\Gamma) \cup \alpha)$ is a union of hanging trees and hence is disjoint from immersed path based at v. □

Lemma 7.3. *If e is an edge of $\mathrm{Core}(\Gamma, v)$ then the induced map of*

$$(\Gamma \backslash e, v) \to (\Gamma, v)$$

is not surjective on π_1.

Proof. Notice that the map is an immersion and thus injective on π_1. By definition, e lives on some immersed path γ that starts and ends at v. Hence $[\gamma] \notin \pi_1(\Gamma \backslash e, v)$, as each element in π_1 is uniquely represented by a immersed path based at v. □

Lemma 7.4. *Let Γ be a finite graph and H a subgroup of $\pi_1(\Gamma, v)$. Let $\rho : (\Gamma_H, v_H) \to (\Gamma, v)$ be the cover corresponding to H. Then*

(1) $\rho|_{\mathrm{Core}(\Gamma_H, v_H)}$ is an immersion with $\mathrm{Im}\left((\rho|_{\mathrm{Core}(\Gamma_H, v_H)})_*\right) = H$.
(2) If $(\Gamma', v') = \mathrm{Core}(\Gamma', v')$ and

$$g : (\Gamma', v') \to (\Gamma, v)$$

is an immersion with $\mathrm{Im}(g_*) = H$, then there is an isomorphism i making the following diagram commute

$$\begin{tikzcd}
\mathrm{Core}(\Gamma_H, v_H) \arrow[dr, "\rho"'] & & (\Gamma', v') \arrow[ll, "i"'] \arrow[dl, "g"] \\
& (\Gamma, v) &
\end{tikzcd}$$

The main statement here is the second part, which says that "if it looks like the core of the H-cover, it is the core of the H-cover."

Proof. We address the two claims one at a time:

(1) $\rho|_{\mathrm{Core}(\Gamma_H, v_H)}$ is an immersion since ρ is a cover. The induced image in π_1 is H, since by definition $\mathrm{Im}(\rho_*) = H$ and cutting off trees does not change π_1.

(2) Since $\mathrm{Im}(\rho_*) = H = \mathrm{image}(g_*)$, covering space theory provides a lift of g to a map $i : (\Gamma', v') \to (\Gamma_H, v_H)$. Since i is a lift of g, it is an immersion. Since immersion maps immersed path to immersed path, and since $(\Gamma', v') = \mathrm{Core}(\Gamma', v')$, we get $i(\Gamma') \subset \mathrm{Core}(\Gamma_H, v_H)$.

Now we consider $\mathrm{Core}(\Gamma_H, v_H)$ as the codomain. We complete the proof as follows.

i **is injective:** Factor i as an inclusion followed by a covering map. The cover must be trivial since i_* is surjective. (In fact it is an isomorphism since $\mathrm{Im}(\rho_*) = H$.)

i **is surjective:** If the image misses a edge, i_* cannot be surjective by Lemma 7.3. □

It would be nice if a graph being equal to its own pointed core was a property that is preserved under folding. It isn't though. So, we are going to use a property that is stronger than the graph being equal to its own pointed core, chosen so this property is preserved under folding. That way, at the end of the folding, we can be sure that if the starting graph had this property then the end graph is equal to its own pointed core.

Lemma 7.5. *Suppose $(\Gamma, v) \to (\Gamma', v')$ is such that*

$(\star)$ *at every vertex $w \neq v$, there is a pair of outward-oriented edges whose image in Γ' are not the same oriented edge.*

Then the same is true after doing a fold to the map.

Proof. Notice that folding only effect pairs of outgoing edges with the same image. □

Non-example

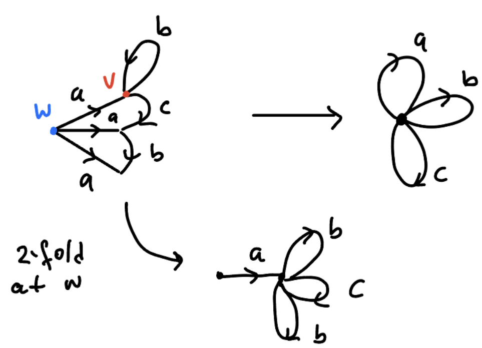

Figure 28. Non-example

Corollary 7.6. *If $f : \Gamma' \to \Gamma$ satisfies $(\star)$ and we fold to get*

$$\begin{array}{ccc} \Gamma' = \Gamma_0 \to \Gamma_1 \to \cdots \to \Gamma_k \xrightarrow{g} \Gamma \\ v' \quad \longrightarrow \quad v_k \to v \end{array}$$

If $\mathrm{Im}(f_*) = H$, then $g : \Gamma_k \to \Gamma$ is isomorphic to the core of the H cover of Γ.

Proof. The last lemma shows $\Gamma_k \to \Gamma$ satisfies $(\star)$ condition. A graph is its own pointed core if no vertex except maybe the basepoint is a leaf. So Γ_k has no leaves other than maybe v_k. So $(\Gamma_k, v_k) = \mathrm{Core}(\Gamma_k, v_k)$. g is an immersion with $\mathrm{Im}(g_*) = H$. By Lemma 7.4 we have the result. ☐

Remark 7.7. If $w_1, \ldots, w_k$ are reduced words in F_n, then the usual map $\Gamma \to R_n$ shown in Figure 23 from a subdivided rose to a rose satisfies the $(\star)$ condition. In particular, the result of folding that map does not depend on how the folding is done, since no matter how the folds are performed one obtains the core of the subgroup generated by the w_i.

Exercise 6. *Suppose $H < F_n$ is a finitely generated subgroup. Show H is of finite index if and only if the core of the H cover of R_n is a cover. (See Hint A.2.)*

Exercise 7. *Show, for any homomorphism $h : F_n \to F_n$, there exists a splitting $F_n = A * B$ such that $A \subset \ker h$ and $B \cap \ker h = \{e\}$. (See Hint A.3.)*

8. Finite generation (09/15, LS, QS)

The goal of this section is to prove a theorem of Nielsen which gives a very explicit generating set of $\mathrm{Aut}(F_n)$.

Definition 8.1 (Signed permutation automorphism)**.** Let $F_n = \langle a_1, \ldots, a_n\rangle$. Then an automorphism $\phi \in \mathrm{Aut}(F_n)$ is said to be a *signed permutation automorphism* (or sometimes just "signed permutation") if there exists a $\sigma \in S_n$ such that for all $i = 1, \ldots, n$,

$$\phi(a_i) = a_{\sigma(i)}^{\pm}.$$

Theorem 8.2 (Nielsen, First Version)**.** *The automorphism group* $\mathrm{Aut}(F_n)$ *is generated by signed permutations and the automorphism given by*

$$\phi_1 : F_n \to F_n, \quad a_i \mapsto \begin{cases} a_1 a_2, & i = 1 \\ a_i, & i > 1 \end{cases}$$

Remark 8.3. This theorem looks inherently asymmetrical, which can be fixed in the following way: if we conjugate ϕ_1 by a signed permutation, we get all automorphisms which are of the following form:

$$\forall i \neq i_0, \quad a_i \mapsto a_i$$

$$a_{i_0} \mapsto \begin{cases} a_{i_0} a_j \\ a_{i_0} a_j^{-1} \\ a_j a_{i_0} \\ a_j^{-1} a_{i_0}. \end{cases}$$

Note that ϕ_1 corresponds to $i_0 = 1$. All the automorphisms obtained this way can be thought of "elementary" matrices. In particular Theorem 8.2 is reminiscent of the fact that GL_n is generated by elementary matrices.

Definition 8.4 (Whitehead automorphism)**.** An automorphism $f \in \mathrm{Aut}(F_n)$ is said to be a *Whitehead automorphism* if it is either a signed permutation or there exists a *multiplier* $m \in \{a_i^{\pm}, i = 1, \dots n\}$ such that for all i we have

$$f(a_i) = \begin{cases} a_i \\ a_i m \\ m^{-1} a_i \\ m^{-1} a_i m \end{cases}$$

Exercise 8. *An endomorphism $f \in End(F_n)$ which has a multiplier in the above sense is an automorphism if and only if $f(m) = m$. (See Hint A.4.)*

Example 8.5. The conjugation by any element of F_n is a product of Whitehead automorphisms, in particular $Inn(F_n) \subset \langle \text{Whitehead} \rangle$.

Example 8.6. The following is a Whitehead automorphism of F_3 with multiplier a_2:

$$\begin{aligned} a_1 &\mapsto a_2^{-1} a_1 \\ a_2 &\mapsto a_2 \\ a_3 &\mapsto a_3 a_2. \end{aligned}$$

Example 8.7. The following is **not** a Whitehead automorphism of F_3:

$$\begin{aligned} a_1 &\mapsto a_2^{-1} a_1 \\ a_2 &\mapsto a_2 \\ a_3 &\mapsto a_3 a_2^{-1}. \end{aligned}$$

Remark 8.8. Every Whitehead automorphism is a product of conjugates of ϕ_1 by signed permutations. In particular, the group generated by signed permutations and ϕ_1 contains the group generated by all Whitehead automorphisms. So the following version of Nielsen's Theorem implies the first version above.

Theorem 8.9 (Nielsen, Second Version)**.** *The automorphism group* $\mathrm{Aut}(F_n)$ *is generated by all Whitehead automorphisms.*

The proof of this theorem will require the following definition:

Definition 8.10 (Change of maximal tree automorphism)**.** A *change of maximal tree automorphism of $F_n = \langle a_1, \dots, a_n \rangle$* is any automorphism of the following form:

Let Γ be a graph and T_1, T_2 two maximal trees. For $i \in \{1, 2\}$, pick an orientation of the edges of $\Gamma \backslash T_i$ and a bijection of these edeges to $\{a_1, \dots, a_n\}$. This yields an isomorphism

$$\psi_i : \pi_1(\Gamma) \xrightarrow{\cong} F_n.$$

Then the composition

$$\begin{tikzcd} & \pi_1(\Gamma) \arrow[dr, "\psi_2"] & \\ F_n \arrow[ur, "\psi_1^{-1}"] & & F_n \end{tikzcd}$$

is called a *change of maximal tree automorphism.*

Example 8.11. Note that the definition of a change of maximal tree automorphism **did not** exclude the case that $T_1 = T_2$. The change of maximal tree automorphisms where $T_1 = T_2$ are exactly the signed permutation automorphisms.

Example 8.12. Consider the graph in Figure 29.

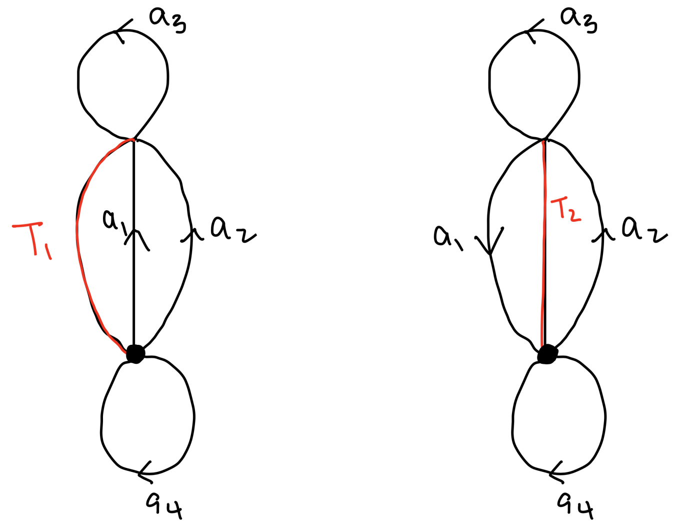

FIGURE 29

Then we have

$$\begin{aligned} a_1 &\mapsto a_1 \\ a_2 &\mapsto a_2 a_1 \\ a_3 &\mapsto a_1^{-1} a_3 a_1 \\ a_4 &\mapsto a_4 \end{aligned}$$

In particular, this automorphism is Whitehead with multiplier a_1! This example is not misleading, as we shall see in the next proposition that all change of maximal tree automorphisms are products of Whitehead automorphisms.

Proposition 8.13. *Let Γ be a graph and denote by T_1, T_2 maximal trees of Γ with $T_2 = (T_1 - e_1) \cup e_2$. Then choose the following data:*

(1) orient the edges of $\Gamma \backslash (T_1 \cup T_2)$ and label them $a_2, \dots, a_n$.
(2) orient e_1 so that it points towards *the base point in T_1*
(3) orient e_2 so that it points away from *the base point in T_2.*

When we use T_1, we label e_2 by a_1 and when we use T_2, we label e_1 by a_1. Then the change of maximal tree automorphism is Whitehead with multiplier a_1.

Lemma 8.14. *The graph $T_1 - e_1 = T_2 - e_2$ has two components which we call A, B respectively and both e_1 and e_2 join A to B.*

Note that one of the two components could be a single vertex. (We are removing the interior of the edges, not their vertices; or one can say that we're removing the edges in a combinatorial sense.)

Proof. This follows from the defining feature of a tree, which is that every edge separates. □

Proof of Proposition 8.13. First we note that contracting a tree does not change the fundamental group. Hence we may contract both components of $T_1 - e_1 = T_2 - e_2$, which we denote by A and B respectively. This gives a graph with two vertices $[A], [B]$. Up to signed permutation, we may assume that the basepoint is contained in A. The graph now looks like Figure 30, which the reader will note is very similar to Example 8.12.

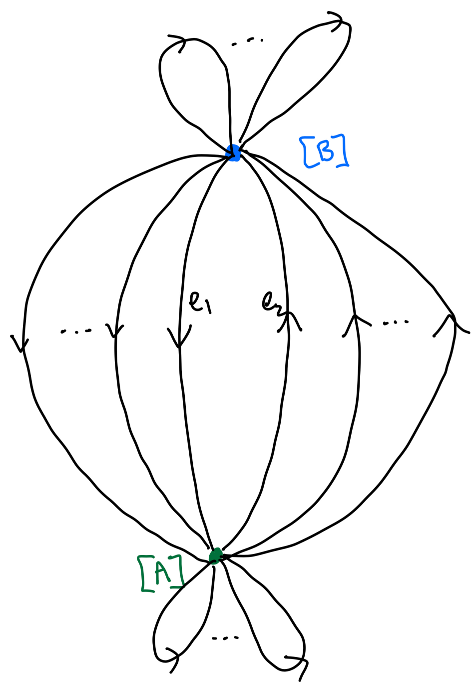

FIGURE 30

We show the change of maximal tree automorphism is Whitehead by checking what it does to generators. The generators of F_n correspond to the edges not on T_1, which fall into five classes. The image of a generator a_i can be read off determining the loop that reads a_i using T_1 and associated data, and seeing what that loops reads when one uses T_2 and associated data. The result is as follows:

(1) if a_i is the label of a loop attached to $[A]$: $a_i \mapsto a_i$
(2) if a_i is the label of a loop attached to $[B]$: $a_i \mapsto a_1^{-1} a_i a_1$
(3) if a_i is the label of the edge e_2: $a_1 \mapsto a_1$
(4) if a_i is the label of an edge going from $[A]$ to $[B]$: $a_i \mapsto a_i a_1$
(5) if a_i is the label of an edge going from $[B]$ to $[A]$: $a_i \mapsto a_1^{-1} a_i$

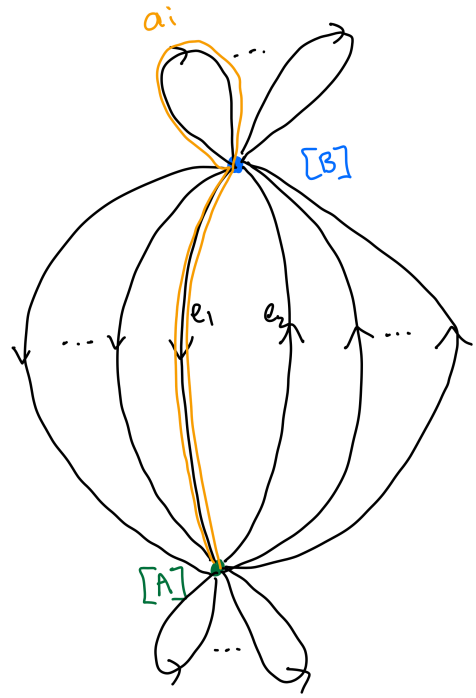

FIGURE 31. The orange loop reads a_i using T_1 and its associated data, and it reads $a_1^{-1}a_ia_1$ using T_2 and its associated data.

The second case is illustrated in Figure 31.

Thus we see that the change of maximal tree morphism is a Whitehead morphism. □

9. Finite generation, $Out(F_2)$ (09/18, NL, LS)

Exercise 9. *If Γ and Γ' are maximal trees in a finite graph G, then there exist maximal trees $\Gamma = \Gamma_0, \ldots, \Gamma_k = \Gamma'$ such that Γ_{i+1} is obtained from Γ_i by removing one edge and adding another. (See Hint A.5.)*

Corollary 9.1. *Any change of maximal tree automorphism is a product of Whitehead automorphisms.*

As a result, the previous versions of Nielsen's theorem are equivalent to the following version.

Theorem 9.2 (Nielsen, Third Version)**.** *$Aut(F_n)$ is generated by change of maximal tree automorphisms.*

Proof. Let $f \in Aut(F_n)$. We will use the usual morphism of graphs $\Gamma \to R_n$ associated to f, pictured in Figure 32, from a subdivided rose Γ to the rose R_n.

Factor this map as

$$\Gamma = \Gamma_0 \to \cdots \to \Gamma_k \to R_n$$

where $\Gamma_i \to \Gamma_{i+1}$ is a fold and the last map $g : \Gamma_k \to R_n$ is an immersion. Note that g is an isomorphism since f is surjective and $Im(f_*) = Im(g_*)$, and $w \in F_n$ is in $Im(g_*)$

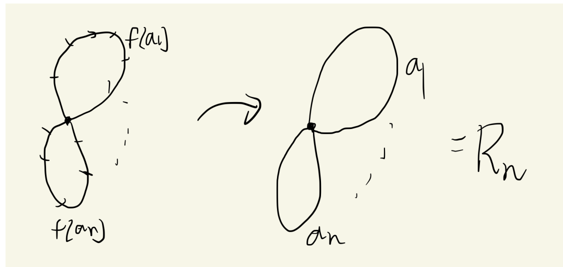

FIGURE 32

implies that it lifts to Γ_k. Also note that all the folds in this factorization are type 1 or type 2, since type 3 and 4 folds cause the rank to go down.

Recall from the proof of Lemma 4.3 that if $\Gamma_i \to \Gamma_{i+1}$ is a type 1 or type 2 fold, we can pick a maximal tree in Γ_i and Γ_{i+1} so that the map $F_n \to \pi_1(\Gamma_i) \to \pi_1(\Gamma_{i+1}) \to F_n$ is the identity (where the first and last maps are obtained by collapsing the maximal trees we have picked for Γ_i and Γ_{i+1}) respectively. Doing this for every consecutive pair of maps, we obtain

$$\begin{tikzcd}
\pi_1(\Gamma_0) \arrow[rr] & & \pi_1(\Gamma_1) \arrow[rr] \arrow[dl] & & \pi_1(\Gamma_2) \arrow[r] \arrow[dl] & \dots \arrow[rr] & & \pi_1(\Gamma_k) \arrow[r] \arrow[dl] & \pi_1(R_n) \\
F_n \arrow[u] & F_n \arrow[r] & F_n \arrow[u] & F_n \arrow[r] & F_n \arrow[u] & & F_n \arrow[r] & F_n \arrow[u] &
\end{tikzcd}$$

where all the horizontal maps in the second row are change of maximal tree automorphisms and the maps $F_n \to \pi_1(\Gamma_i) \to \pi_1(\Gamma_{i+1}) \to F_n$ are the identity. □

We now discuss $Out(F_2)$. Let T be a torus and Σ be the punctured torus $T - p$ so $\pi_1(\Sigma) = F_2$. Recall that

$$MCG(\Sigma) \hookrightarrow Out(F_2) \twoheadrightarrow GL(2, \mathbb{Z}).$$

This sequence exists in any rank, but we will show that in rank 2 the maps are isomorphisms.

Lemma 9.3. *The map $MCG(\Sigma) \hookrightarrow Out(F_2)$ is surjective.*

Proof. Let us represent Σ by drawing the torus as the square with edges with identified and by removing the point corresponding to the four vertices. By Nielsen's theorem, it suffices to show that $MCG(\Sigma)$ contains the signed permutation automorphisms and the map ϕ_1 defined by

$$a_1 \mapsto a_1, \qquad a_2 \mapsto a_1 a_2.$$

The signed permutations are generated by rotation by $\frac{\pi}{2}$ and reflection. For example, reflecting vertically yields the map $a_1 \mapsto a_1$ and $a_2 \mapsto a_2^{-1}$, as indicated in Figure 33. The map ϕ_1 is induced by the induced action of

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix},$$

on $\Sigma = (\mathbb{R}^2 - \mathbb{Z}^2)/\mathbb{Z}^2$ as indicated in Figure 34. □

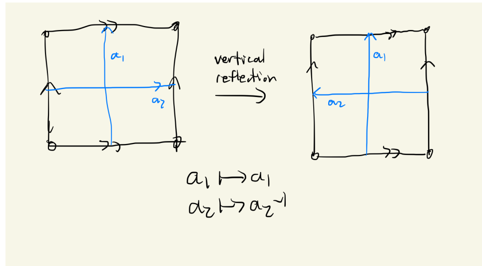

FIGURE 33

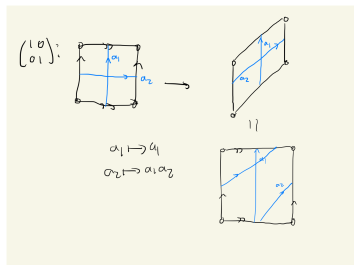

FIGURE 34. This figure has a typo: the bottom left entry of the matrix should be 1.

Lemma 9.4. *The composite $MCG(\Sigma) \to GL(2, \mathbb{Z})$ is injective.*

Proof. Take $[\varphi] \in MCG(\Sigma)$ which maps to the identity in $GL(2, \mathbb{Z})$. The map $\varphi : T - \{p\} \to T - \{p\}$ is induced by a map $\overline{\varphi} : T \to T$. Its action on $\pi_1(T) = \mathbb{Z}^2$ is the identity. Then by [Hat02, Proposition 1B.9], $\overline{\varphi}$ is homotopic to the identity, so $[\varphi]$ is the identity. □

Corollary 9.5. *$Out(F_2) \twoheadrightarrow GL(2, \mathbb{Z})$ is an isomorphism.*

Definition 9.6. Let $IA_n \subseteq Out(F_n)$ be the kernel of $Out(F_n) \to GL(n, \mathbb{Z})$. This group is referred to as the "identity on abelianization" or the "Torelli" subgroup.

We have just shown $IA_2 = \{id\}$. This is far from the case with IA_3.

Example 9.7. The map on F_3 given by $a_1 \mapsto a_2a_1a_2^{-1}, a_2 \mapsto a_2, a_3 \mapsto a_3$ is not inner and is in IA_3.

In fact, IA_3 is not finitely presented, and whether this is true for $IA_n, n \geqslant 4$ remains an open question.

The situation for higher genus surfaces is also quite different.

Example 9.8. Let Σ be the punctured genus 2 surface. We can visually represent Σ as an octagon with its edges identified missing its center.

It turns out that every element of $MCG(\Sigma)$ sends the loop $c = [a_1, b_1][a_2, b_2]$ to a conjugate of c, but there exists $\varphi \in Aut(\pi_1(\Sigma))$ where $\varphi(c)$ is not conjugate to c. Therefore $MCG(\Sigma) \hookrightarrow Out(\pi_1(\Sigma))$ is not onto.

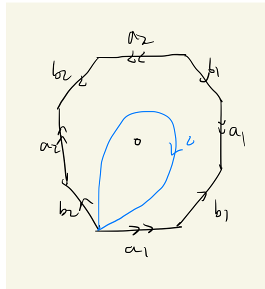

FIGURE 35

10. Introduction to outer space (09/20, SK, KL)

Today, we'll be introducing Outer Space (a space that $\mathrm{Out}(F_n)$ acts on). We continue to follow resources of Bestvina, now additionally using [Bes14].

Definition 10.1. A **metric** on a graph Γ is a function $\mathrm{Edges}(\Gamma) \to (0, \infty)$. The image of an edge is called its **length**, and the sum of the lengths is called the **volume.**

Remark 10.2. If Γ has k edges, the space of metrics on Γ is $(0, \infty)^k$. The volume 1 metrics comprise an open simplex $\{(\ell_1, \dots, \ell_k) \in \mathbb{R}^k : \sum_{i=1}^k \ell_i = 1, \ell_i > 0\}$.

The diagram below shows examples of these spaces where $k = 2$ or $k = 3$ (note that the latter is missing its vertices and edges).

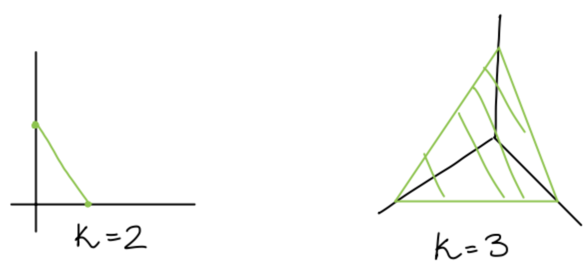

FIGURE 36. Examples of open simplices

Eventually, Outer Space will be put together from simplices as in the above diagram.

Definition 10.3. A **marking** of a graph Γ is a homotopy equivalence $f : R_n \to \Gamma$. An **inverse marking** is a homotopy equivalence $g : \Gamma \to R_n$.

We now think about how we might get inverse markings in practice.

Lemma 10.4. *Let Γ be a graph such that $\pi_1(\Gamma) \cong F_n$, and let T be a maximal tree of Γ. Orient the edges of $\Gamma - T$ and label them by a basis $w_1, \dots, w_n$ of F_n. Then the map $g : \Gamma \to R_n$ that collapses T and maps the remaining edges according to their labels is an inverse marking.*

See the example below of an inverse marking, where the colors indicate the map from Γ to R_2.

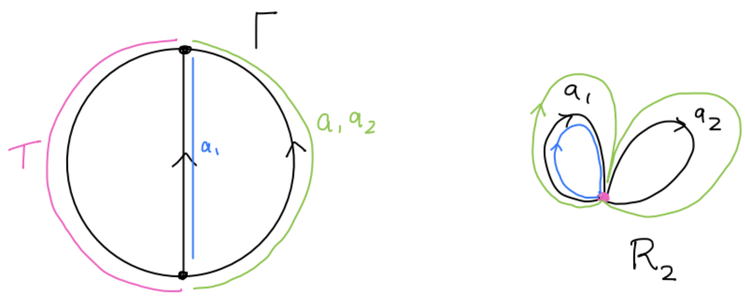

FIGURE 37. Example of an inverse marking

We now prove Lemma 10.4. Recall that Whitehead's theorem tells us, heuristically, that if a map between reasonable spaces induces isomorphisms on homotopy groups, then the map is a homotopy equivalence.

Proof of Lemma 10.4. Pick any basepoint v of Γ (note that the choice of basepoint doesn't matter). Then $f_* : \pi_1(\Gamma, v) \to \pi_1(R_n, f(v))$ is an isomorphism because a basis of $\pi_1(\Gamma, v)$ maps to a basis of $\pi_1(R_n)$. Whitehead's theorem tells us that f is a homotopy equivalence. □

Exercise 10. *Prove that two maps $f, g : (\Gamma, v) \to (\Gamma', v')$ between graphs are homotopic if and only if $f_* = g_*$. (This will be proved in class next lecture.)*

Example 10.5. Let w_1, w_2 be a basis for F_2. Consider 3 maximal trees, T_1, T_2, T_3, of the theta graph:

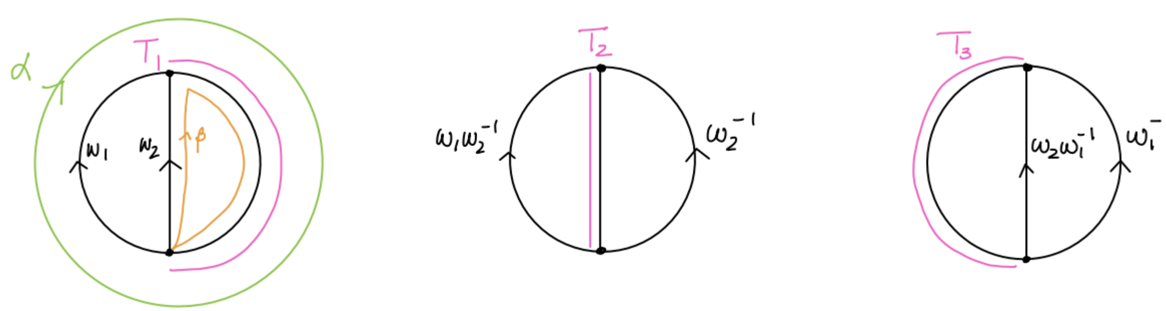

FIGURE 38. Example of homotopic inverse marking

Then the maps to R_2 described in the above diagram all describe homotopic inverse markings. We can use the previous exercise to check this. Note that green $[\alpha]$ and orange $[\beta]$ generate $\pi_1(\Gamma)$. In all three cases,

$$[\alpha] \mapsto w_1$$
$$[\beta] \mapsto w_2.$$

We now define outer space.

Definition 10.6. The (Culler-Vogtmann) **outer space $\mathbf{CV}_n$** is defined as

$$\{(\Gamma, \ell, h) : \Gamma \text{ is a finite graph with all vertices of valence} \geqslant 3, \quad \ell \text{ is a metric on } \Gamma \text{ of volume } 1,\ h : \Gamma \to R_n \text{ is an inverse marking}\}/\sim,$$

where $(\Gamma, \ell, h) \sim (\Gamma', \ell', h')$ if there exists an isometry $\rho : \Gamma \to \Gamma'$ such that $h' \circ \rho$ is freely homotopic (i.e. no basepoint is preserved) to h:

$$\begin{tikzcd} \Gamma \arrow[rr, "\rho"] \arrow[dr, "h"'] & & \Gamma' \arrow[dl, "h'"] \\ & R_n & \end{tikzcd}$$

Remark 10.7. The outer space CV_n has a nice topology that we will define soon.

10.1. $\mathbf{CV}_2$.

Example 10.8. There are 3 graphs with all vertices of valence $\geqslant 3$ and $\pi_1 \cong F_2$:

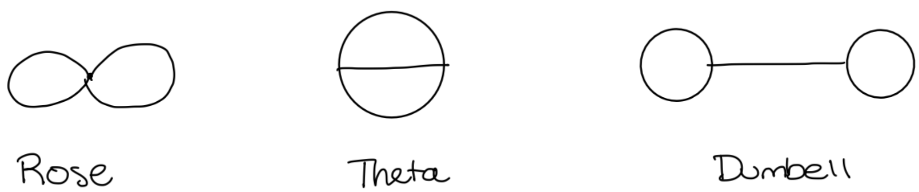

FIGURE 39. Graph with all vertices of valence at least 3

Say we have a marked theta graph:

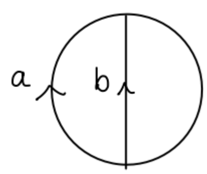

FIGURE 40. Marked theta graph

This gives an open 2-simplex in CV_2:

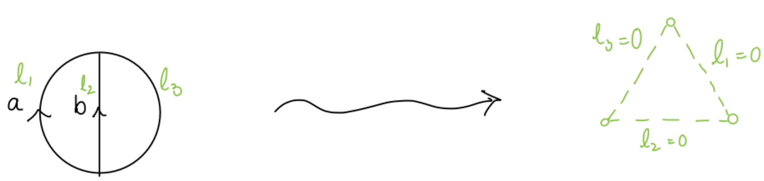

FIGURE 41. Simplex of marked theta graph

We'll now give a "warm up" discussion of what might be going on at the boundary of this open simplex. Since we haven't even defined the topology, this will necessarily be a heuristic discussion, but hopefully it will help to motivate a more rigorous discussion afterwards.

For example, as $\ell_3 \to 0$, the above theta graph becomes the rose on the left in the figure below, which corresponds to the pink point in the open 2-simplex:

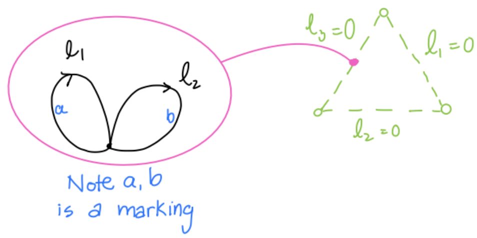

FIGURE 42. A graph that lives on this 2-simplex

In particular, each point of the open 1-simplex $\ell_3 = 0$ lying on boundary of the open 2-simplex corresponds to a rose with edges of length ℓ_1 and ℓ_2 with $\ell_1 + \ell_2 = 1$. (And similarly for each point of the other two open 1-simplices on the boundary of the open 2-simplex.) Interior points (ℓ_1, ℓ_2, ℓ_3) of the 2-simplex correspond to theta graphs with edges of lengths $\ell_1, \ell_2, \ell_3 > 0$ respectively. Below is a diagram that shows where a few different graphs lie on the 2-simplex.

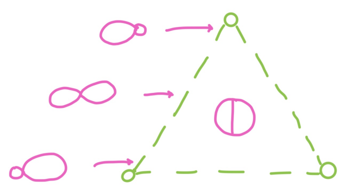

FIGURE 43. More graphs that live on this 2-simplex

Finally, note that we do not include the vertices, since if we collapse 2 edges of the θ graph, the resulting graph does not have fundamental group isomorphic to F_2. More formally, a sequence of triples (ℓ_1, ℓ_2, ℓ_3) where $\ell_1 \to 0$ and $\ell_2 \to 0$ cannot converge in CV_2; it goes to ∞.

Example 10.9. Consider a dumbbell:

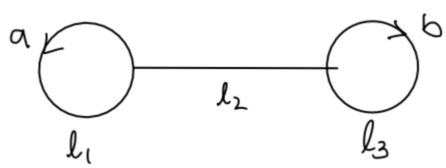

FIGURE 44. Dumbbell graph

Its simplex has only one open edge in its closure, as the limit as $\ell_2 \to 0$ exists in CV_2, but the limits as $\ell_1 \to 0$ and $\ell_3 \to 0$ does not.

Figure 45. The simplex of the dumbbell graph

11. CV_2, markings (09/22, YM, UP)

In the last lecture examples of (open) simplices in CV_2 were given. In this lecture, additional discussion of CV_2 will be given, and more tools related to markings will be provided.

Recall that for theta graphs it is valid to collapse any one edge while keeping the fundamental group intact, which gives a 2-simplex with 3 of its (open) edges present in outer space. We can't collapse more than one edge without changing π_1, so the 2-simplex has the vertices missing. For dumbbell graphs, only one edge of the 2-simplex is present in outer space, as collapsing the loops will kill a generator and result in a different fundamental group. Therefore for each theta graph it is connected to three edges, but each dumbbell graph is only connected to one edge. Every edge of CV_2 corresponds to a marked rose, and is attached to three 2-simplices as follows:

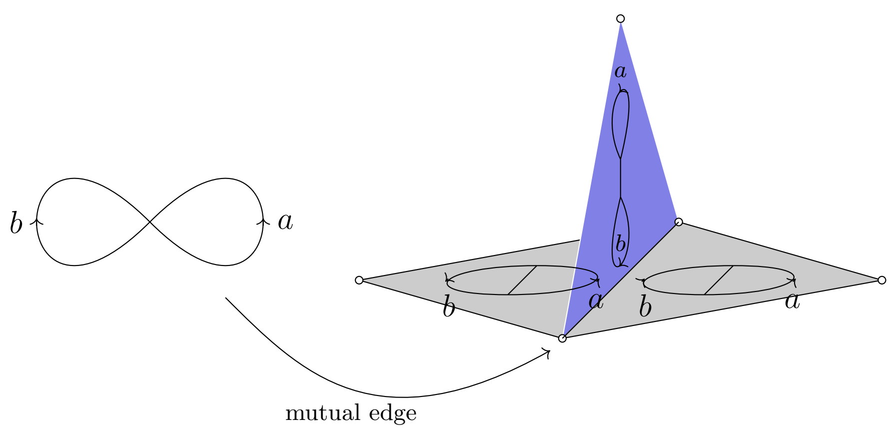

Figure 46. Every edge in CV_2 is adjacent to three 2-simplices

None of the vertices are in this complex. Notice that two of the edges in the purple triangle representing the dumbbell graph are not included in the space.

Definition 11.1. The **Reduced Outer Space (of rank 2)** is CV_2 without the dumbbell graphs.

Remark 11.2. Outer space is **not** a manifold.

Reduced CV_2 looks like a 3-regular tree of triangles. Another common representation of CV_2 is using "distorted" triangles inscribing on a circle, as follows:[1]

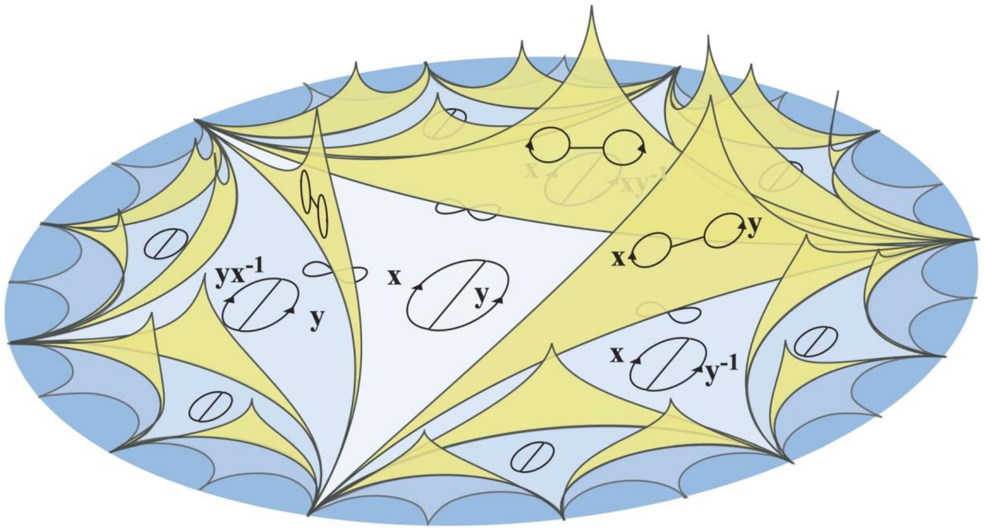

FIGURE 47. CV_2

The only intersection of these triangles with the circle are their vertices, which we have shown are not in CV_2. Any point in the disc represents a graph with a certain marking in reduced CV_2.

Remark 11.3. Although reduced CV_2 is a manifold, reduced outer space in higher rank is **not** a manifold.

We now turn to an interlude on markings, first proving a special case of [Hat02, Proposition 1B.9]:

Lemma 11.4. *Let $f, g : (\Gamma, x) \to (\Gamma', x')$ be morphisms of finite graphs with $f_* = g_*$. Then f and g are homotopic through maps sending x to x'.*

Proof. Let $u : (\hat{\Gamma}, \hat{x}) \to (\Gamma, x)$ and $u' : (\hat{\Gamma}', \hat{x}') \to (\Gamma', x')$ be the universal covers. Then $f \circ u$ and $g \circ u$ can be lifted to maps $\hat{f}$ and $\hat{g}$:

$$\begin{tikzcd} (\hat{\Gamma}, \hat{x}) \arrow[r, shift left, "\hat{f}"] \arrow[r, shift right, "\hat{g}"'] \arrow[d, "u"] & (\hat{\Gamma}', \hat{x}') \arrow[d, "u'"] \\ (\Gamma, x) \arrow[r, shift left, "f"] \arrow[r, shift right, "g"'] & (\Gamma', x') \end{tikzcd}$$

For $\alpha \in \pi_1(\Gamma, x)$ and $\rho \in \hat{\Gamma}$, both $\hat{f}, \hat{g}$ satisfy

$$\hat{f}(\alpha\rho) = f_*(\alpha)\hat{f}(\rho), \qquad \hat{g}(\alpha\rho) = f_*(\alpha)\hat{g}(\rho)$$

[1] This image is retrieved from Vogtmann's article [Vog08].

Then define a straight line homotopy from $\hat{f}$ to $\hat{g}$,

$$\hat{H} : \hat{\Gamma} \times [0,1] \to \hat{\Gamma}',$$

by $H(p,t)$ travelling along the geodesic from $\hat{f}(p)$ to $\hat{g}(p)$ with constant speed parameterized by $[0,1]$. Due to equivariance, we get an induced map $H : \Gamma \times [0,1] \to \Gamma'$ with $H_0 = f$ and $H_1 = g$, which is the desired homotopy. $\square$

Recall that for a given map $s : [0,1] \to X$, we get an isomorphism on the fundamental group $I_s : \pi_1(X, s(0)) \to \pi_1(X, s(1))$ by extending loops:

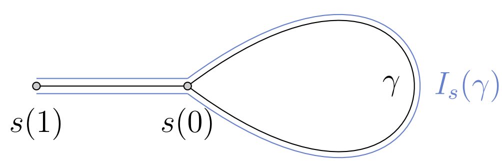

FIGURE 48. Isomorphism Characterized by Homotopy

If $s(0) = s(1)$, then this results in a conjugation by s. Recall also that the universal cover of (X, x) can be defined as paths starting at x modulo homotopy rel base points. Therefore s could also give a homeomorphism between the model of the universal cover based at $s(0)$ and the universal cover based at $s(1)$.

Lemma 11.5. *Let $f, g : \Gamma \to \Gamma'$ be maps of graphs. Suppose there exists a continuous map $s : [0,1] \to \Gamma'$ s.t. $s(0) = f(x), s(1) = g(x)$; and suppose that $g_* = I_s \circ f_*$, i.e. the following diagram commutes. Then f and g are homotopic.*

$$\begin{tikzcd}
 & \pi_1(\Gamma, x) \arrow[dl, "f"'] \arrow[dr, "g"] & \\
\pi_1(\Gamma', f(x)) \arrow[rr, "I_s \ \cong"] & & \pi_1(\Gamma', g(x))
\end{tikzcd}$$

Exercise 11. *Prove the above lemma. (See Hint A.6.)*

Lemma 11.6. *Let (Γ, v) be a graph and (X, x) a space. Then for any $\rho : \pi_1(\Gamma, v) \to \pi_1(X, x)$ there exists a $f : (\Gamma, v) \to (X, x)$ with $f_* = \rho$.*

Proof. This is clear via collapsing a maximal tree $T \subset \Gamma$, and sending each petal α to a loop in $\rho(\alpha)$. □

Lemma 11.7 (Whitehead's Theorem for Graphs)**.** *Let $f : (\Gamma, v) \to (\Gamma', v')$ be a graph morphism with $f_* : \pi_1(\Gamma, v) \to \pi_1(\Gamma', v')$ an isomorphism. Then f is a homotopy equivalence.*

Proof. Pick $g : (\Gamma', v') \to (\Gamma, v)$ with $g_* = (f_*)^{-1}$. It suffices to show that $f \circ g$ and $g \circ f$ are homotopic to id. This follows from Lemma 11.6 as $f_* \circ g_* = g_* \circ f_* = \mathrm{id}$, i.e. they induce the identity map in π_1. □

Remark 11.8. The results above show that homotopy equivalences on pointed graphs are in bijection to isomorphisms from π_1 to F_n.

12. Markings, the topology on CV_n (09/25, SZ, SY)

Definition 12.1. Let $f, g : \Gamma \to \Gamma'$. Fix a base point $v \in \Gamma$. We say

$$f_* : \pi_1(\Gamma, v) \to \pi_1(\Gamma', f(v))$$

and

$$g_* : \pi_1(\Gamma, v) \to \pi_1(\Gamma', g(v))$$

are *conjugate* if there exists a path s such that

$$g_* = I_s \circ f_*.$$

Remark 12.2. If $g(v) = f(v)$, the maps are conjugate if and only if there exists $\alpha \in \pi_1(\Gamma', f(v))$ such that $g_* = \alpha f_* \alpha^{-1}$.

In general, if we fix some arbitrary arc s_0, f_* and g_* are conjugate if and only if there exists an α such that $g_* = \alpha(I_{s_0} \circ f_*)\alpha^{-1}$.

Lemma 12.3. *f and g are homotopic if and only f_* and g_* are conjugate.*

Proof. If s exists, this follows from Lemma 11.5.

Conversely, if there exists homotopy

$$H : \Gamma \times [0, 1] \to \Gamma'$$

then take $s(\cdot) = H(v, \cdot)$. See Figure 49. □

Exercise 12. *Finish this proof.*

Corollary 12.4. *The following is a bijection:*

$$\{\textit{inverse markings } \Gamma \to R_n\}/\textit{homotopy} \longleftrightarrow \{\textit{isomorphisms } \pi_1(\Gamma) \to F_n\}/\textit{conjugation}$$

The bottom line is that one can use identifications of $\pi_1(\Gamma)$ with F_n up to conjugacy instead of markings as we defined them, if desired.

Now, we go over some basics on the topology of CV_n.

Proposition 12.5. *Fix a graph Γ and an inverse marking h. The map from the open simplex of metrics to CV_n defined by*

$$\left[\ell : E\Gamma \to (0, \infty), \sum \ell(e) = 1\right] \mapsto (\Gamma, \ell, h) \in CV_n$$

is injective.

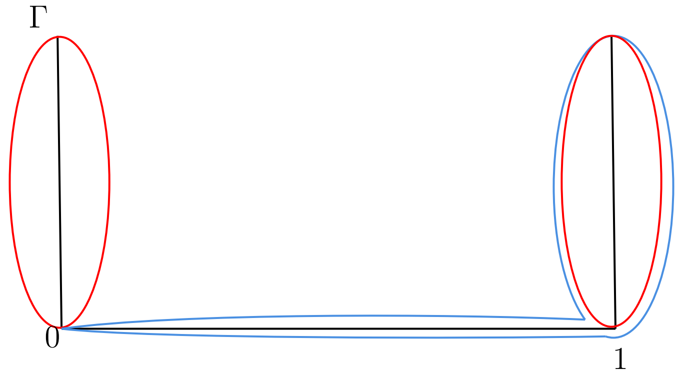

FIGURE 49. Intuitively, the red curve is homotopic to the blue curve

Proof. Say $\ell \neq \ell'$ but $(\Gamma, \ell, h) \simeq (\Gamma, \ell', h)$. There exists isometry

$$\rho : (\Gamma, \ell) \to (\Gamma, \ell')$$

such that $h \circ \rho \simeq h$. Every free homotopy class in Γ contains a unique immersed loop. Since $h \circ \rho \simeq h$, ρ must map each (oriented) immersed loop to itself, and the result follows from the following exercise. □

Exercise 13. *Let Γ be a finite graph with all vertices of valence $\geqslant 3$. Let $\rho : \Gamma \to \Gamma$ be a graph isomorphism that takes each (oriented) immersed loop to itself. Then $\rho = id$. (See Hint A.7.)*

Definition 12.6. For (Γ, h) a graph with an inverse marking, let $\Sigma(\Gamma, h)$ be the set of functions $\ell : E\Gamma \to [0, \infty)$ such that $\sum \ell(e) = 1$, and the set of edges e with $\ell(e) = 0$ is a forest.

Remark 12.7. $\Sigma(\Gamma, h)$ is a union of open simplices of different dimensions. There is a natural map $\Sigma(\Gamma, h) \to CV_n$, obtained by collapsing the forest.

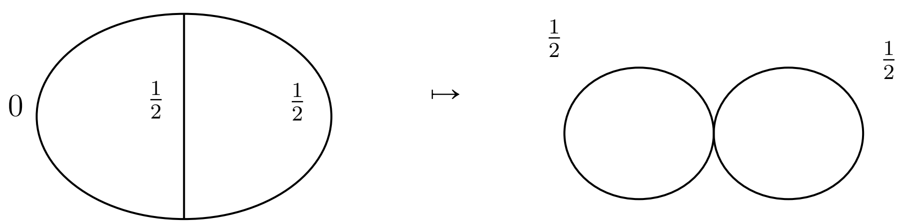

FIGURE 50. Example of collapsing a forest: this is a homotopy equivalence, so a marking of the left graph induces a marking of the right graph

Lemma 12.8. *This map $\Sigma(\Gamma, h) \to CV_n$ is injective.*

Proof. $\Sigma(\Gamma, h)$ is a union of open simplices, each of which maps injectively to CV_n. So it suffices to show that the images of different simplices are disjoint. This follows from the following exercise. □

Exercise 14. *Let F_1, F_2 be distinct forests in Γ. Then there exists an immersed loop α with $\alpha \cap F_1$ and $\alpha \cap F_2$ having different number of edges. (This means, in the simplices corresponding to collapsing F_1 and to collapsing F_2, the immersed loop α has a different number of edges. See Hint A.8 and Figure 51.)*

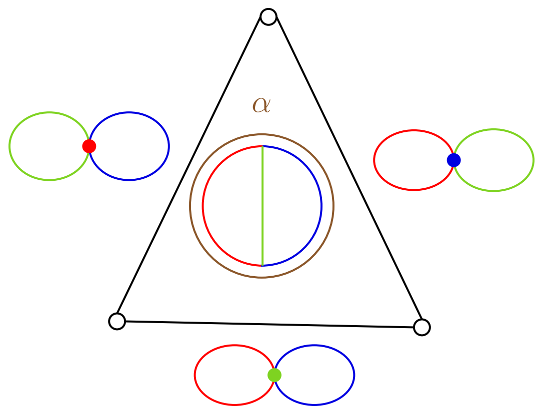

FIGURE 51. Example: Each edge (1-simplex) of the triangle corresponds to collapsing one of the three edges in the "Θ graph". In this picture, the image of brown loop α has either one or two edges in different 1-simplices.

Definition 12.9. The topology on CV_n is defined so $U \subseteq CV_n$ is open if its preimage in each $\Sigma(\Gamma, h)$ is open.

From the above, we have come very close to showing there exists a simplicial complex X and a subcomplex X_0 such that CV_n is homeomorphic to $X - X_0$. That is true, and we will take it for granted going forward.

13. THE FAREY GRAPH (09/27, SY, SL)

For this lecture we won't follow any particular source, but one reference is [Hat22].

Definition 13.1. $\begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{Z}^2$ is called *imprimitive* if $\exists\ m > 1,\ m \in \mathbb{Z}$, $\begin{pmatrix} a_0 \\ b_0 \end{pmatrix} \in \mathbb{Z}^2$, with $\begin{pmatrix} a \\ b \end{pmatrix} = m \begin{pmatrix} a_0 \\ b_0 \end{pmatrix}$. Otherwise, it is called *primitive*. Equivalently, $\begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{Z}^2$ is primitive if and only if $\gcd(|a|, |b|) = 1$. (Recall: $\gcd(n, 0) = 0,\ \forall\ n \in \mathbb{N}$.)

Remark 13.2. If $g \in \mathrm{GL}(2, \mathbb{Z})$, then $\begin{pmatrix} a \\ b \end{pmatrix}$ is primitive $\Leftrightarrow g \begin{pmatrix} a \\ b \end{pmatrix}$ is primitive.

Exercise 15. $\begin{pmatrix} a \\ b \end{pmatrix}$ *is primitive* $\Leftrightarrow \exists\ c, d \in \mathbb{Z}$ *with* $\begin{pmatrix} a & c \\ b & d \end{pmatrix} \in \mathrm{GL}(2, \mathbb{Z})$. *(Hint: Euclidean algorithm.)*

Remark 13.3. Primitive elements of $\mathbb{Z}^2$ mod negation are in bijection to $\mathbb{Q} \cup \{\infty\}$ via

$$\pm \begin{pmatrix} a \\ b \end{pmatrix} \mapsto \frac{a}{b}.$$

Instead of primitive elements, people often speak instead about fractions in lowest terms, with $\frac{0}{1}, \frac{1}{1}, \frac{1}{0}$ being considered in lowest terms.

Definition 13.4. The *Farey graph* is the graph with one vertex for each primitive element of $\mathbb{Z}^2$ mod negation, and an edge from $\pm\begin{pmatrix} a \\ b \end{pmatrix}$ to $\pm\begin{pmatrix} c \\ d \end{pmatrix}$ if $\det\begin{pmatrix} a & c \\ b & d \end{pmatrix} = \pm 1$.

Remark 13.5. Often one uses $\mathbb{Q} \cup \{\infty\}$ to label the vertices.

Suppose we have an edge from $\pm\begin{pmatrix} a \\ b \end{pmatrix}$ to $\pm\begin{pmatrix} c \\ d \end{pmatrix}$, then F contains the following:

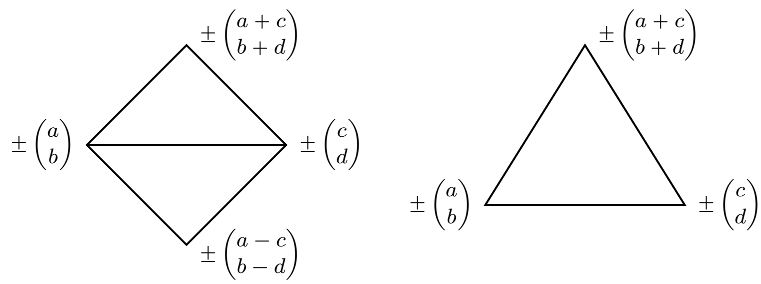

FIGURE 52. Left: two triangles adjacent at an edge. Right: Just the top triangle.

We can justify the inclusion of the edges in this figure by noting that

$$\begin{pmatrix} a & a+c \\ b & b+d \end{pmatrix} = \begin{pmatrix} a & c \\ b & d \end{pmatrix}\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} a & a-c \\ b & b-d \end{pmatrix} = \begin{pmatrix} a & c \\ b & d \end{pmatrix}\begin{pmatrix} 1 & 1 \\ 0 & -1 \end{pmatrix}.$$

Remark 13.6. The top triangle is more symmetric than it looks, since

$$\begin{pmatrix} c \\ d \end{pmatrix} = \begin{pmatrix} a+c \\ b+d \end{pmatrix} + \begin{pmatrix} -a \\ -b \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} a+c \\ b+d \end{pmatrix} + \begin{pmatrix} -c \\ -d \end{pmatrix}.$$

So each vector can be written as a sum of the other two if we pick suitable signs. This also holds for the bottom triangle.

Note also, that there are 4 sums of the form

$$\pm\begin{pmatrix} a \\ b \end{pmatrix} \pm \begin{pmatrix} c \\ d \end{pmatrix},$$

so 2 after moding out by negation. Both of these are represented in our local picture above.

Definition 13.7. A triangle (aka. 3-cycle) in F is *good* if it is of the form of the left triangle in Figure 53.

Remark 13.8. By definition, each edge is part of 2 good triangles.

Remark 13.9. Later we will see all triangles are good.

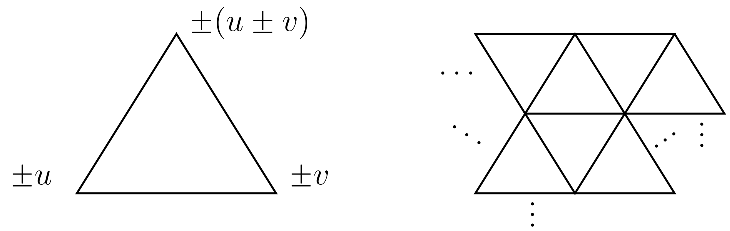

FIGURE 53. Left: A good triangle. Right: a collection of adjacent good triangles.

Big picture idea: We want to understand F by "exploring" using triangles.

Remark 13.10. Every vertex is a part of an edge. If $\begin{pmatrix} a \\ b \end{pmatrix}$ is primitive, then there exists $\begin{pmatrix} c \\ d \end{pmatrix}$ such that

$$\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = \pm 1.$$

Hence, we get and edge from $\pm \begin{pmatrix} a \\ b \end{pmatrix}$ to $\pm \begin{pmatrix} c \\ d \end{pmatrix}$. (See Exercise 14.)

Lemma 13.11. *For any two edges, $e_1, e_2 \in E(F)$, there exists a sequence of good triangles $T_1, T_2, \cdots, T_n$, such that*

- *e_1 is an edge of T_1,*
- *e_2 is an edge of T_n, and*
- *T_i and T_{i+1} share an edge.*

Proof sketch. Say the vertices of e_1 are $\begin{pmatrix} a_1 \\ b_1 \end{pmatrix}$ and $\begin{pmatrix} c_1 \\ d_1 \end{pmatrix}$ and the vertices of e_2 are $\begin{pmatrix} a_2 \\ b_2 \end{pmatrix}$ and $\begin{pmatrix} c_2 \\ d_2 \end{pmatrix}$.

We can find a sequence

$$M_1, M_2, \ldots, M_k \in \left\{ \begin{pmatrix} 1 & \pm 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ \pm 1 & 1 \end{pmatrix}, \begin{pmatrix} 0 & \pm 1 \\ \pm 1 & 0 \end{pmatrix}, \begin{pmatrix} \pm 1 & 0 \\ 0 & \pm 1 \end{pmatrix} \right\}$$

such that

$$\begin{pmatrix} a_2 & c_2 \\ b_2 & d_2 \end{pmatrix} = \begin{pmatrix} a_1 & c_1 \\ b_1 & d_1 \end{pmatrix} \cdot M_1 \cdot M_2 \cdots M_k.$$

This follows from the general fact that $GL(n, \mathbb{Z})$ is generated by elementary matrices, but also keep in mind that the $n = 2$ case of this fact used here is easier than the $n > 2$ case.

Keeping in mind that

$$\begin{pmatrix} a & a+c \\ b & b+d \end{pmatrix} = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

etc, we get the result. $\square$

Remark 13.12. If we use $\mathbb{Q} \cup \{\infty\}$ instead of primitive vectors, good triangles look like:

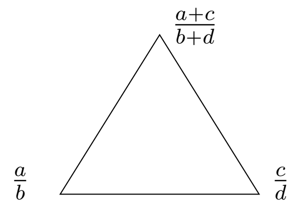

FIGURE 54

Lemma 13.13. *If there is an edge from* $\pm\begin{pmatrix} a \\ b \end{pmatrix}$ *to* $\pm\begin{pmatrix} c \\ d \end{pmatrix}$ *and* $\frac{a}{b} < \frac{c}{d}$, $b, d \geqslant 0$, *then*

$$\frac{a+c}{b+d} \in \left(\frac{a}{b}, \frac{c}{d}\right) \quad ; \quad \frac{a-c}{b-d} \notin \left(\frac{a}{b}, \frac{c}{d}\right).$$

Proof. For example

$$\frac{a+c}{b+d} > \frac{a}{b} \Leftrightarrow ad - bc < 0,$$

and $ad - bc = bd\left(\frac{a}{b} - \frac{c}{d}\right) < 0$. Similar estimates prove the remaining statements. (For some estimates one needs to consider cases depending on the sign of $b - d$.) $\square$

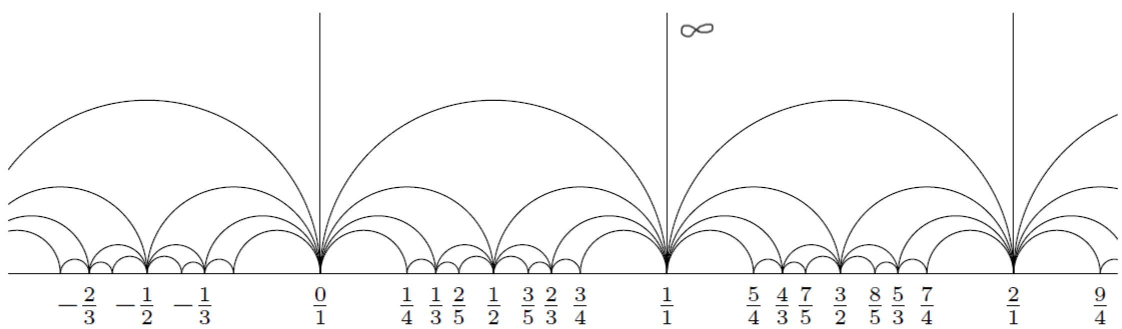

FIGURE 55. Farey graph on upper half plane

Corollary 13.14. *If we draw F with vertices in $\mathbb{Q} \cup \{\infty\}$, and edges as semicircles in the upper half plane, then the edges do not cross and all complementary regions are triangles.*

See Figure 55.[2]

Proof. We know every edge e is part of 2 triangles. The last lemma says they are on opposite sides of e.

Start with one triangle and iteratively add adjacent good triangles. Lemma 13.11 says you get all of F in this way. $\square$

[2]This image is retrieved from Dr Ian Short's webpage.

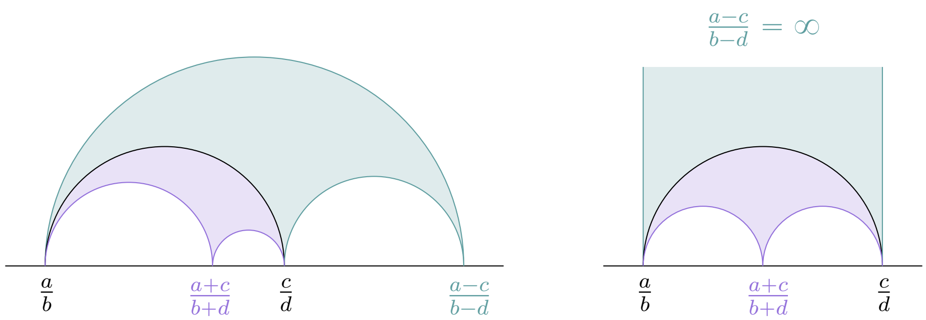

FIGURE 56. Adjacent triangles lie on opposite sides of the shared edge

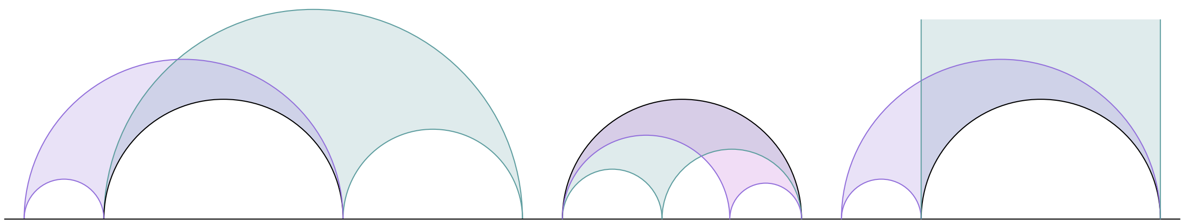

FIGURE 57. Adjacent triangles cannot overlap, so the pictures in this figure never occur

If you use stereographic projection, one gets the different but equivalent picture shown in Figure 58.[3]

Corollary 13.15. *The following are true:*

- *All triangles are good.*
- *F is planar.*
- *The dual graph is the 3-regular tree.*

Remark 13.16. All vertices have ∞ valences.

14. CV_2, cv_2, and the Farey complex (09/29, YL, ZH)

Recall the set-up of the Farey graph. Its vertices are primitive elements $\begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{Z}^2$ and there is an edge between $\begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{Z}^2$ and $\begin{pmatrix} c \\ d \end{pmatrix} \in \mathbb{Z}^2$ if $\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = \pm 1$. We view the primitive elements as fractions, up to a sign, and imagine them as points on the circle as in Figure 58.

Definition 14.1. The **Farey Complex** $\mathcal{F}_\Delta$ is $\mathcal{F}$ with its triangles glued in.

Thus $\mathcal{F}_\Delta$ is a simplicial complex of dimension 2. The above results give the following, where $(\mathcal{F}_\Delta)_0$ denotes the zero skeleton of $\mathcal{F}_\Delta$, i.e. the set of all its vertices:

Corollary 14.2. $\mathcal{F}_\Delta - (\mathcal{F}_\Delta)_0$ *is homeomorphic to an open disk.*

[3]This image is retrieved from Ian Agol's paper.

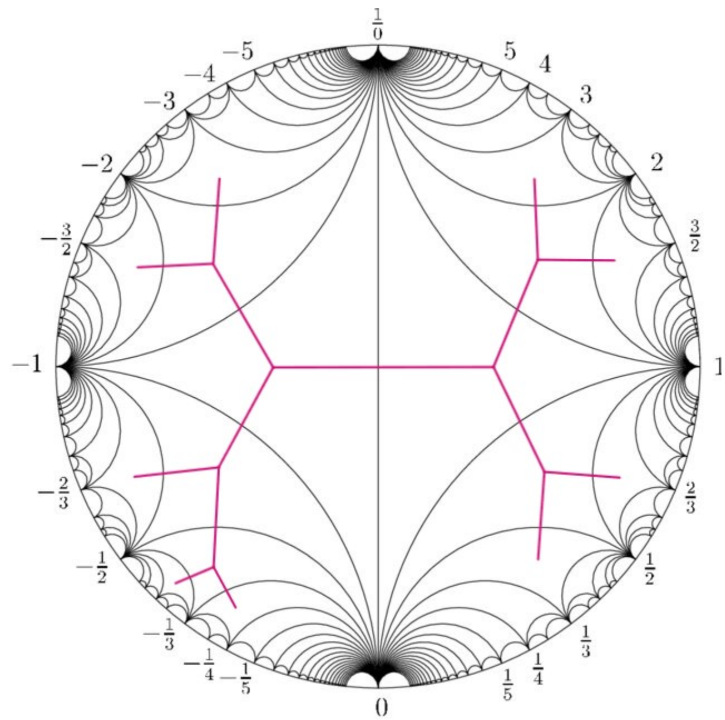

FIGURE 58. Farey graph and dual graph on unit disk

Definition 14.3. Reduced outer space cv_n is the subset of CV_n where the graph has no separating edges.

Remark 14.4. $cv_2 = CV_2 \backslash \{\text{dumbbells}\}$.

In general, we know the following about cv_n:

Lemma 14.5. *cv_n is a deformation retract of CV_n.*

Proof. First note that the set of separating edges is a forest. The deformation retract is just linearly contracting each edge of the forest, while rescaling the lengths of the remaining edges to keep the volume constant. See Figure 59. □

Our goal is to prove the following, which finally provides a rigorous justification for the picture of CV_2 presented in Figure 47.

Theorem 14.6. *cv_2 is homeomorphic to $\mathcal{F}_\Delta \backslash (\mathcal{F}_\Delta)_0$.*

To prove this theorem, we will use the abelianization map $\text{Ab} : F_2 \to \mathbb{Z}^2$, and, crucially, the fact that $\text{Out}(F_2) \to GL(2, \mathbb{Z})$ is an isomorphism. We begin by relating bases of F_2 to bases of $\mathbb{Z}^2$.

Definition 14.7. A **basis of** $\mathbb{Z}^2$ is a pair $v, w \in \mathbb{Z}^2$ such that the map $\mathbb{Z}^2 \to \mathbb{Z}^2$ sending $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ to v and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ to w is an isomorphism. Equivalently, every point in $\mathbb{Z}^2$ is a linear combination of v, w with integer coefficients.

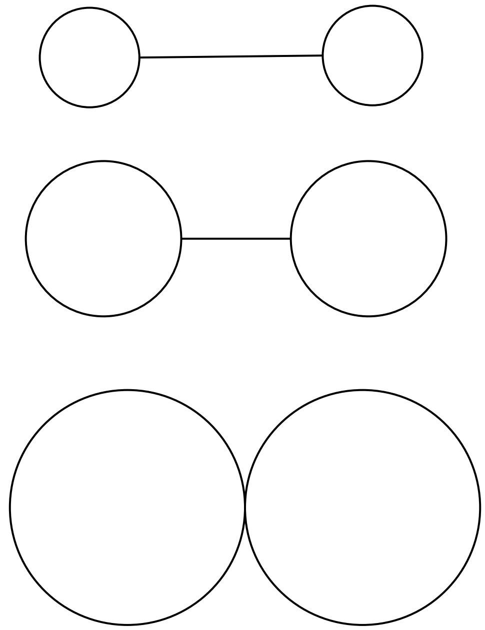

FIGURE 59

Remark 14.8. The pair $\begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} c \\ d \end{pmatrix}$ forms a basis if and only if $\det \begin{pmatrix} a & c \\ b & d \end{pmatrix} = \pm 1$, since the inverse of $\begin{pmatrix} a & c \\ b & d \end{pmatrix}$ in $GL(2, \mathbf{Q})$ is

$$\frac{1}{\det \begin{pmatrix} a & c \\ b & d \end{pmatrix}} \begin{pmatrix} d & -c \\ -b & a \end{pmatrix}.$$

Lemma 14.9. *The map Ab induces a bijection between conjugacy classes of bases of F_2 and bases of $\mathbb{Z}^2$.*

Proof. Surjectivity follows from $Out(F_2) \to GL(2, \mathbb{Z})$ being surjective and the fact that there is only one $GL(2, \mathbb{Z})$-orbit of a basis of $\mathbb{Z}^2$.

To show injectivity, suppose v, w and v', w' are two bases of F_2, with $\mathrm{Ab}(v) = \mathrm{Ab}(v')$ and $\mathrm{Ab}(w) = \mathrm{Ab}(w')$. Define $\phi \in \mathrm{Aut}(F_2)$ by $\phi(v) = v'$ and $\phi(w) = w'$. Then $\mathrm{Ab}(\phi) = Id \in GL(2, \mathbb{Z})$. Since $\mathrm{Out}(F_2) \cong GL(2, \mathbb{Z})$, we know that ϕ is an inner automorphism, i.e. there exists a $g \in F_2$ such that for all $x \in F_2$, we have $\phi(x) = gxg^{-1}$. Therefore, the two bases are g conjugates of each other. □

Definition 14.10. A **primitive element** of F_n is an element contained in some basis of F_n.

Remark 14.11. Primitive elements of F_2 map to primitive elements of $\mathbb{Z}^2$. For example, a_1^2 and $a_1 a_2 a_1^{-1} a_2^{-1}$ are not primitive elements.

Lemma 14.12. *The map* Ab *induces a bijection between conjugacy classes of primitive elements of F_2 and primitive elements of $\mathbb{Z}^2$.*

Proof. Surjectivity follows from the proof of the previous lemma.

To show injectivity, let $v, v' \in F_2$ be primitive, with $Ab(v) = Ab(v')$. First we can find $w, w' \in F_2$ such that v, w and v', w' are two bases of F_2. Define $\phi \in Aut(F_2)$ such that $\phi(v) = v'$ and $\phi(w) = w'$. Without loss of generality, assume $v = a_1$ and $w = a_2$. Then we have $\mathrm{Ab}(v) = \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \mathrm{Ab}(v')$. In addition, we have

$$\mathrm{Ab}(\phi) = \begin{pmatrix} 1 & n \\ 0 & \pm 1 \end{pmatrix}$$

for some $n \in \mathbb{Z}$ since $\det(\mathrm{Ab}(v'), \mathrm{Ab}(w')) = \pm 1$.

A preimage of $\mathrm{Ab}(\phi)$ is the map $\psi \in \mathrm{Aut}(F_2)$ sending a_1 to a_1 and a_2 to $a_2^{\pm 1} a_1^n$. Since $\mathrm{Out}(F_2) \cong GL(2, \mathbb{Z})$, we have $\phi = g\psi g^{-1}$, for some $g \in F_2$. In addition, since $\psi(v) = v$ and $\phi(v) = v'$, we have

$$v' = \phi(v) = g\psi(v)g^{-1} = gvg^{-1},$$

and so v and v' are conjugate. □

We warn however that many elements of F_2 are not primitive but map to primitive elements of $\mathbb{Z}^2$.

Exercise 16. *Show $a_1^2 a_2^3$ is not primitive, even though its image in $\mathbb{Z}^2$ is $(2, 3)$, which is primitive. (See Hint A.9.)*

We have now characterized the vertices of $\mathcal{F}$ in terms of F_2. Next, we characterize the edges.

Lemma 14.13. *Let $\mathcal{C}_1, \mathcal{C}_2$ be conjugacy classes of primitive elements in F_2. Then $\mathrm{Ab}(\mathcal{C}_1)$ and $\mathrm{Ab}(\mathcal{C}_2)$ are connected in $\mathcal{F}$ if and only if there exist $v_1 \in \mathcal{C}_1$ and $v_2 \in \mathcal{C}_2$ such that v_1, v_2 is a basis of F_2.*

Proof. If v_1, v_2 is a basis of F_2, then $\mathrm{Ab}(v_1) = \mathrm{Ab}(\mathcal{C}_1), \mathrm{Ab}(v_2) = \mathrm{Ab}(\mathcal{C}_2)$ is a basis of $\mathbb{Z}^2$ by Lemma 14.9. Thus they are connected in the Farey graph.

Conversely, suppose that $\mathrm{Ab}(\mathcal{C}_1)$ is connected to $\mathrm{Ab}(\mathcal{C}_2)$ in $\mathcal{F}$. Then $\mathrm{Ab}(\mathcal{C}_1), \mathrm{Ab}(\mathcal{C}_2)$ is a basis of $\mathbb{Z}^2$. By Lemma 14.9, there is a basis v_1, v_2 of F_2 such that the conjugacy class of v_i is $\mathcal{C}_i$ for $i = 1, 2$. □

Remark 14.14. In summary, we can alternatively define $\mathcal{F}$ as the graph with a vertex for each primitive element of F_2 up to conjugacy and inversion, and an edge if you can make a basis with one element from each vertex.

To draw the connection between cv_2 and the Farey graph, it's important to first recall from Example 10.5 that if w_1, w_2 be a basis for F_2, the following markings are all homotopic:

Key Observation: the edge labels in Figure 60 are in $\{w_1^{\pm 1}, w_2^{\pm 1}, (w_1 w_2^{-1})^{\pm 1}\}$. The abelianization of these elements are the vertices of a Farey triangle:

We can label each of the three missing vertices of the simplex of metrics by one of

$$\{w_1^{(-1)^k}, w_2^{(-1)^k}, (w_1 w_2^{-1})^{(-1)^k}\},$$

based on which one has length approaching 0 as you approach that missing vertex. Now map the simplex to $\mathcal{F}_\Delta$ by mapping the missing vertices to the corresponding

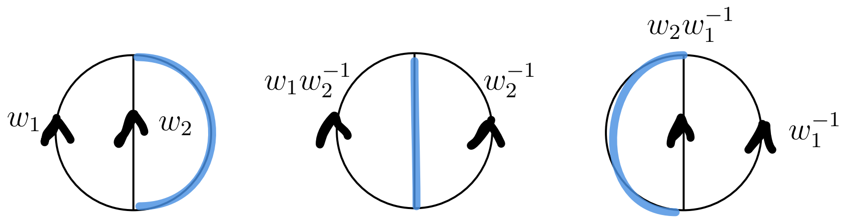

FIGURE 60

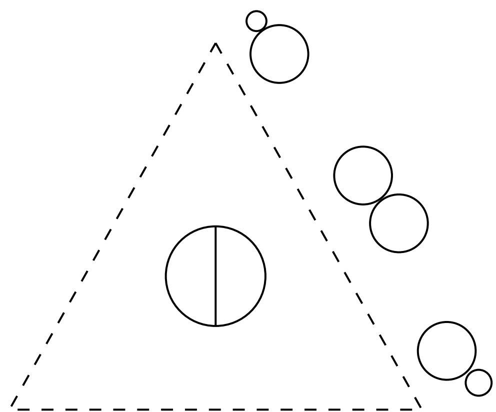

FIGURE 61

vertices of $\mathcal{F}_\Delta$ and extend linearly. This gives a map $cv_2 \to \mathcal{F}_\Delta \backslash (\mathcal{F}_\Delta)_0$. This map is a homeomorphism because it is a covering map and $\mathcal{F}_\Delta \backslash (\mathcal{F}_\Delta)_0$ is simply-connected.

15. The action on CV_n (10/02, NL, KL)

Lemma 15.1. *The largest and smallest dimensional simplices in CV_n have dimension $3n-4$ and $n-1$ respectively, and they correspond to the trivalent graphs and roses respectively.*

Proof. If Γ has E edges and V vertices, then the dimension of its cell is $E-1$. We have $V-E = 1-n$, and $1 \leqslant V \leqslant \frac{2E}{3}$. Then $V=1$ and $V=\frac{2E}{3}$ correspond to the smallest and largest dimensional simplices respectively. $\square$

Here is another perspective we can take on this lemma. Observe that

(1) if Γ has an edge e that isn't a loop, contracting e gives a graph corresponding to a smaller simplex;
(2) if Γ is not trivalent, it has a vertex of valence at least 4. It can be replaced with an edge in multiple ways, giving a new graph with larger simplex. See Figure 62.

Remark 15.2. Thus, codimension 1 simplices are expected to be adjacent to 3 maximum dimensional simplices. In Figure 62, if the original graph on the left is in cv_n and v is a non-separating vertex, all 3 are in cv_n.

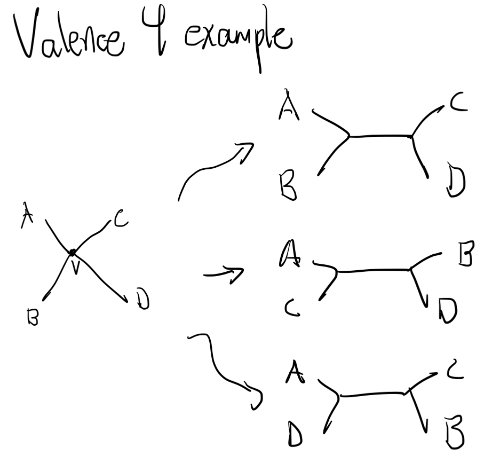

FIGURE 62

Exercise 17. CV_n *is locally finite (i.e., each simplex is in the boundary of only finitely many others).*

Now we turn our attention to the action of $\mathrm{Out}(F_n)$. The results we proved on markings imply that $\mathrm{Out}(F_n)$ is isomorphic to the group of homotopy equivalences of R_n. With this implicit, an element $\varphi \in \mathrm{Out}(F_n)$ acts on CV_n by $[(\Gamma, l, h)] \mapsto [(\Gamma, l, \varphi \circ h)]$. Alternatively, we can view a marking as an isomorphism $\pi_1(\Gamma) \to F_n$ up to conjugation, and the action of φ is by post-composition with φ.

Lemma 15.3. *Let* $(\Gamma, l, h) \in \mathrm{CV}_n$. *The map* $\mathrm{Isom}(\Gamma, l) \to \mathrm{Out}(F_n)$ *defined by* $\varphi \mapsto h_* \circ \varphi_* \circ h_*^{-1}$ *is an isomorphism onto* $\mathrm{Stab}(\Gamma, l, h)$.

This map is just considering the action on π_1 of the graph, and using h to translate that to F_n.

Proof. By Exercise 13, the map is injective. To see that the image lies in $\mathrm{Stab}(\Gamma, l, h)$, note that $h_*\varphi_*h_*^{-1}(\Gamma, l, h) = (\Gamma, l, \psi)$ where $\psi_* = h_*\varphi_*h_*^{-1}h_* = h_*\varphi_*$. But then

$$\begin{tikzcd} \Gamma \arrow[d, "h"'] & \Gamma \arrow[l, "\varphi"'] \arrow[dl, "\psi"] \\ R_n & \end{tikzcd}$$

commutes up to homotopy because it commutes on π_1. Therefore, $[(\Gamma, l, h)] = [(\Gamma, l, \psi)]$ in CV_n, and so $h_*\varphi_*h_*^{-1} \in \mathrm{Stab}(\Gamma, l, h)$. □

Exercise 18. *Check surjectivity.*

Corollary 15.4. *The action of* $Out(F_n)$ *on* CV_n *has finite stabilizers.*

Exercise 19. *Show* $Out(F_n) \curvearrowright CV_n$ *has no kernel if* $n \geqslant 3$. *If* $n = 2$, *then the kernel is the subgroup* $\langle a_1 \mapsto a_1^{-1}, a_2 \mapsto a_2^{-1} \rangle \cong \mathbb{Z}/2\mathbb{Z}$.

Exercise 20. *Show there are only finitely many orbits of simplices in* CV_n*.*

The following result, often called the "Nielsen Realization Theorem," is important, but we will not prove it. See [Besb, Lecture 4] for a proof using a non-trivial result of Stallings.

Theorem 15.5. *Every finite subgroup of* $\mathrm{Out}(F_n)$ *fixes a point in* CV_n*.*

Exercise 21. *The group* $\mathrm{Out}(F_n)$ *has only finitely many conjugacy classes of finite subgroups.*

Proposition 15.6. *The group* $\mathrm{Out}(F_n)$ *is virtually torsion-free.*

(Recall that a group is said to virtually have property P if the group has a subgroup of finite index with property P.)

Proof. We will show the kernel of the composite map

$$\mathrm{Out}(F_n) \to \mathrm{GL}(n, \mathbb{Z}) \to \mathrm{GL}(n, \mathbb{Z}/3\mathbb{Z})$$

is torsion-free. Suppose that φ is in this kernel, and is torsion. Then $\langle\varphi\rangle$ is a finite subgroup of $\mathrm{Out}(F_n)$, and so by Theorem 15.5, it fixes a point $[(\Gamma, l, h)]$ in CV_n. It follows from Lemma 15.3 that φ is represented by an isometry of the graph Γ. Since φ is in the kernel of the composite map, this isometry acts trivially on $H_1(\Gamma, \mathbb{Z}/3\mathbb{Z})$. But the map $\{\text{oriented embedded cycles}\} \to H_1(\Gamma, \mathbb{Z}/3\mathbb{Z})$ is injective, and non-trivial isometries cannot preserve all embedded loops by the proof of Exercise 13. □

Our next goal is to sketch a proof of the contractibility of CV_n.

Theorem 15.7. CV_n *is contractible.*

We will also define the spine of CV_n, a contractible subset that produces a compact space when we quotient out by it.

16. Introduction to greedy folding (10/04, UP, YW)

Our goal now is to prove that CV_n is contractible. We'll show that CV_n deformation retracts to the simplex of the rose R_n. The trajectories to this simplex will be via "greedy folding paths".

Definition 16.1. A continuous map $\phi : (\Gamma, l) \to (\Gamma', l')$ between metric graphs is *linear on edge* (or just *linear* for short) if:

(1) Γ and Γ' have all vertices of valence $\geqslant 3$.
(2) ϕ_* is an isomorphism on π_1.
(3) for each each edge of (Γ, l) there exists $\sigma > 0$ such that the edge, parameterised by unit speed, maps to an immersed (i.e., non-backtracking) path in (Γ', l') of speed σ.

If the slope of all edges are the same we will say ϕ is *constant slope*.

Remark 16.2. ϕ does not have to be a graph morphism.

Remark 16.3. $\phi : (\Gamma, l) \to (\Gamma', l')$ is constant slope of slope σ if $\phi : (\Gamma, \sigma l) \to (\Gamma', l')$ is constant slope of slope 1.

Remark 16.4. ϕ is constant slope of slope 1 iff ϕ restricted to each edge is a local isometry.

Example 16.5. Let (Γ', l', x') be a pointed metric graph, and fix an isomorphism $\pi_1(\Gamma', x') \cong F_n$. Let l be the metric on the rose R_n so petal i has the length of (the immersed representative of) a_i in Γ' (based at the point x'). Then there is a natural map $(R_n, l) \to (\Gamma', l')$ of slope 1 that sends petal i to the representative of a_i. Figure 63 shows an example with $n = 3$.

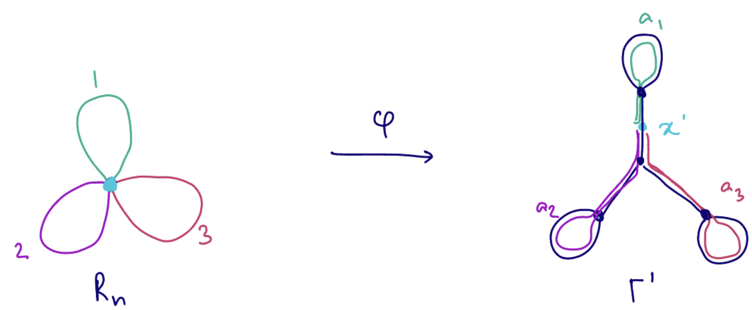

FIGURE 63

Note here that x' need not be a vertex.

Definition 16.6. We will say that a constant slope map ϕ is *foldable* if there are two (small) segments leaving a vertex in Γ with the same image in Γ'.

Example 16.7. In the example drawn in Figure 64, the green segments of the graph on the left represent segments with the same image under a map ϕ and thus ϕ is a foldable map. As illustrated, we can "fold these segments together". (Perhaps a better name would be "zip" rather than "fold", since the folding process is very like zipping up a zipper, but the term "fold" is standard.)

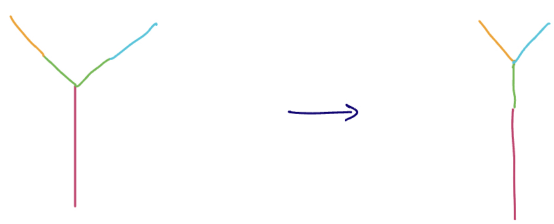

FIGURE 64

Proposition 16.8. *Every constant slope map $\phi : \Gamma \to \Gamma'$ is either a homeomorphism (scale by σ) or is foldable.*

Remark 16.9. This is a version of the fact that "every graph morphism that is an immersion and is surjective on π_1 is a homeomorphism", which applies whenever the target graph is finite and doesn't have vertices of valence 1. (That fact is discussed for example in the proof of Theorem 6.3 and follows from Lemma 5.2.)

We'll see two proofs of this proposition.

First proof of Proposition 16.8. Say Γ has vertex set V, and Γ' has vertex set V'. Define

$$V'_{\text{new}} := V' \cup \phi(V) \subset \Gamma'$$
$$V_{\text{new}} := V \cup \phi^{-1}(V'_{\text{new}}) \subset \Gamma.$$

Note that $\phi(V_{\text{new}}) \subset V'_{\text{new}}$ and $\phi^{-1}(V'_{\text{new}}) \subset V_{\text{new}}$. Declare these to be the new vertex sets; some vertices may have valence 2 but that is ok. With these vertex sets, ϕ is a graph morphism. If ϕ is not foldable, this morphism is an immersion. Since ϕ_* is surjective, ϕ is a homeomorphism. □

Before doing the second proof, we note an easy lemma.

Lemma 16.10. *Let $F : X \to Y$ be a map of sets that is equivariant with respect to actions $G \curvearrowright X$ and $G \curvearrowright Y$ (i.e. $F(gx) = gF(x)$). Let F be the induced map $f : X/G \to Y/G$. Then f injective $\implies$ f is injective.*

Proof. Say F is injective, and suppose

$$[F(x_1)] = f([x_1]) = f([x_2]) = [F(x_2)].$$

By definition, there exists $g \in G$ such that $F(x_1) = gF(x_2) = F(gx_2)$. Injectivity of F implies $x_1 = gx_2$, and thus $[x_1] = [x_2]$. □

We apply this as follows.

Second proof of Proposition 16.8. Say ϕ is not a homeomorphism. Since ϕ_* is surjective, ϕ is surjective. (This is because in a finite graph where every vertex has valence $\geqslant 2$, deleting a vertex reduces the rank of its π_1). So ϕ is not injective.

Consider a lift to universal covers $\tilde{\phi} : \tilde{\Gamma} \to \tilde{\Gamma}'$. Then the lemma implies that $\tilde{\phi}$ is not injective. Pick $x_1, x_2 \in \tilde{\Gamma}$ with $\tilde{\phi}(x_1) = \tilde{\phi}(x_2)$. Since $\tilde{\Gamma}'$ is a tree, the image of the geodesic from x_1 to x_2 in $\tilde{\Gamma}$ must backtrack as shown in Figure 65. This shows the map is foldable. □

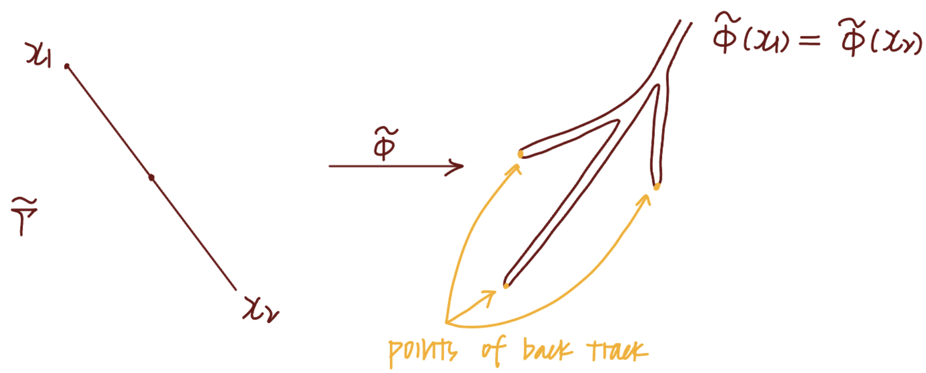

FIGURE 65

Given a map $\phi : (\Gamma, l) \to (\Gamma', l')$ that is foldable, we can fold it "for time ϵ" by identifying all pairs of segments of length ϵ that leave the same vertex and have the same image.

Before we do this carefully, we'll draw some pictures.

Example 16.11. Figure 66 depicts a map $\phi : \Gamma \to \Gamma'$, and segments of the same colour have images of the corresponding colour in Γ'. The third (bottom) graph is obtained after folding for some time ϵ.

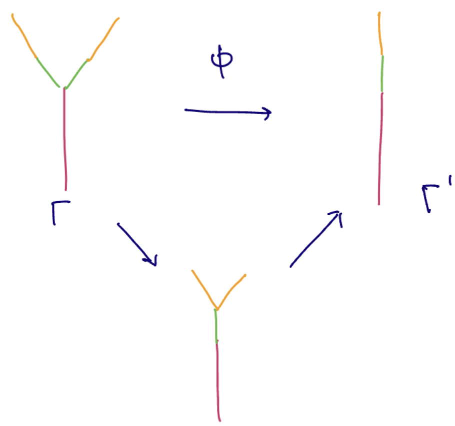

FIGURE 66

If we continue to fold at all vertices with equal speed, "events" occur where the topology of the graph (but not π_1) may change - for example, two valence 3 vertices may be identified and form a valence 4 vertex. After an event, we can "re-start" the process and continue folding if desired.

Example 16.12. Figure 67 shows a map $\phi : \Gamma \to \Gamma'$, as well as various intermediate graphs obtained at different stages of the folding process.

Example 16.13. Figure 68 depicts another example where folding changes the topology of a graph Γ.

17. Greedy folding paths (10/06, YW, HT)

Definition 17.1. Given a constant slope map (with slope 1) $\Phi : (\Gamma, \ell) \to (\Gamma', \ell')$, the first interval of the greedy folding is a path of metric graphs Γ_t, $t \in [0, t_1]$ where

(1) $\Gamma_0 = \Gamma$,
(2) the map Φ factors through each Γ_t,
(3) $\forall t < t_1$, and for all ϵ small enough depending on t, $\Gamma_{t+\epsilon}$ is obtained from Γ_t by, at every vertex, identifying all collections of segments of length ϵ of edges leaving that vertex when these segments have the same image in Γ', and
(4) t_1 is either the first time the topology changes or the time where $\Gamma_{t_1} = \Gamma'$.

Exercise 22. *Check that this is well defined, in that there is a unique path with these properties. (See Hint A.10.)*

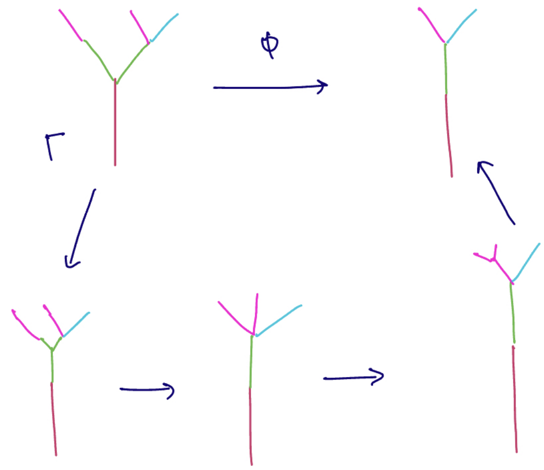

FIGURE 67

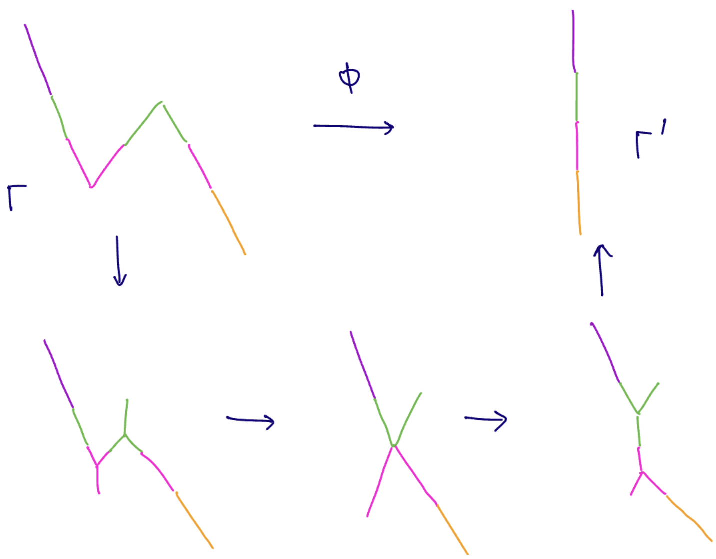

FIGURE 68

Remark 17.2. If $s < t$, $vol(\Gamma_s) - vol(\Gamma_t) \geqslant t - s$, so $t_1 \leqslant vol(\Gamma) - vol(\Gamma')$. In words, the volume is decreasing at a certain rate depending on how much folding is occurring, so in particular this process can't go on for too long.

Definition 17.3. The greedy folding path is defined iteratively via the above definition. Once the path is defined up to time t_n, we do one more 'interval of folding', traveling across some simplex until the topology changes.

An example of greedy folding can be seen in Figure 69.

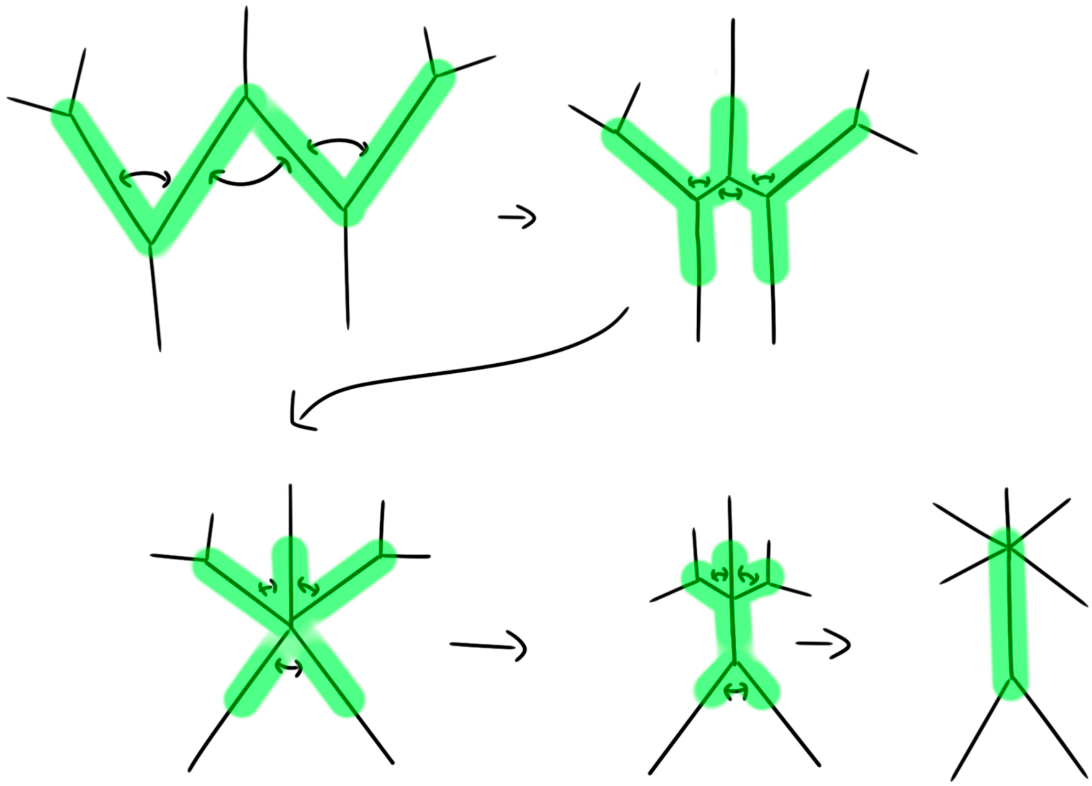

FIGURE 69

Lemma 17.4. *This iterative process terminates after finitely many steps. In particular, we get a genuine path from* Γ *to* Γ'*, passing through finitely many simplices (so the topology of the partially-folded graph changes only finitely many times).*

A quick disclaimer: to get the path to end of Γ', we might need to chop off hanging trees as we go. This always represents a simple deformation retraction, so isn't such a big deal for us.

Here's a sketch of the proof: by subdividing edges (as in last class), we can assume $\Phi : \Gamma \to \Gamma'$ is a morphism. Consider the poset $\mathcal{P}$ of all graphs obtained by doing Stallings (discrete) folds to this morphism. The poset structure is defined such that performing an additional Stallings fold on a graph in $\mathcal{P}$ always produces a graph that is a successor within the poset. Necessarily $\mathcal{P}$ is finite.

At each time t, we consider an arbitrary maximal element $\tilde{\Gamma}$ of the poset $\mathcal{P}_t \in \mathcal{P}$ of elements of $\mathcal{P}$ such that the morphism $\Gamma \to \Gamma_t$ factors through the morphism $\Gamma \to \tilde{\Gamma}$. We should think of this poset of containing all elements from which we can get to Γ_t with some additional (continuous) folding.

In other words, we have

$$\begin{tikzcd} \Gamma \arrow[r] \arrow[d] & \Gamma_t \\ \tilde{\Gamma} \arrow[ur] & \end{tikzcd}$$

Moreover, there does not exist a pair of edges in $\tilde{\Gamma}$ with same image in Γ_t, since otherwise we could fold that pair and $\tilde{\Gamma}$ would not be maximal.

So Γ_t is $\tilde{\Gamma}$ plus some 'partial folds'. The number of homeomorphism type you can set from partial folds is bounded via an explicit function of the rank of the graph. (See Exercises 27 and 28.) So if the greedy folding construction goes at least k steps, we have

$$k \leqslant (\text{function of the graph rank}) \cdot (\#E(\Gamma') - \#E(\Gamma))$$

Note that $\#(E(\Gamma') - \#E(\Gamma))$ is the number of moves required to totally fold Φ.

A precise statement on the existence of greedy folding paths can be found in [BF14a, Proposition 2.2].

Remark 17.5. Skora gives a uniform descriptions of Γ_t. The idea is that if you have an open, surjective quotient map $q : X \to Y$ with discrete fibers and X, Y are path connected, locally path connected metric spaces, you can try to interpolate between X, Y by first defining equivalence relations $\sim_t$ where $x \sim_t x'$ if $q(x) = q(x')$ and there is a path γ from x to x' with $q(\gamma) \subset \overline{B_t(q(x))}$. If X_t is X modulo $\sim_t$, then varying t from 0 to some finite quantity should give a continuous folding from X to Y 'through' the graphs X_t.

This can be applied to our situation after passing to universal covers, but it takes work to see it defines a folding path. (See Exercise 26.)

This idea is shown in Figure 70.

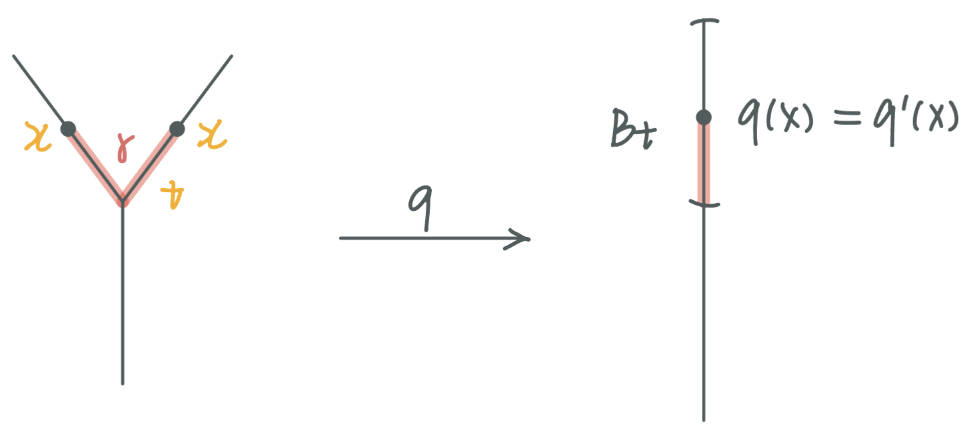

FIGURE 70

We now give some exercises on folding.

Exercise 23. *Suppose $\phi : \Gamma \to \Gamma'$ is constant slope of slope 1, and that all edges of both Γ and Γ' have length 2, and that ϕ is a graph morphism. In both domain and range, add a new vertex in the midpoint of each. Call the old vertices even and the new vertices odd, and now every edge is length 1 and ϕ is still a graph morphism.*

Show that the result Γ_1 of greedy folding by time 1 can be obtained by a collection of Stallings folds based at even vertices. (These are full folds, as we did before talking about CV_n.) In Γ_1, we can define even vertices to be pre-images of even vertices in Γ', and odd vertices to be pre-images of odd vertices of Γ'.

Show that the result Γ_2 of greedy folding $\Gamma_1 \to \Gamma$ by time 1 can be obtained by a collection of Stallings folds based at odd vertices.

Similarly show that, as the process continues, you alternate between folding at even vertices and odd vertices, and at every stage the result can be described by a sequence of Stallings folds. Hence the process must terminate after finitely many times.

Exercise 24. *Suppose that $\psi : \Gamma \to \Gamma'$ is constant slope and slope 1, and all edges have rational lengths. Rescale, subdivide, and use the previous exercise to give a different proof that greedy folding terminates after only finitely many intervals.*

Remark 17.6. It is apparently possible to use the above two exercises and rational approximation to show that greedy folding paths exist, although this proof doesn't immediately tell you as much about the path.

Exercise 25. *Suppose $\phi : \Gamma \to \Gamma'$ is constant slope of slope 1, and Γ_t is the greedy folding path. Consider the image in Γ' of the vertex set of Γ_t; this is a collection of points in Γ' that move around. Show that*

(1) when not colliding with each other or vertices, these points move at speed 1
(2) when a point hits a vertex, it can pass through or continue along a different edge.
(3) when two points hit each other on the middle of an edge, they either have to bounce off each other or pass through each other.

Can you use this to give a different proof that only finitely many events can happen along the folding path?

Exercise 26. *Consider Skora's formula and show the following:*

(1) for tiny times, it coincides with our greedy folding,
(2) applying Skora's formula for time t_1 and then starting again and applying it for time t_2 is the same as applying it for time $t_1 + t_2$,
(3) so for all times it coincides with our greedy folding.

(See Hint A.11.)

Exercise 27. *Consider a graph morphism $m : A \to B$. Fix an edge e in B. WLOG assume e has length 1. Consider the subgraph $m^{-1}(e)$ of A. Consider two oriented edges of this subgraph equivalent if their tail vertices are the same, and if they map to e with the same orientation. So equivalent edges can be folded. Let $G_1, \ldots, G_k$ be the equivalence classes. (Note that each unoriented edge has two orientations, and these won't be equivalent.)*

For $t = (t_1, ..., t_k) \in [0,1]^k$, define the graph A_t to be the result of folding each G_i for time t_i. Show that only finitely many homeomorphism types of graphs occur as A_t for some t, and that the finite bound depends only on k. (See Hint A.12.)

Exercise 28. *Use this to prove the claim above that the number of homeomorphism types of partial folds is bounded by a function of the rank. Hint: since we have subdivided so our map is a graph morphism, the number of edges isn't bounded by a function of the rank. But the number of edges where folding can occur is linearly bounded by the rank. And, since we only consider partial folds, the folding that happens over different edges in the codomain are independent of each other.*

The final ingredient for contractibility is:

Lemma 17.7. *One can continuously pick a point on each graph in CV_n, so that on the simplex of the rose you pick the vertex.*

More formally, there is a bundle $\hat{CV_n} \to CV_n$ with a natural topology, and the lemma indicates this bundle has a section. But we won't be so formal.

Proof. Consider $(\Gamma, \ell, h) \in CV_n$. For each $\gamma \in F_n \cong \pi_1(\Gamma)$, there is a geodesic in $\widetilde{\Gamma}$ called the 'axis', where $\tilde{\Gamma}$ is the universal cover of Γ. (One way to define it is the preimage of the unique immersed loop in $\widetilde{\Gamma}/\langle\gamma\rangle$. For other definitions see [Wil].)

If $\mathrm{Axis}(a_1)$ and $\mathrm{Axis}(a_2)$ overlap, pick the point to be the image in Γ of the midpoint of the overlap as shown in Figure 71.

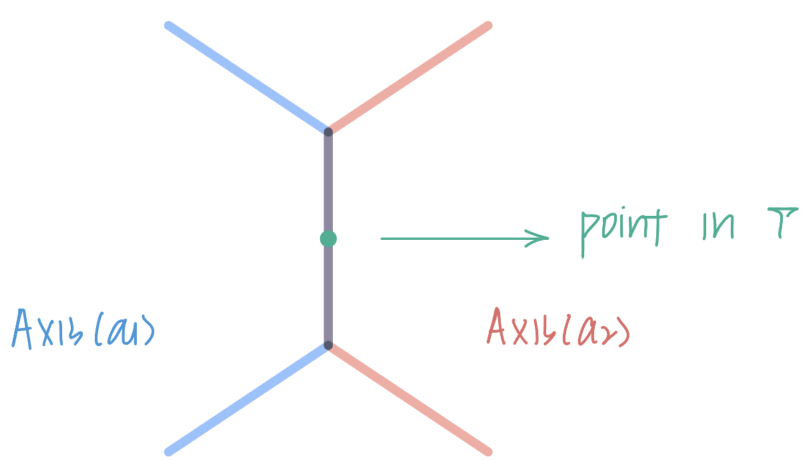

FIGURE 71

Otherwise, use midpoint of geodesic for $\mathrm{Axis}(a_1)$ to $\mathrm{Axis}(a_2)$ as shown in Figure 72. □

Remark 17.8. We only have the identification of $\pi_1(\Gamma)$ and F_n up to conjugacy. You might think that means our use of a_1, a_2 in the previous proof is problematic. But if you replace (a_1, a_2) with (ga_1g^{-1}, ga_2g^{-1}), the axes both move by g, and the midpoint

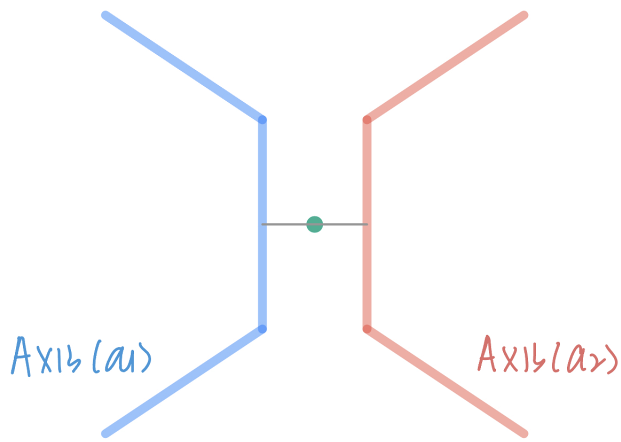

FIGURE 72

m moves by g. But m and $g(m)$ have the same image downstairs in Γ since g is a Deck transformation.

18. CONTRACTIBILITY OF CV_n, SPINE, VCD (10/09, KS, RE)

Now we are going to give a proof sketch of Theorem 15.7. Originally it is due to [CV86], but instead we are roughly following [Sko].

Proof sketch. We will show that CV_n deformation retracts to the simplex of the standard rose.

By Lemma 17.7, we can continuously pick a point in the graph on all of CV_n. Given $(\Gamma, l, h) \in CV_n$, as in Example 16.5 consider a metric on R_n where the loop i has a length equal to the length of the immersed representative of a_i in Γ based at the chosen point (and then normalize to have volume 1). Consider the constant slope map from that rose to Γ. The greedy folding path is a path from (Γ, l, h) to the given rose. It turns out that following this path defines a deformation retract. $\square$

Remark 18.1. It is non-trivial to check continuity. For details, see [Sko] or [Cla05].

Next, we are going to introduce a general construction of spines. Let X be a subcomplex of a simplicial complex $\bar{X}$. Let Y be the union of all simplices of barycentric subdivision of $\bar{X}$ disjoint from X. Then

(1) $\bar{X} - X$ deformation retracts to Y, and
(2) $\dim Y \leqslant \dim \bar{X} - (i+1)$ if the i-skeleton of $\bar{X}$ is in X.

Example 18.2. In Figure 73, $\bar{X}$ is a triangle, X is the set of its vertices, and Y is the spine. Green arrows show the deformation retraction.

Exercise 29. *Give exact formulas for deformation retraction in Example 18.2.*

Definition 18.3. Y is called *spine* of $\bar{X} - X$.

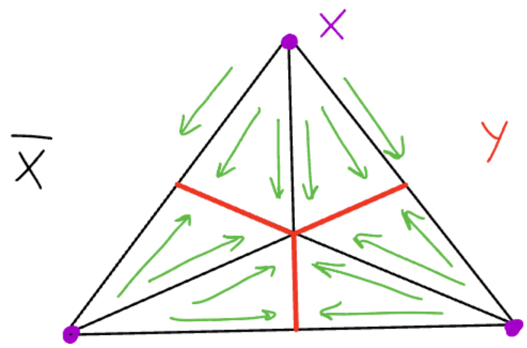

FIGURE 73

Following (2), the spine K_n of CV_n is a simplicial complex of dimension

$$(3n-4)-((n-2)+1)=2n-3.$$

From the contractibility of CV_n and Whitehead's theorem, it follows that its spine K_n is also contractible. The fact that the $\mathrm{Out}(F_n)$ action on CV_n has finitely many orbits of simplices (see Exercise 20) implies that $K_n/\mathrm{Out}(F_n)$ is compact.

If the action was free (i.e. has no fixed points) the quotient would be a finite simplicial complex and a $K(\mathrm{Out}(F_n),1)$ (i.e. the space with fundamental group $\mathrm{Out}(F_n)$ and has a contractible universal cover). This is not quite true, but the fact that $\mathrm{Out}(F_n)$ is virtually torsion free (i.e. there is a finite index subgroup with no torsion) makes it possible to conclude the following.

Corollary 18.4. $\mathrm{Out}(F_n)$ *is finitely presented (i.e. is the fundamental group of a finite simplicial complex).*

Next, we are going to give some facts about *virtual cohomological dimension.* (We won't give a definition though.) Virtual cohomological dimension is a group invariant

$$\mathrm{vcd}:\{\text{groups}\}\to\{0,1,\dots,\infty\} \tag{1}$$

with the following properties:

- $\mathrm{vcd}(\pi_1$ of a finite simplicial complex$)\leqslant k$ if the complex has dimension k and has contractible universal cover;
- If H is a subgroup of G, $\mathrm{vcd}(H)\leqslant \mathrm{vcd}(G)$ with equality if H is finite index in G;
- $\mathrm{vcd}(\mathbb{Z}^n)=n$.

Corollary 18.5. $\mathrm{vcd}(\mathrm{Out}(F_n))=2n-3$.

Remark 18.6. So, $\mathbb{Z}^{2n-2}$ is not a subgroup of $\mathrm{Out}(F_n)$.

Proof. The discussion of K_n shows that $\mathrm{vcd}(\mathrm{Out}(F_n)) \leqslant 2n-3$. Automorphisms of the form

$$\begin{aligned} a_1 &\mapsto a_1, \\ a_2 &\mapsto a_2 a_1^{p_1}, \\ a_3 &\mapsto a_1^{p_2} a_3 a_1^{p_3}, \\ &\vdots \\ a_n &\mapsto a_1^{p_{2n-4}} a_n a_1^{p_{2n-3}}, \end{aligned}$$

where $p_i \in \mathbb{Z}$, show that there is a $\mathbb{Z}^{2n-3}$ subgroup of $\mathrm{Out}(F_n)$. □

Exercise 30. *In general, the spine has one vertex for each open simplex in $\overline{X}$ not in X, and a collection of vertices are contained in a simplex of the spine iff they can be reordered so the open simplices in $\overline{X}$ they correspond have the property that each is in the boundary of the previous one.*

Exercise 31. *Let P_n be the simplicial complex with vertices corresponding to equivalence classes of marked graphs (like CV_n but with no metric), and where a collection of vertices span a simplex iff they can be re-ordered so that each can be obtained from the last via a forest collapse. Show P_n is homeomorphic to the spine K_n of outer space. (This gives a combinatorial description of the spine outer space.)*

19. The asymmetric metric (10/11, SL, SZ)

For the next several lectures we'll follow [Bes14] closely.

Definition 19.1. Given $(\Gamma, \ell, h), (\Gamma', \ell', h') \in CV_n$, a *difference of marking map* is a Lipschitz map $\phi : \Gamma \to \Gamma'$ such that $h' \circ \phi$ is homotopic to h.

$$\begin{tikzcd} \Gamma \arrow[rr, "\phi"] \arrow[dr, "h"'] & & \Gamma' \arrow[dl, "h'"] \\ & R_n & \end{tikzcd}$$

Figure 74. Diagram for difference of marking

Definition 19.2. Let

$$\sigma(\phi) = \sup_{x \neq x'} \frac{d(\phi(x), \phi(x'))}{d(x, x')}$$

be the Lipschitz constant of ϕ.

Definition 19.3. Given $(\Gamma, \ell, h), (\Gamma', \ell', h') \in CV_n$, define

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

where ϕ ranges over difference of marking maps.

Remark 19.4. By abuse of notation, we will write $d(\Gamma, \Gamma')$.

Remark 19.5. By the Arzelà-Ascoli Theorem, we have that a sequence of L-Lipschitz functions has a subsequence that converges to a function that is L-Lipschitz. Hence, the infimum is realized.

Recall from Definition 16.1 that a map $\phi : \Gamma \to \Gamma'$ is *linear on edges* if each edge e maps to an immersed path with constant slope (depending on e).

Remark 19.6. Every $\phi : \Gamma \to \Gamma'$ can be homotoped to a unique linear map with the same values on vertices. We can do this by "tightening" the map on each edge. Before tightening the image of the edge is some arbitrary path, and afterwards it is the unique immersed, constant speed path joining the original endpoints. This process does not increase the Lipschitz constant.

Definition 19.7. A difference of markings $\phi : \Gamma \to \Gamma'$ is called *optimal* if $\log \sigma(\phi) = d(\Gamma, \Gamma')$ and it is linear.

The previous remarks show that an optimal map always exists.

We now show that d is an asymmetric metric, and is moreover $Out(F_n)$ invariant.

Lemma 19.8. *For all $\Gamma_1, \Gamma_2, \Gamma_3 \in CV_n$, we have the following:*

(1) $d(\Gamma_1, \Gamma_3) \leqslant d(\Gamma_1, \Gamma_2) + d(\Gamma_2, \Gamma_3)$;
(2) $d(\Gamma_1, \Gamma_2) \geqslant 0$ with equality if and only if $\Gamma_1 = \Gamma_2$;
(3) If $f \in Out(F_n)$, $d(f\Gamma_1, f\Gamma_2) = d(\Gamma_1, \Gamma_2)$.

Proof. (1): This follows from the general fact

$$\sigma(\psi \circ \phi) \leqslant \sigma(\psi)\sigma(\phi).$$

(2): Let $\phi : \Gamma \to \Gamma'$ be optimal. If $d(\Gamma, \Gamma') < 0$, this contradicts the fact that

$$vol(\Gamma) = vol(\Gamma')$$

and

$$vol(\Gamma') \leqslant \sigma(\phi) vol(\Gamma).$$

Note that ϕ must be surjective. If it's not, ϕ_* can't be surjective, and we have that ϕ is a difference of markings so ϕ_* must be surjective.

If $d(\Gamma, \Gamma') = 0$, then $\sigma(\phi) = 1$. Since $vol(\Gamma) = vol(\Gamma')$, it is constant slope. Since ϕ is a homotopy equivalence, ϕ must be an isometry. ☐

Exercise 32. *Check this and prove Lemma 19.8 (3).*

Definition 19.9. Given a conjugacy class of some $\alpha \in F_n$, let $\ell_\alpha(\Gamma)$ denote the length of the unique immersed loop in the free homotopy class associated to α.

We get a lower bound on distance via:

Lemma 19.10.

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

Proof. This follows from

$$\ell_\alpha(\Gamma') \leqslant \sigma(\phi)\ell_\alpha(\Gamma)$$

for an optimal map ϕ. ☐

Definition 19.11. If $\ell_\alpha(\Gamma') = \sigma(\phi)\ell_\alpha(\Gamma)$, we call α a *witness*.

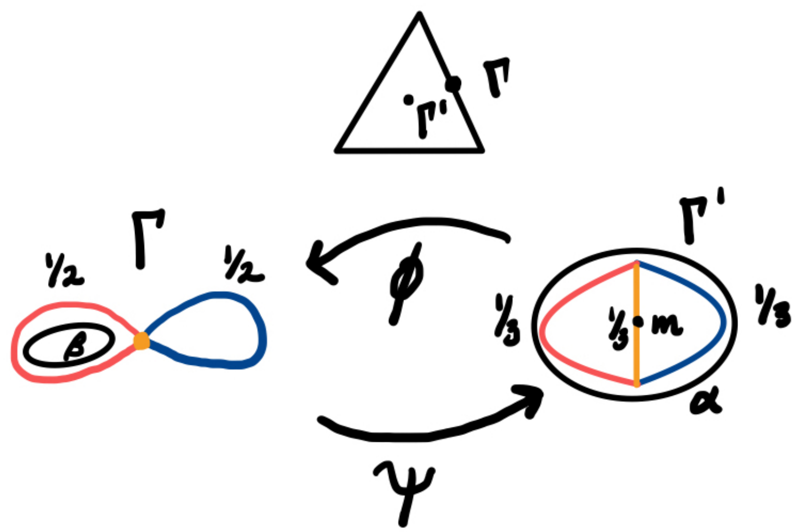

FIGURE 75

Example 19.12. Consider Γ and Γ' in Figure 75, with h given by the identity and h' given by collapsing the middle edge. Consider the linear map $\phi : \Gamma' \to \Gamma$ that collapses the middle edge. Then,

$$\sigma(\phi) = \frac{\frac{1}{2}}{\frac{1}{3}} = \frac{3}{2},$$

so $d(\Gamma', \Gamma) \leqslant \log \frac{3}{2}$. Note that

$$\frac{\ell_\alpha(\Gamma)}{\ell_\alpha(\Gamma')} = \frac{1}{\frac{2}{3}} = \frac{3}{2}.$$

Hence, $d(\Gamma', \Gamma) \geqslant \log \frac{3}{2}$, so we have $d(\Gamma', \Gamma) = \log \frac{3}{2}$.

There is a unique linear map $\psi : \Gamma \to \Gamma'$ that sends the vertex to m. We have,

$$\sigma(\psi) = \frac{\frac{2}{3}}{\frac{1}{2}} = \frac{4}{3},$$

and

$$\frac{\ell_\beta(\Gamma')}{\ell_\beta(\Gamma)} = \frac{\frac{2}{3}}{\frac{1}{2}} = \frac{4}{3}.$$

Hence, $d(\Gamma, \Gamma') = \log \frac{4}{3}$.

Remark 19.13. Note that $d(\Gamma', \Gamma) \neq d(\Gamma, \Gamma')$. d is an "asymmetric metric", which means exactly that it satisfies all the axioms for a metric except symmetry. One could "symmetrize" by considering

$$d_{\text{sym}}(\Gamma, \Gamma') = d(\Gamma, \Gamma') + d(\Gamma', \Gamma).$$

This is a metric but is surprisingly useless.

FIGURE 76

Exercise 33. *Consider C_ϵ in Figure 76. Show that $d(C_\epsilon, C_{1/2})$ goes to infinity as ϵ goes to zero, but $d(C_{1/2}, C_\epsilon)$ is bounded independent of ϵ.*

We conclude that d is *really* not symmetric, and d is not complete in the sense that closed balls are not compact.

Proposition 19.14. *(Witnesses Exist) For all $\Gamma, \Gamma' \in CV_n$, there exists α such that*

$$\frac{\ell_\alpha(\Gamma')}{\ell_\alpha(\Gamma)} = e^{d(\Gamma,\Gamma')}.$$

Moreover, there exists α in Γ homeomorphic to S^1, $S^1 \vee S^1$, or a dumbbell.

Corollary 19.15.

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

with α either ranging over all conjugacy classes of the finitely many loops of the above form.

Remark 19.16. We can algorithmically compute this.

20. WITNESSES EXIST (10/13, YM, SY)

Recall that the asymmetric metric mentioned in the previous lecture:

$$d(\Gamma, \Gamma') := \inf_{\phi:\Gamma\to\Gamma'} \log \sigma(\phi) \overset{*}{\geqslant} \sup \log \frac{\ell_\alpha(\Gamma')}{\ell_\alpha(\Gamma)}$$

where if $(*)$ is an equality then α is called a *witness*. It is also discussed that the infimum is actually a minimum, and that maps which realize the minimal and are linear on edges are called optimal maps.

The main goal of this lecture is to prove the following proposition:

Proposition 20.1. *For any Γ, Γ', a witness exists.*

Definition 20.2. For a point v in Γ (often a vertex), $T_v(\Gamma)$ is defined to be the set of directions at v.

If v is a vertex, then $|T_v(\Gamma)|$ is the valence of v; otherwise $|T_v(\Gamma)| = 2$. Note that $T_v(\Gamma)$ is a version of a unit tangent space, but sometimes it is called just the tangent space at v.

Example 20.3. For the following graph, $|T_{v_1}(\Gamma)| = 3$; and $|T_{v_2}(\Gamma)| = 2$.

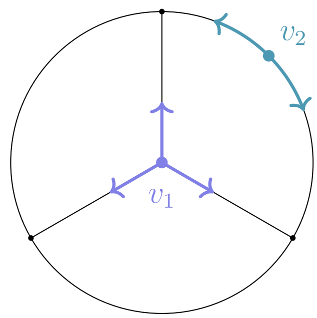

FIGURE 77. T_v of a graph with v being a vertex or on edges

Analogously, we could define "derivative maps" as follows.

Definition 20.4. Let $\phi : \Gamma \to \Gamma'$ be a map of graphs which is linear on edges. Define $\phi_* : T_v\Gamma \to T_{\phi(v)}\Gamma'$ to be the induced map from directions at v to directions at $\phi(v)$.

Definition 20.5. A **train-track structure** on Γ is an equivalence relation on T_v for all vertices $v \in \Gamma$. The equivalence classes of the train-track structure of a vertex are referred to as **gates**.

Remark 20.6. We usually draw directions belonging to the same gate as tangent (or nearly tangent). It's helpful to imagine a railroad switch, where a train can go from a track A into either of two tangent tracks B, C (so A splits into B and C). But a train cannot turn directly from B to C.

Remark 20.7. A graph morphism $\phi : \Gamma \to \Gamma'$ naturally induces a train-track structure, where $d_1 \sim d_2$ if $\phi_*(d_1) = \phi_*(d_2)$.

Remark 20.8. In this train track structure, directions belong to the same gate if they are eligible to be folded together.

Definition 20.9. An immersed path in Γ is **legal** if at each vertex of it the entering gate at it of the path is different from the exiting gate.

Definition 20.10. Let $\phi : \Gamma \to \Gamma'$ be an optimal map. Then subgraph $\Delta = \Delta_\phi$ which is the union of the edges where ϕ has slope $\sigma(\phi)$ is called the **tension graph of** ϕ.

Lemma 20.11. *Let α be an immersed loop in Δ_ϕ that is legal (w.r.t. the train-track structure induced by ϕ). Then $\frac{\ell_\alpha(\Gamma)}{\ell_\alpha(\Gamma')} = \sigma(\phi)$, so α is a witness.*

Proof. This is just re-formalization of the definitions: α being legal implies that $\phi(\alpha)$ is immersed; and $\alpha \subseteq \Delta$ implies that the length of $\phi(\alpha)$ is σ times the length of α, i.e. α is a witness. □

Lemma 20.12. *Let Δ be a graph with a train-track structure s.t. every vertex has at least 2 gates. Then there is a legal loop in one of the following form:*

1) *Embedded loop.*
2) *Embedded loop except 1 point.*
3) *Embedded loop except a segment.*

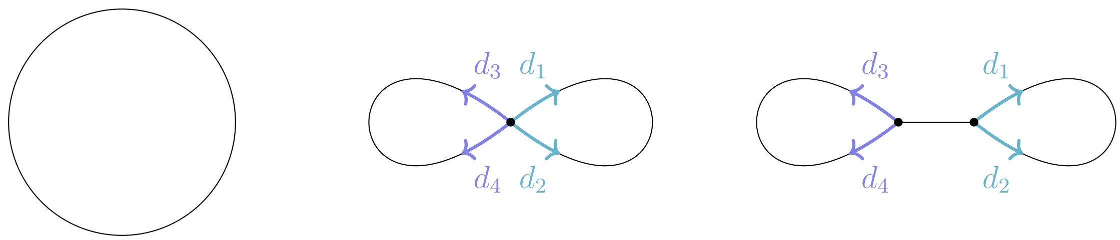

FIGURE 78. Embedded loops in a train-track structure, with $d_1 \sim d_2$, $d_3 \sim d_4$

Sketch of Proof. Start with an oriented edge e_1. Given oriented edge e_i, we define the oriented edge e_{i+1} as follows.

a) If there is an oriented edge $e_k, k < i$ that is leaving the tail of e_i and is in a different gate from e_i, we define $e_{i+1} = e_k$ and we terminate the construction.
b) If there is an oriented edge $e_k, k < i$ that is entering the tail of e_i and is in a different gate from e_i, we define e_{i+1} to be e_k with the opposite orientation.
c) Otherwise we pick e_{i+1} to be any oriented edge that is leaving the tail of e_i and is in a different gate from e_i.

By construction the path is legal. If we terminate after pick e_{i+1} to be $e_k, k < i$, then $(e_k, e_{k+1}, \ldots, e_i)$ is a legal loop. Since every vertex in this graph has at least two gates and the graph is finite, by pigeonhole there must exist a legal loop, i.e. the process above terminates.

One can check that this loop is always described by one of the three cases. $\square$

Remark 20.13. A different approach to the lemma is to first show a legal loop exists (using pigeonhole), and then show an appropriately minimal legal loop must have one of the three allowed forms. The sense in which the loop should be minimal is first that its length should be as small as possible, and second that the number of distinct unoriented edges it crosses (counted without multiplicity) should be minimal.

Lemma 20.14. *If Δ_ϕ has a vertex that only attains one gate in Δ_ϕ, then ϕ can be perturbed to ψ s.t. $\Delta_\psi \subsetneq \Delta_\phi$.*

Proof. Let v be a vertex that only has one gate in Δ_ϕ. Now consider ψ_0 which is defined to be the same as ϕ outside of a small neighbourhood of v but with a perturbation around v so that v now maps to a point ε away from $\phi(v)$ along the edge which corresponds to the image of the one gate at v in Δ. In the figure the purple edges

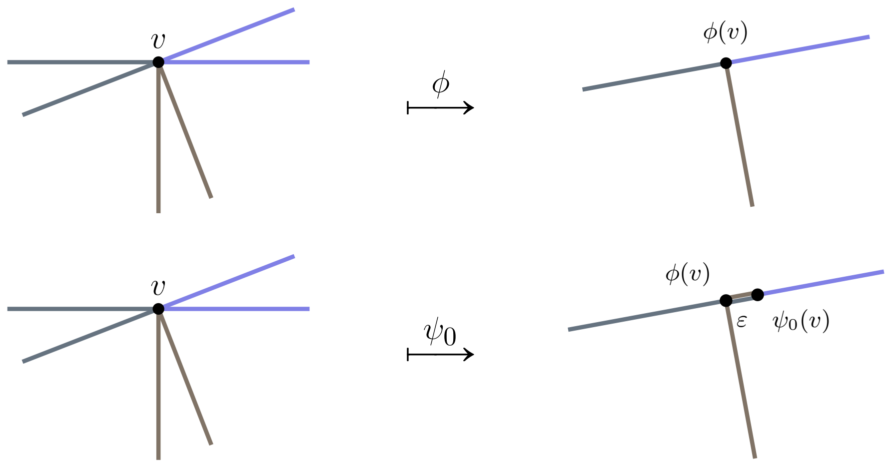

FIGURE 79

are included in Δ, edges of the same color denote the edges in the same gate; and the mapping is specified via correspondence of colors. Let ψ be the modification on ψ_0 s.t. it is linear. Then in ψ it is of slope σ on edges not adjacent to v. Edges not in Δ_ϕ cannot be part of the witness as ε can be made small and the slopes vary continuously. It is clear that $\Delta_\psi \subsetneq \Delta_\phi$ as edges adjacent to v are in Δ_ϕ but are not in Δ_ψ. □

With these lemmas it is straightforward to show the proposition:

Proof of Proposition 20.1. Choose an optimal map s.t. the tension graph is as small as possible. By definition of optimal map and tension graph, the tension graph cannot be empty. Lemma 20.14 shows that all vertices have at least 2 gates. Lemma 20.12 shows that there exists a legal immersed loop α. By Lemma 20.11 α is the desired witness. □

21. Geodesics in CV_n (10/18, HT, YW)

Definition 21.1. Let $\Phi : \Gamma \to \Gamma'$ be a graph map that is linear on edges. Then Δ_Φ is a subgraph of Γ on which Φ has maximal slope.

Using the ideas from last class, we can prove the following straightforward lemma. It is very useful for showing a map is optimal.

Lemma 21.2. *If $\Phi : \Gamma \to \Gamma'$ is linear on edges and Δ_Φ has a legal loop, then Φ is optimal (i.e. the Lipschitz constant is minimized).*

This more or less follows from Lemma 19.10, letting α be the legal loop obtained above.

We'll now discuss geodesics in CV_n. Since these are geodesics in an asymmetric metric, we begin by giving basic definitions.

Definition 21.3. A geodesic in an asymmetric metric space X is a map $\gamma : [a, b] \to X$ satisfying the isometry condition

$$d(\gamma(s), \gamma(t)) = t - s$$

whenever $a \leqslant s \leqslant t \leqslant b$.

Since we require $s \leqslant t$, reversing a geodesic might not yield a geodesic.
We need a more useful characterization.

Lemma 21.4. *Fix $\gamma : [a, b] \to X$, and assume that*

$$d(\gamma(s), \gamma(u)) = d(\gamma(s), \gamma(t)) + d(\gamma(t), \gamma(u))$$

for all $s < t < u$ with $s, t, u \in [a, b]$. Then γ is a reparameterization of a geodesic in the sense that there exists a monotone homeomorphism $R : [a', b'] \to [a, b]$ such that $\gamma(R(t))$ is a geodesic.

Proof. Set $a' = 0$ and $b' = d(\gamma(a), \gamma(b))$, and then define $R^{-1}(s) = d(\gamma(a), \gamma(s))$. Checking that R and R^{-1} are monotone homeomorphisms is fairly straightforward. Now, for $t_1 \leqslant t_2$, the assumption gives that

$$d(\gamma(R(t_1)), \gamma(R(t_2))) = d(\gamma(a), \gamma(R(t_2))) - d(\gamma(a), \gamma(R(t_1))).$$

This is $t_2 - t_1$ by definition of $R^{-1}(s)$. $\square$

Checking this condition on CV_n turns out to not be so bad.

Lemma 21.5. *Suppose $\Phi : \Gamma \to \Gamma'$ is an optimal map, and factors as*

$$\begin{tikzcd} \Gamma \arrow[rr, "\Phi"] \arrow[dr, "\Phi_1"'] & & \Gamma' \\ & \Gamma'' \arrow[ur, "\Phi_2"'] & \end{tikzcd}$$

where Φ_1 and Φ_2 are also optimal. Assume further that some α is a witness for Φ, Φ_1, and Φ_2. Then

$$d(\Gamma, \Gamma') = d(\Gamma, \Gamma'') + d(\Gamma'', \Gamma').$$

Proof. Since α is a witness, by definition α is legal and lies in Δ_Φ, Δ_{Φ_1}, and Δ_{Φ_2}. So each map sends the immersed representative of α in its domain to the immersed representative of α in its codomain. Additionally Φ stretches α by $e^{d(\Gamma,\Gamma')}$, Φ_1 stretches α by $e^{d(\Gamma,\Gamma'')}$, and Φ_2 stretches α by $e^{d(\Gamma'',\Gamma')}$. Since $\Phi = \Phi_2 \circ \Phi_1$, necessarily

$$e^{d(\Gamma,\Gamma')} = e^{d(\Gamma,\Gamma'')} \times e^{d(\Gamma'',\Gamma')}$$

and we can recover the desired equality. $\square$

Note that (with the setup of the previous lemma) if Φ, Φ_1, Φ_2 are linear and (the immersed representative in the appropriate graph of) α is legal and lies in Δ_Φ, Δ_{Φ_1}, and Δ_{Φ_2}, then Φ, Φ_1, and Φ_2 are all automatically optimal.

Exercise 34. *Consider a single simplex of CV_n. Then every straight line segment within that simplex can be reparameterized to be a geodesic. (See Hint A.13.)*

We can now relate geodesics to greedy folding paths.

Corollary 21.6. *If $\Phi : \Gamma \to \Gamma'$ is an optimal map and $\Delta_\Phi = \Gamma$, the greedy folding path from Φ to Φ' is a geodesic.*

Proof. Since $\Delta_\Phi = \Gamma$, Φ is necessarily constant-slope, so greedy folding makes sense. Now, we know a witness α exists for Φ, and since α remains immersed under Φ, it is never folded. It therefore remains a witness throughout the folding process, and combining the previous two lemmas yields the desired result. □

We already know CV_n is connected, and we can now show that it is connected via geodesics.

Corollary 21.7. *For any $\Gamma, \Gamma' \in CV_n$, there exists a geodesic γ that starts at Γ and ends at Γ'.*

Proof sketch. There exists some map from Γ to Γ' (say, deform Γ to get a rose, and then fold the rose to get Γ'). As a consequence, there exists an optimal map $\Phi : \Gamma \to \Gamma'$. If $\Delta_\Phi = \Gamma$, we can apply the previous corollary. Otherwise, adjust (the metric on) Γ by making edges in Δ_Φ longer and making all other edges shorter. This can be done linearly, and the linear adjustment forms a straight line in the simplex containing Γ. At some point, this adjustment will force Δ_Φ to pick up another edge. Either now $\Delta_\Phi = \Gamma$, or we can continue this adjustment procedure, resulting in a concatenation of straight lines. Eventually we get $\Delta_\Phi = \Gamma$. After that point one can use a greedy folding path. Any witness for the original optimal map stays a witness for the whole path, so this concatenation is a geodesic. □

Here's an example of this process (see Figure 80). Starting with the graph in the top left, we rescale to obtain the graph on the top right, then greedily fold to obtain the graph on the bottom.

Remark 21.8. It turns out that finding geodesics in CV_n is not so difficult, and one can make do with a procedure that's substantially simpler than the one we've presented so far. One can subdivide so an optimal map becomes a graph morphism (still after making the tension graph everything), and then fold "in any order". (Lemma 21.5 gives that points obtained by any folding will be on a geodesic.)

A general problem with CV_n (or a feature, depending on one's perspective) is that its geodesics are very non-unique. This isn't so weird: for instance, the taxicab metric on a square grid also has lots of geodesics between almost every pair of points (see Figure 81).

There are a few possible perspectives on why geodesics are not unique.

(1) First, optimal maps are usually not unique (although this is really just a consequence of the next two points).
(2) Second, folding is non-unique, and in fact fairly arbitrary. As mentioned earlier, the greedy folding we've been using is far from the only option.
(3) For a given $\Phi : \Gamma \to \Gamma'$, the metric on CV_n only really cares about what happens on Δ_Γ - whatever occurs on the 'slack' part of the graph is irrelevant.

Exercise 35. *Construct two (or more) distinct geodesics from the left graph to the right graph shown in Figure 82.*

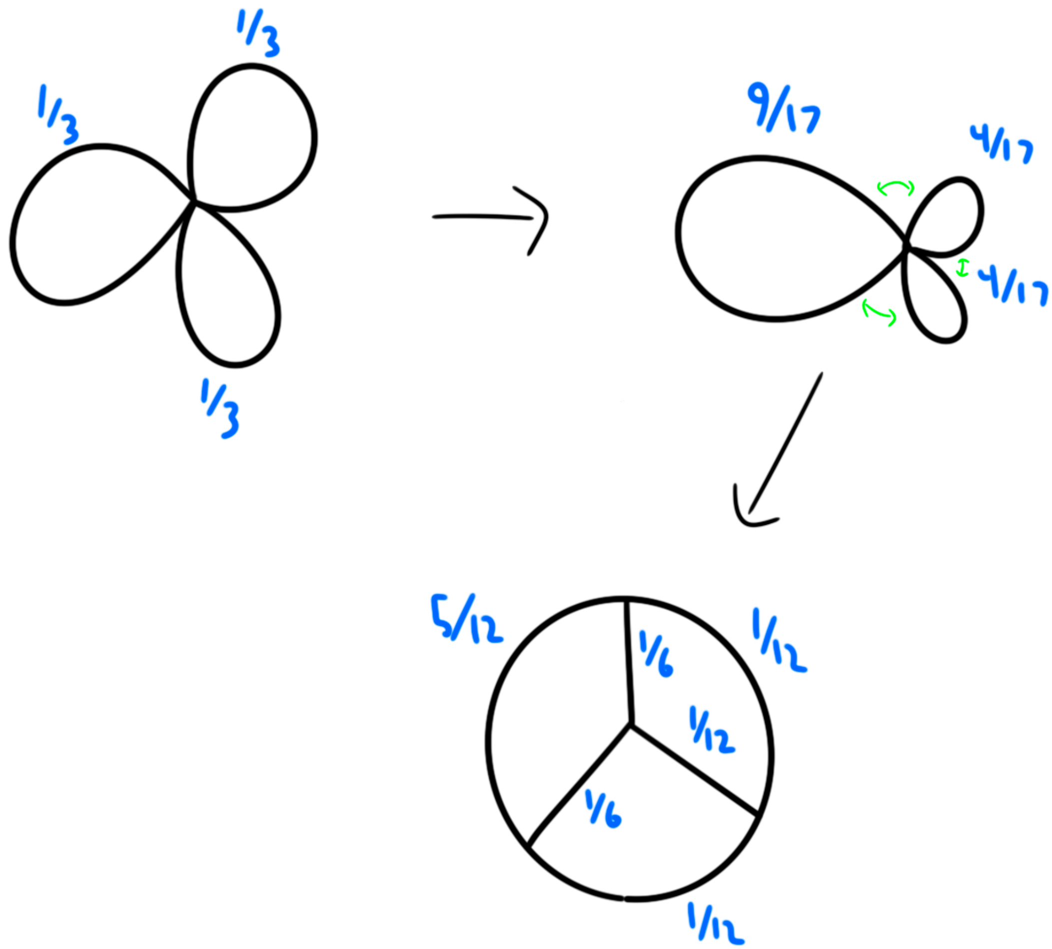

Figure 80

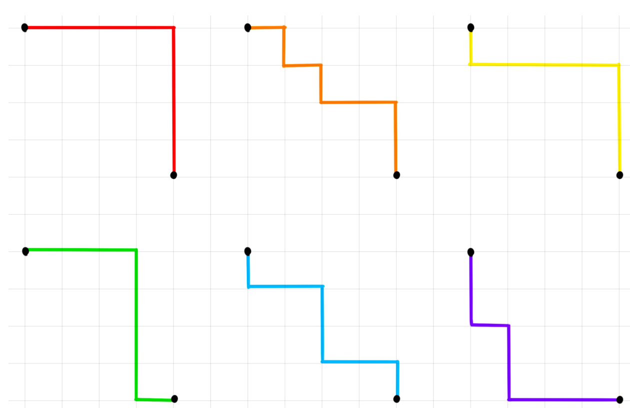

Figure 81

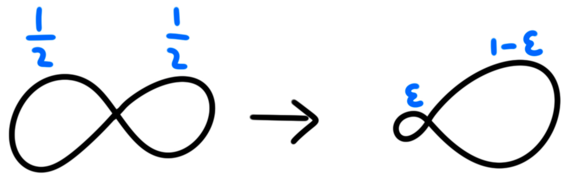

FIGURE 82

22. Translation length, train track maps (10/20, KL, ZH)

Our next topic concerns how we think about classifying isometries of a metric space. We will start with some general definitions and will then apply this to $\mathrm{Out}(F_n)$.

Definition 22.1. Let (X, d) be an asymmetric metric space, and let $\Phi \in \mathrm{Isom}(X)$. The **displacement function** is the function

$$\begin{aligned} D = D_\Phi \colon X &\to [0, \infty) \\ x &\mapsto d(x, \Phi(x)). \end{aligned}$$

The **translation length** is $\tau(\Phi) := \inf(D_\Phi)$.

Exercise 36. *Show that:*

(1) For any $x \in X$, the limit

$$\lim_{n\to\infty} \frac{1}{n} d(x, \Phi^n(x))$$

exists. (Hint: use Fekete's lemma.)

(2) The limit in (1) is independent of the choice of $x \in X$. (Hint: Use the triangle inequality.)

Definition 22.2. The limit in (1) is called the stable translation length and is denoted by $\widehat{\tau}(\Phi)$.

Lemma 22.3. *We have $\widehat{\tau}(\Phi) \leqslant \tau(\Phi)$.*

Proof. For any $\varepsilon > 0$, there exists an $x \in X$ such that $d(x, \Phi(x)) \leqslant \tau + \varepsilon$. For every positive integer n, we have

$$\begin{aligned} d(x, \Phi^n(x)) &\leqslant \sum_{i=1}^{n} d(\Phi^{i-1}(x), \Phi^i(x)) && \text{(by the triangle inequality)} \\ &= n d(x, \Phi(x)) && \text{(since } \Phi \text{ is an isometry)} \\ &\leqslant n(\tau + \varepsilon). \end{aligned}$$

We conclude that $\widehat{\tau} \leqslant \tau + \varepsilon$ for any $\varepsilon > 0$ and hence that $\widehat{\tau} \leqslant \tau$. ☐

We classify isometries as follows:

Definition 22.4. An isometry $\Phi \in \mathrm{Isom}(X)$ is called

- **elliptic** if Φ has a fixed point;

- **hyperbolic** if $\tau > 0$ and it is realized;
- **parabolic** if τ is not realized.

One uses this classification for matrices acting on symmetric spaces, isometries of manifolds, and the mapping class group acting on Teichmüller space. Our goal is to understand it for $\Phi \in \mathrm{Out}(F_n)$ acting on outer space CV_n.

We begin by describing when $\Phi \in \mathrm{Out}(F_n)$ is elliptic:

Lemma 22.5. *An element $\Phi \in \mathrm{Out}(F_n)$ is elliptic if and only if it has finite order.*

Proof. If $\Phi \in \mathrm{Out}(F_n)$ is elliptic, it fixes a point. The stabilizer of a point in CV_n is the *finite* group of isometries of the graph. Hence $\langle\Phi\rangle$ is contained in a finite group, and Φ must have finite order.

For the reverse direction, if Φ has finite order, then by the Nielsen Realization Theorem (Theorem 15.5), Φ fixes a point. Therefore, it is elliptic. ☐

Hyperbolic isometries are more difficult to describe. Our next goal will be to generate some examples.

Definition 22.6. Let $\phi : \Gamma \to \Gamma$ be optimal. If

1. $\Delta_\phi = \Gamma$ (i.e. ϕ is constant slope),
2. every vertex of Γ has at least two gates, and
3. ϕ maps legal paths to legal paths,

then ϕ is called a **train track map**. Its slope is called the **dilatation**. (See Remark 28.10 for a comment on which train track structure is used.)

Note that to check whether ϕ satisfies 3, it suffices to check that:

1. the induced map on gates at every vertex is injective, and
2. ϕ maps edges to legal paths.

This is because an illegal turn in the image of a legal path is either at the image of a vertex on that path, or in the image of an edge on that path.

Remark 22.7. Condition 3 guarantees that the image under ϕ^n of a legal path is legal for any positive integer n. In particular, the image under ϕ^n of an edge is always locally injective; note that this key feature of a train track map can be formulated without referring to train tracks.

Definition 22.8. We say that $\phi : \Gamma \to \Gamma$ represents $\Phi \in \mathrm{Out}(F_n)$ if, by using the identification of $\pi_1(\Gamma)$ with F_n provided by a marking, we have $\phi_* = \Phi$. (In this situation people sometimes also write $\phi : \Gamma \to \Phi\Gamma$.)

Proposition 22.9. *Suppose that $\phi : \Gamma \to \Gamma$ represents $\Phi \in \mathrm{Out}(F_n)$ and that ϕ is optimal. Suppose that ϕ has an invariant subgraph $\Gamma_0 \subset \Delta_\phi$ such that $\phi|_{\Gamma_0}$ is a train track map. Then*

$$d(\Gamma, \Phi(\Gamma)) = \tau(\Phi) = \log(\sigma),$$

where $\sigma = \sigma(\phi)$ is the Lipschitz constant of ϕ. In particular, ϕ is hyperbolic (if $\sigma > 1$) or elliptic (if $\sigma = 1$).

Remark 22.10. The special case where $\Gamma_0 = \Delta_\phi = \Gamma$ is already important: train track maps with slope greater than 1 are hyperbolic.

Proof of Proposition 22.9. By definition, since ϕ is optimal, we have $d(\Gamma, \phi(\Gamma)) = \log\sigma$. Of course, $\tau \leqslant \log\sigma$, and so it suffices to show that $\tau \geqslant \log\sigma$. Since $\tau \geqslant \widehat{\tau}$, it suffices to show that $\widehat{\tau} \geqslant \log\sigma$. By definition,

$$\widehat{\tau} = \lim_{n\to\infty} \frac{1}{n} d(x, \Phi^n(x)),$$

so it suffices to show that

$$d(\Gamma, \Phi^n(\Gamma)) \geqslant n\log\sigma$$

for every $n \in \mathbb{N}$. Note that ϕ^n gives a linear map from Γ to $\Phi^n(\Gamma)$ that is σ^n-Lipschitz. By results from Lecture 19, there is a witness α for ϕ contained in Γ_0. Any power of ϕ maps α to an immersed loop (since any power of ϕ maps legal paths to legal paths). So since $\Gamma_0 \subset \Delta_\phi$, we have α is also a witness for ϕ^n. Therefore, ϕ^n is optimal and $d(\Gamma, \Phi^n(\Gamma)) = n\log\sigma$. ☐

Example 22.11. Let $\Phi \in \mathrm{Out}(F_2)$ be defined by $a \mapsto b$ and $b \mapsto ab$. Let $\Gamma = R_2$, and let $\phi : \Gamma \to \Gamma$ be the map representing Φ that maps the vertex of Γ to itself. One question we might ask is: can we choose lengths $\ell(a)$ and $\ell(b)$ so that ϕ is a train track map of slope λ, for some constant λ?

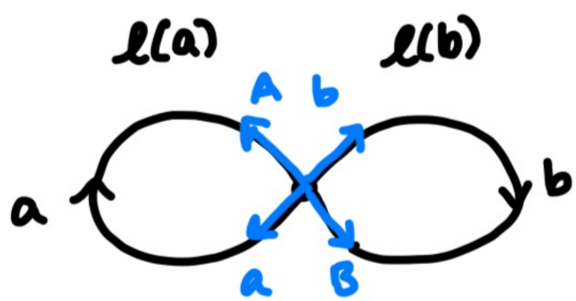

FIGURE 83

For this to be the case, since $a \mapsto b$ we would need $\ell(b) = \lambda\ell(a)$. Likewise, since $b \mapsto ab$, we would need $\ell(a) + \ell(b) = \lambda\ell(b) = \lambda^2\ell(a)$. This latter condition implies $1 + \lambda = \lambda^2$, which yields $\lambda = \frac{1+\sqrt{5}}{2}$. We probably also want to impose the condition $\ell(a) + \ell(b) = 1$, which is satisfied only if we choose $\ell(a) = \frac{1}{\lambda}$ and $\ell(b) = \frac{1}{\lambda^2}$. With these lengths, ϕ is constant slope in $\Delta_\phi = \Gamma$.

For the gates, we will use the descriptors a, A, b, B, pictured in blue in Figure 83. Then

$$a \mapsto b, b \mapsto a, A \mapsto B, \text{ and } B \mapsto B.$$

The gates are then $\{a\}, \{b\}, \{A, B\}$, and we have

$$\phi : \{a\} \mapsto \{b\}, \{b\} \mapsto \{a\}, \text{ and } \{A, B\} \mapsto \{A, B\}.$$

In particular, ϕ is injective on gates. We conclude that ϕ is indeed a train track map of slope λ.

23. Parabolic automorphisms (10/23, YL, SK)

We continue with the last example from last class, and try to build hyperbolic automorphisms in a more general setting:

Example 23.1. Suppose $\phi : \Gamma \to \Gamma$ sends vertices to vertices and each edge to an immersed path. Let $e_1, ..., e_k$ be the unoriented edges of Γ. The **transition matrix** M is a $k \times k$ matrix with entry M_{ij} being the number of times $\phi(e_j)$ crosses e_i. (In Example 22.11, the transition matrix is $M = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$.) Assume M has a power with all positive entries. Then by the Perron-Frobenius Theorem, $\exists \lambda, \ell_1, ..., \ell_k > 0$, such that

$$M \begin{pmatrix} \ell_1 \\ \vdots \\ \ell_k \end{pmatrix} = \lambda \begin{pmatrix} \ell_1 \\ \vdots \\ \ell_k \end{pmatrix}$$

If we give each e_i the length ℓ_i, then (after replacing ϕ with its linearization for this metric) ϕ has a constant slope of λ. If further every vertex has at least 2 gates and ϕ sends legal paths to legal paths, then ϕ is a train track map and in particular is hyperbolic.

The next is an example of a parabolic automorphism.

Example 23.2. Let $\Phi \in \mathrm{Out}(F_2) : a_1 \mapsto a_1, a_2 \mapsto a_1 a_2$. For each $\epsilon > 0$ small, consider the graph morphism corresponding to Φ described by Figure 84, such that

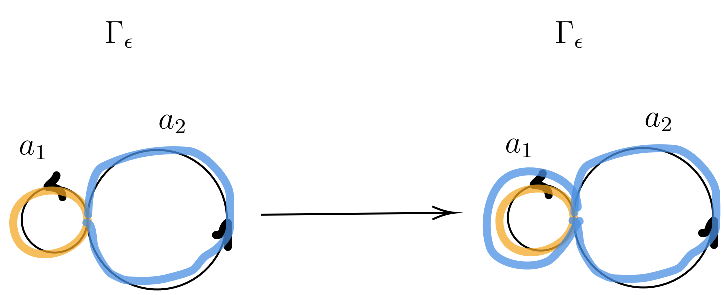

Figure 84

$l(a_1) = \epsilon, l(a_2) = 1 - \epsilon$. This is a map representing Φ with Lipschitz constant $\sigma = \frac{1}{1-\epsilon}$, and when $\epsilon \to 0$ we have $\sigma \to 1$. Thus, $d(\Gamma_\epsilon, \Phi(\Gamma_\epsilon)) \to 0$ as $\epsilon \to 0$, which implies $\tau_\Phi = 0$. Hence, Φ cannot be hyperbolic. Further, since Φ has infinite order, it cannot be elliptic. Thus it is parabolic.

Definition 23.3. $\Phi \in \mathrm{Out}(F_n)$ is **reducible** if it can be represented by a $\phi : \Gamma \to \Gamma$, such that for some subgraph $\Gamma_0 \subsetneq \Gamma$ that is not a forest, we have $\phi(\Gamma_0) \subset \Gamma_0$. Otherwise, Φ is called **irreducible**.

Remark 23.4. It is implicit above that Γ is a finite graph without leaves. The definition requires that Γ_0 have some of the topology of the graph but not all of it. The invariant subgraph gives Φ a structure analogous to a block upper triangular matrix. (If Γ_0 isn't connected we should actually compare to matrices with blocks that are permuted. These ideas will be made more precise later.)

Definition 23.5. $\Phi \in \mathrm{Out}(F_n)$ is called **fully irreducible** if for every $k \geqslant 1$, Φ^k is irreducible.

Theorem 23.6. *Every parabolic element of* $\mathrm{Out}(F_n)$ *is reducible.*

Definition 23.7. The **ϵ-thick part** CV_n^ϵ of CV_n is the subset where all immersed loops have length at least ϵ.

The idea of the above definition is that there are no short loops, but we can have short edges.

Exercise 37. *Show that* $CV_n^\epsilon/\mathrm{Out}(F_n)$ *is compact. (See Hint A.14.)*

We now prove a lemma helpful in proving Theorem 23.6. Assume Φ is parabolic and pick $\Gamma_k \in CV_n$ such that $d(\Gamma_k, \Phi(\Gamma_k)) \to \tau_\Phi$.

Lemma 23.8. $\forall \epsilon > 0$*, there are only finitely many* Γ_k *in* CV_n^ϵ*.*

Proof. Suppose there exists a fixed $\epsilon > 0$ such that there are infinitely many $\Gamma_k \in CV_n^\epsilon$. Passing to a subsequence, WLOG assume all $\Gamma_k \in CV_n^\epsilon$. By compactness of $CV_n^\epsilon/\mathrm{Out}(F_n)$ we may pass to another subsequence and assume WLOG Γ_k converges in $CV_n^\epsilon/\mathrm{Out}(F_n)$.

Thus, for each k we can find a $\Psi_k \in \mathrm{Out}(F_n)$, and we can find a $\Gamma_\infty \in CV_n$, such that $\Psi_k\Gamma_k \to \Gamma_\infty$. By the triangle inequality, as $k \to \infty$ we have the limit

$$\begin{aligned} d(\Psi_k^{-1}\Gamma_\infty, \Phi\Psi_k^{-1}\Gamma_\infty) &\leqslant d(\Psi_k^{-1}\Gamma_\infty, \Gamma_k) + d(\Gamma_k, \Phi\Gamma_k) + d(\Phi\Gamma_k, \Phi\Psi_k^{-1}\Gamma_\infty) \\ &= d(\Gamma_\infty, \Psi_k\Gamma_k) + d(\Gamma_k, \Phi\Gamma_k) + d(\Psi_k\Gamma_k, \Gamma_\infty) \\ &\to 0 + \tau + 0 = \tau. \end{aligned}$$

Thus, $d(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty) \to \tau$.

Exercise 38. *For any* $B > 0$ *and* $\Gamma \in CV_n$*, we have* $|\{\Phi : d(\Gamma, \Phi\Gamma) \leqslant B\}| < \infty$*. (See Hint A.15.)*

The exercise above implies that after passing to a subsequence, $\Psi_k\Phi\Psi_k^{-1}\Gamma_\infty$ is constant. Thus,

$$d(\Gamma_\infty, \Psi_k\Phi\Psi_k^{-1}\Gamma_\infty) = \tau = d(\Psi_k^{-1}\Gamma_\infty, \Phi\Psi_k^{-1}\Gamma_\infty),$$

showing that actually τ is realized and giving a contradiction. □

24. Parabolics, hyperbolics (10/25, SZ, YM)

Theorem 24.1. *Every parabolic* $\Phi \in \mathrm{Out}(F_n)$ *is reducible.*

Proof. Pick $\{\Gamma_k\}_k \subseteq CV_n$ such that $d(\Gamma_k, \Phi\Gamma_k) \to \tau$. Define $\Gamma_k(\delta)$ to be the union of essential loops of length less or equal to δ in Γ_k. (A loop is called essential if it is not homotopically trivial.) The proof is mainly in the following lemma:

Lemma 24.2. *There exists k and δ such that $\Gamma_k(\delta)$ is not empty and not Γ_k and* $\mathrm{Core}(\phi(\Gamma_k(\delta))) \subseteq \Gamma_k(\delta)$.

Proof. Pick ε very small such that

$$(6n-6)\varepsilon(e^{\tau+1})^{3n+100} \ll 1.$$

(Note $6n-6$ is twice the maximal number of edges. Essentially loops may enter an edge a bit and then backtrack, so $\Gamma_k(\delta)$ can contain part of an edge on one side of the edge and part on the other side; this is why we use twice the number of edges.) By Lemma 23.8, we can fix $\Gamma_k \notin CV_n^\varepsilon$. Suppose Φ is represented by $\phi : \Gamma_k \to \Gamma_k, \sigma(\phi) < e^{\tau+1}$.

Since every (partial) edge of $\Gamma_k(\varepsilon)$ is contained in a loop of length at most ε, each (partial) edge of $\Gamma_k(\varepsilon)$ has length at most ε. Hence

$$\mathrm{vol}(\Gamma_k(\varepsilon)) \leqslant \varepsilon(6n-6) \ll (e^{\tau+1})^{-(3n+100)}.$$

Since $\Gamma_k(\delta)$ is the union of essential loops of length $\leqslant \delta$ and ϕ is $e^{\tau+1}$-Lipschitz, $\phi(\Gamma_k(\delta))$ is the union of essential loops of length $\leqslant \sigma(\phi)\delta$, and we get

$$\phi(\Gamma_k(\delta)) \subseteq \Gamma_k(e^{\tau+1}\delta).$$

In the sequence

$$\mathrm{Core}(\Gamma_k((e^{\tau+1})^\ell\varepsilon)), \quad \ell = 0, 1, \ldots$$

there are at most $3n-3$ changes, corresponding to the maximal number of edges. Therefore there exists $i < 3n+10$ such that

$$\mathrm{Core}(\Gamma_k((e^{\tau+1})^i\varepsilon)) = \mathrm{Core}(\Gamma_k((e^{\tau+1})^{i+1}\varepsilon)).$$

So, for this i, we have

$$\mathrm{Core}(\phi(\Gamma_k((e^{\tau+1})^i\varepsilon))) \subseteq \mathrm{Core}(\Gamma_k((e^{\tau+1})^{i+1}\varepsilon)) = \mathrm{Core}(\Gamma_k((e^{\tau+1})^i\varepsilon)) \subseteq \Gamma_k((e^{\tau+1})^i\varepsilon),$$

proving the result for $\delta = (e^{\tau+1})^i\varepsilon$. □

The choices of $3n+100$ or $3n+10$ are somewhat arbitrary and incorporate a healthy safety margin.

Back to the proof of the theorem. Using a homotopy equivalence of Γ that restricts to a deformation retract of $\Gamma(\delta)$ onto $\mathrm{Core}(\Gamma(\delta))$, we can homotope ϕ to ϕ' with $\phi'(\Gamma(\delta)) \subset \Gamma(\delta)$. □

Recall that the dilatation λ of a hyperbolic Φ is defined to be $e^{\tau(\phi)}$.

Theorem 24.3. *Let $\Phi \in \mathrm{Out}(F_n)$ be hyperbolic with dilatation λ. Suppose $\Gamma \in CV_n$ is such that*

$$d(\Gamma, \Phi(\Gamma)) = \log \lambda.$$

Say Φ is represented by $\phi : \Gamma \to \Gamma$. Possibly after perturbing both Γ and the optimal map, we have

(1) $\phi(\Delta) \subseteq \Delta$,
(2) ϕ sends legal paths in Δ to legal paths,
(3) all vertices in Δ have at least 2 gates.

(i.e. $\phi|_\Delta$ is a train track map.)

In this statement and in its proof, the perturbation will always perturb the property $d(\Gamma, \Phi(\Gamma)) = \log\lambda$. For the proof we will follow [Bes11].

Remark 24.4. The "converse" is true by Proposition 22.9.

Recall that Lemma 20.14 states that if Δ_ϕ has a vertex that only has one gate, then ϕ can be perturbed to ψ such that $\Delta_\psi \subsetneq \Delta_\phi$. (This perturbation effects only the optimal map; the metric graph is left unchanged.)

Now let Φ be hyperbolic, and let $\phi : \Gamma \to \Gamma$ be optimal representing Φ with $\sigma(\phi) = \lambda$. We will keep these properties when we perturb in the following lemma.

Lemma 24.5. *If all vertices of Δ_ϕ have at least 2 gates and $\phi(\Delta) \nsubseteq \Delta$, we can perturb to $\phi' : \Gamma' \to \Gamma'$ such that Γ and Γ' are homeomorphic and $\Delta_{\phi'} \subsetneq \Delta_\phi$.*

Proof. Suppose e is an edge of Δ_ϕ and $\phi(e)$ is not contained in Δ_ϕ. Perturb the metric on Γ by scaling Δ by $1+\varepsilon$ and the complement of Δ by $\frac{1}{1+\varepsilon'}$, with ε' chosen depending on ε to keep the volume 1. Let Γ' be Γ with this new metric and ϕ' be the linearization of ϕ for this metric.

Since $\varepsilon, \varepsilon' \approx 0, \Delta_{\phi'} \subseteq \Delta_\phi$ (since the slope of an edge varies continuously in the deformation). Also,

$$\text{slope}(e) = \frac{\text{length } \phi'(e)}{\text{length } e}.$$

Comparing to before the deformation, the denominator is scaled by $1 + \varepsilon$ and the numerator is scaled by a smaller amount. Therefore $e \notin \Delta_{\phi'}$. □

25. Nice maps for hyperbolic automorphisms (10/27, YW, JG)

Theorem 25.1. *Suppose Φ is a hyperbolic with dilatation $\lambda = e^{\tau_\Phi}$, and $d(\Gamma, \Phi\Gamma) = \log\lambda$. (All perturbation in proof will keep this property.)*

After perturbing, there exists $\phi : \Gamma \to \Gamma$ that is optimal such that:

(1) $\phi(\Delta) \subseteq \Delta$,
(2) ϕ(legal path in Δ) = legal path, and
(3) Δ has no 1-gate vertex.

We know from last time that if the first or the third condition does not hold, we can perturb ϕ or the metric to make Δ smaller (no change to topology of Γ).

Lemma 25.2. *Suppose $\phi : \Gamma \to \Gamma$ is optimal with $\sigma(\phi) = \lambda$.*

If there exists e an edge of Δ with $\phi(e)$ not legal, we can deform to $\phi' : \Gamma' \to \Gamma'$ such that either

(1) Γ' has more edges than Γ or
(2) Γ, Γ' are homeomorphic and $\Delta_{\phi'} \subsetneq \Delta_\phi$.

Proof. Say $\phi(e)$ makes an illegal turn at $\phi(p), p \in e$, as in Figure 85.

Let Γ' be the result of folding a tiny bit at $\phi(p)$ (just the two directions of the illegal turn) as shown in Figure 86.

Note ϕ induces a map $\phi'_0 : \Gamma' \to \Gamma'$ by the definition of illegal turn. Let ϕ' be the linearization of ϕ'_0.

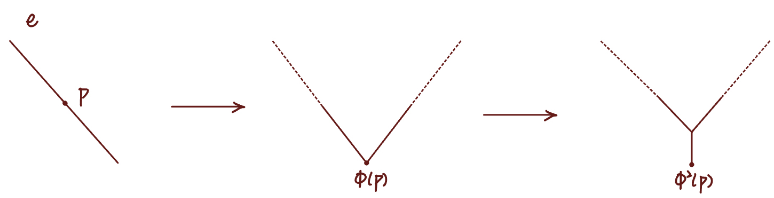

FIGURE 85

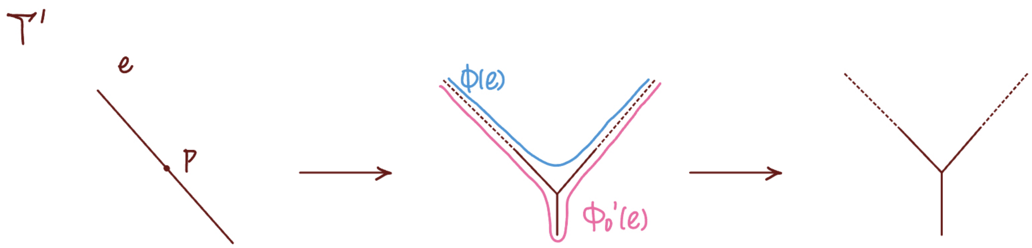

FIGURE 86

We need to rescale Γ' to be vol 1, but since Γ' is both the domain and codomain, this does not effect slopes.

If the folding changes the topology, we have 1 as shown in Figure 87.

FIGURE 87

If not, we may have the example as shown in Figure 88.

Then we note $e \notin \Delta_{\phi'}$. $\square$

We have now dealt with every possibility except for a legal turn mapping to an illegal turn.

Lemma 25.3. *Suppose $\phi : \Gamma \to \Gamma$ as before, ϕ maps edges to legal paths, and ϕ_* maps a legal turn to an illegal turn. Suppose Δ has no 1-gate vertices. Then we can perturb to get $\phi' : \Gamma' \to \Gamma'$ as before with*

(1) Γ' has more edges than Γ, or

FIGURE 88

(2) Γ', Γ *are homeomorphic, and* $\Delta = \Delta'$*, but the total number of gates is greater in* Γ *than* Γ'.

Proof. If we have Figure 89, fold a bit at $\phi(v)$ to get Γ', ϕ' as before as shown in

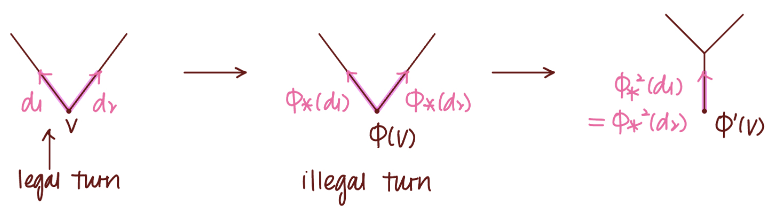

FIGURE 89

Figure 90.

FIGURE 90

If topology changes, we have 1; otherwise, the topology does not change and we would have Figure 91. But $d_1 \sim d_2$, so the number of gates has gone down. $\square$

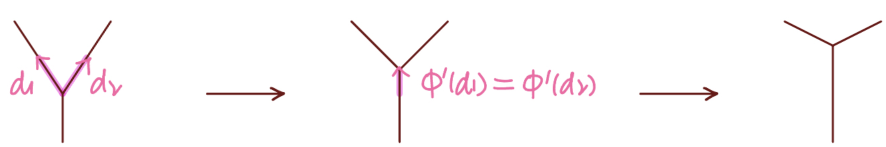

FIGURE 91

We can now conclude as follows

Proof of Theorem 25.1. : Pick a perturbation (which we rename Γ for convenience) and and optimal map $\phi : \Gamma \to \Gamma$ according to the following priorities:

(1) Γ has as many edges as possible,
(2) Δ has as few edges as possible, and
(3) Γ has as few gates as possible.

The lemmas then give that ϕ has all the desired properties. □

We now turn to the question of "growth", aiming to ask an analogous question to "how does $\|A^k v\|$ grow as $k \to \infty$ when A is a matrix and v is a vector and $\| \cdot \|$ is a norm." Instead of using a norm, we'll make use of a metric graph.

Recall that $\ell_\alpha(\Gamma)$ = length of the immersed representative of α in Γ and

$$e^{-d(\Gamma',\Gamma)} \leqslant \frac{\ell_\alpha(\Gamma')}{\ell_\alpha(\Gamma)} \leqslant e^{d(\Gamma,\Gamma')},$$

since $\ell_\alpha(\Gamma') \leqslant \ell_\alpha(\Gamma) e^{d(\Gamma,\Gamma')}$ and $\ell_\alpha(\Gamma) \leqslant \ell_\alpha(\Gamma') e^{d(\Gamma',\Gamma)}$. Note that the upper and lower bounds do not depend on α.

Lemma 25.4. *If Φ is hyperbolic and Γ is arbitrary, then there exists c such that $\ell_{\Phi^k(\alpha)}(\Gamma) \leqslant c\lambda^k$, where $\lambda = e^\tau$ is the dilatation.*

Proof. Find $\psi : \Gamma' \to \Gamma'$ that is optimal and has slope λ to represent Φ. The lemma is true for Γ', so it is true for Γ. □

We will now use this lemma to give an example of a parabolic with positive translation length.

Example 25.5. Consider $\Phi \in Out(F_4)$ such that

$$\begin{aligned} a_1 &\longmapsto a_2 \\ a_2 &\longmapsto a_1 a_2 \\ a_3 &\longmapsto a_4 \\ a_4 &\longmapsto a_3 a_1 a_4 \end{aligned}$$

Note that, restricted to $\langle a_1, a_2 \rangle$, this is our first example of a hyperbolic automorphism, studied in Example 22.11 with dilatation $\lambda = \frac{1+\sqrt{5}}{2}$. Restricting to $\langle a_3, a_4 \rangle$ almost gives a second copy of this automorphism, except there is an a_1 in the formula for the image of a_4.

Figure 92 gives a natural map which has max slope λ plus something that goes to 0 as $\epsilon \to 0$; so $\tau \leqslant \lambda$.

As in the proof of Proposition 22.9, one can use stable translation length to show that $\tau \geqslant \lambda$. Hence $\tau = \lambda$.

But $\ell_{\Phi^k}(a_3) \asymp k\lambda^k$ since the same is true for the linearization as shown in Figure 93. (Note that there are no inverses in the formulas for Φ, allowing an unusually good comparison to the abelianization.)

Keeping Lemma 25.4 in mind, we conclude that Φ is parabolic and has positive translation length.

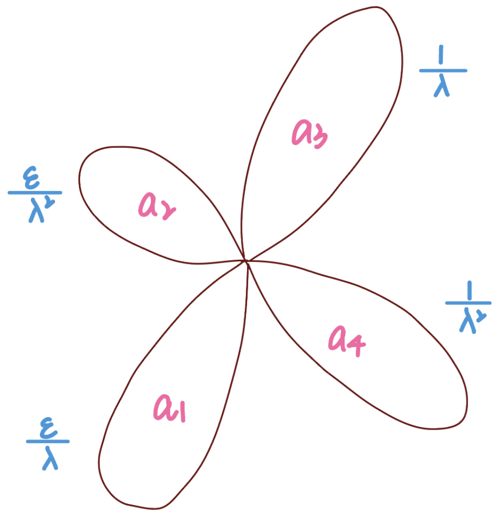

FIGURE 92

$$\left(\begin{array}{cc|cc} 0 & 1 & & 1 \\ 1 & 1 & & \\ \hline & & 0 & 1 \\ & & 1 & 1 \end{array}\right)^k \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

FIGURE 93

26. GROWTH, ALGEBRAIC DEFINITION OF REDUCIBLE (10/30, YL, YM)

Definition 26.1. If $\Phi \in \mathrm{Out}(F_n)$, $\gamma \in F_n, \gamma \neq Id$, define the **growth** of γ under Φ to be

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

where $\Gamma \in CV_n$ is arbitrary.

Exercise 39. *Show that this definition does not depend on the choice of* Γ. *(See Hint A.16.)*

You can compare this definition to: if $\Phi \in GL(n, \mathbb{R}), v \in \mathbb{R}^n$, then $\limsup_{m\to\infty} \frac{\log \|\Phi^m v\|}{m}$ does not depend on the choice of $\| \cdot \|$.

Although it is beyond the scope of this course, to give an idea for what is possible we remark that train track technology can also be used to prove the following [Lev09]:

Theorem 26.2. *$\forall \Phi \in \mathrm{Out}(F_n)$, $\exists \lambda_1, \cdots, \lambda_k > 0$ (all weak Perron) such that $\forall \gamma$ the growth rate $\tau(\Phi, \gamma)$ is either 0 or one of the $\log \lambda_i$. Moreover, $\tau_\Phi = \max(0, \log \lambda_i)$.*

Example 26.3. Let $\Phi \in \mathrm{Out}(F_n)$ be defined by $a_1 \mapsto a_1, a_i \mapsto a_{i-1}a_i, i \geqslant 2$. Then

$$Ab(\Phi) = \begin{pmatrix} 1 & 1 & 0... & 0 \\ 0 & 1 & 1... & 0 \\ ... & & & \\ 0 & 0 & ...1 & 1 \\ 0 & 0 & ...0 & 1 \end{pmatrix}$$

is a Jordan block and $\|Ab(\Phi)^m \begin{pmatrix} 0 \\ \vdots \\ 1 \end{pmatrix}\| \sim m^{n-1}$. Since the automorphism is defined using positive words, this implies that $\ell_{\Phi^m(a_n)}(R_n) \sim m^{n-1}$, where R_n is the standard rose. Thus, if $n > 2$ Φ is "polynomially growing but not linear". This is not surprising comparing to $GL(n, \mathbb{Z})$ but this sort of behavior does not happen in mapping class groups.

Recall that parabolics are reducible. Thus, being irreducible and having infinite order implies being hyperbolic. Our main results on hyperbolics imply that every irreducible is represented by a train track map.

Below we introduce the algebraic reformulation of being irreducible.

Proposition 26.4. *$\Phi \in \mathrm{Out}(F_n)$ is reducible if and only if there exists a free product decomposition*

$$F_n = A_1 * A_2 * \cdots * A_k * C$$

with $k \geqslant 1$ and each A_i is neither trivial nor all of F_n, such that Φ permutes the conjugacy classes $[A_1], \ldots, [A_k]$.

Remark 26.5. Here are some clarifications on the statement

1. The assumption implies that either $k \geqslant 2$ or $C \neq \{id\}$.
2. But if $k \geqslant 2$, C could be $\{id\}$.
3. Nothing is required of $\Phi(C)$.
4. Say $\phi \in \mathrm{Aut}(F_n)$, $\Phi = [\phi]$, and say $\phi(A_i) = g_i A_{\sigma(i)} g_i^{-1}$ for some g_i and for some $\sigma \in S_k$. We do not necessarily have that ϕ preserves conjugacy classes of $A_1 * \cdots * A_k$, since the g_i's may be different.

Lemma 26.6. *Let A be a free factor of F_n (i.e. $\exists B$ such that $F_n = A * B$). Let C be any finitely generated subgroup of F_n containing A. Then A is a free factor of C.*

Proof. Note that there exists a graph Γ and a subgraph Γ_0 such that $\pi_1(\Gamma) = F_n$ and $\pi_1(\Gamma_0) = A$. Pick a basis $x_1, ..., x_n$ of F_n such that $x_1, ..., x_l$ is a basis of A. Consider the core Γ_C of the C-cover of Γ (i.e. take the C cover of Γ, trim hanging trees). Since $A < C$, the inclusion $\Gamma_0 \hookrightarrow \Gamma$ lifts to the inclusion $\Gamma_0 \hookrightarrow \Gamma_C$. Viewing Γ_0 as a subgraph of Γ_C we note that $\pi_1(\Gamma_0) = A$ is a free factor of $\pi_1(\Gamma_C) = C$. $\square$

Almost the same proof also gives stronger conclusions, see for example [Asc, Proposition 2.14.], and a statement called the Kurosh Subgroup Theorem gives even stronger conclusions. For the moment, we need only the following:

Corollary 26.7. *If A, C are both rank k free factors of F_n such that $A < C$ then $A = C$.*

Proof. The lemma above implies that there is a subgroup D, such that $A * D = C$. Thus, $rank(A) + rank(D) = rank(C)$ implies $rank(D) = 0$ and hence $D = \{id\}$ and $A = C$. ☐

27. Reducibility, pseudo-Anosovs (11/01, SY, UP)

Proposition 27.1. *Φ is reducible if and only if $\exists\ F_n = A_1 \star \cdots \star A_k \star C,\ (k \geqslant 1,\ A_i \neq \{id\}, F_n)$ and Φ permutes the $[A_i]$'s.*

A reference for this is [BH92, Remark 1.3, Lemma 1.16], where they use a slightly different but equivalent definition of reducible.

Proof. **First**, suppose Φ is reducible. Pick $\phi: \ \Gamma \to \Gamma$ with an invariant subgraph $Z \subsetneq \Gamma$ that is not a forest and $\phi_* = \Phi$. Let $Z_1, \ldots, Z_k$ be the connected components of Z. Since $\ker \phi_\star = \{\mathrm{id}\}$, ϕ cannot map a non-tree component of Z to a tree component. So, WLOG, none of the Z_i are trees. Fix a maximal tree T of Γ and an orientation of the edges not in T and a basepoint, and let

$$A_i = \langle edges\ in\ Z_i - T\rangle \quad ; \quad C = \langle edges\ in\ \Gamma - (T \cup Z)\rangle$$

Note: A_i is isomorphic to $\pi_1(Z_i)$ and $F_n = A_1 \star \cdots \star A_k \star C$. Let $\sigma: \ \{1, \cdots, k\} \to \{1, \cdots, k\}$ be such that $\phi(Z_i) \subset Z_{\sigma(i)}$.

Claim 0: Fix a path from the image of the basepoint to the basepoint, and use this path to turn ϕ_* into an honest automorphism of $\pi_1(\Gamma)$. With this implicit, for each i there is a g_i such that $\phi_*(A_i) \subset g_i A_{\sigma(i)} g_i^{-1}$.

Claim 1: σ is a permutation.

Otherwise, WLOG, $\phi(Z) \subset \bigcup_{i=2}^{n} Z_i$. This intuitively gives a contradiction since the rank of $A_1 \star \cdots \star A_k$ has smaller rank than the rank of $A_2 \star \cdots \star A_k$. Formally, one can abelianize, and use that $\mathbb{Z}^r$ cannot be a subgroup of $\mathbb{Z}^s$ if $r > s$.

Claim 2: $\Phi([A_i]) = [A_{\sigma(i)}]$.

$\Phi([A_i])$ has a representative in $A_{\sigma(i)}$. If they are not equal, then $\mathrm{Rank}(\Phi(A_i)) < \mathrm{Rank}(A_{\sigma(i)})$. Since a power of σ is the id, we get a contradiction.

Conversely, suppose Φ permutes A_i's in a decomposition as in the statement. Pick $f \in \mathrm{Aut}(F_n)$, $[f] = \Phi$, and g_i s.t. $f(A_i) = g_i A_{\sigma(i)} {g_i}^{-1}$.

Consider a graph with an inverse marking as in Figure 94. Define $\phi: \Gamma \to \Gamma$ s.t. the rose of A_i maps to the rose of $A_{\sigma(i)}$, inducing $x \mapsto g_i f(x) {g_i}^{-1}$ on π_1. Let γ_i be the path (loop) corresponding to g_i, and send e_i to $\gamma_i e_{\sigma(i)}$. On the rose for C, map according to f.

Exercise 40. *Check that $\phi_\star = f$.*

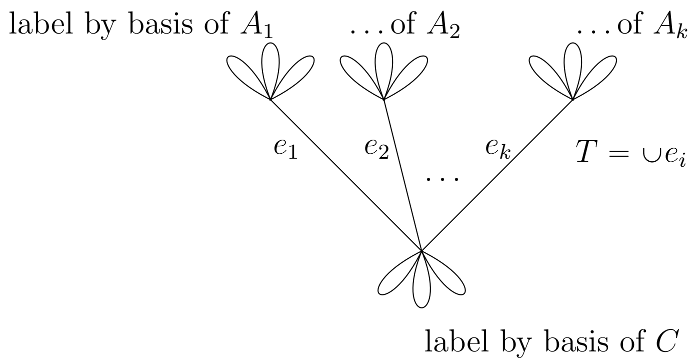

FIGURE 94

This shows Φ is reducible, since Γ has an invariant subgraph Z (formed by the roses for the A_i) that is not a forest. □

Corollary 27.2. *If $f: \Sigma \to \Sigma$ is a homeomorphism of a (punctured) surface that preserves a strict subset of the punctures, then $f_\star \in \mathrm{Out}(\pi_1(\Sigma))$ is reducible.*

Proof. Pick a basis of $\pi_1(\Sigma)$ including loops around punctures in S. The conjugacy classes of these loops get permuted by $f_\star$. □

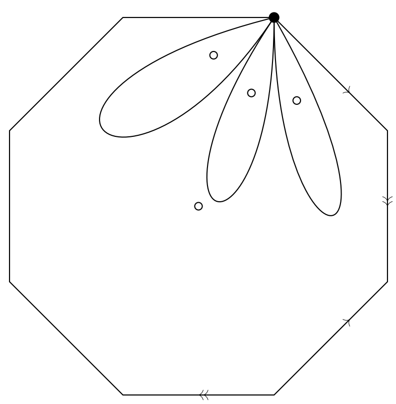

FIGURE 95

Remark 27.3. You cannot find a basis containing a loop around every puncture, since the loops around the punctures would sum to 0 in $H_1 = \mathrm{Ab}(\pi_1)$.

Definition 27.4. $[f] \in \mathrm{MCG}(\Sigma)$ is **pseudo-Anosov (pA)** if it does not preserve a finite union of essential, non-peripheral, simple closed curves (up to isotopy). (Essential means the curve doesn't bound a disk. Non-peripheral means it doesn't bound a disc with one puncture. Simple means no self-intersections.)

Remark 27.5. It is known that:

- "most" mapping classes are pA, and
- $[f]$ is pA if and only if it is hyperbolic for the action on $\mathrm{Teich}(\Sigma)$. (Elliptics are the finite order elements, rest are parabolic.)

We will now present some important examples coming from [BH92, Example 1.4].

Fact 27.6. *If $[f] \in \mathrm{MCG}(\Sigma)$ is a pA, α is a non-peripheral, essential, simple closed curve, and C is a proper free factor of $\pi_1(\Sigma)$, then $f^n(\alpha)$ is not conjugate into C for all n large enough.*

Glimpse of proof. One can take f to "vertically stretch and horizontally contract" (an analogue of "irreducibles have train track maps"). Every vertical line is dense in Σ. Hence if $\Sigma_C \to \Sigma$ is the cover corresponding to $C \subset \pi_1$, vertical lines in Σ_C are proper. Since α is not peripheral, it has a "flat geodesic". $f^n(\alpha)$ is more and more vertical, so it "looks like a vertical line" and cannot close up in Σ_C. ☐

Corollary 27.7. *If $[f] \in \mathrm{MCG}(\Sigma)$ is pA, all periodic conjugacy classes of free factors are of the form: $\langle$loop around a puncture$\rangle$.*

In other words, if A is a free factor and $[f^p(A)] = [A]$, then $A = \langle$*loop around a puncture*$\rangle$.

Proof sketch. Any other free factor A has an α that is not peripheral. Fact 27.6 implies $f^p(\alpha)$ is not conjugate into the free factor A for p large enough. ☐

28. Examples of interesting automorphisms (11/03, YM, AB)

Continue on the discussion of the previous lecture, this gives a further corollary:

Corollary 28.1. *If f is pseudo-Anosov, then f_* is reducible if and only if f preserves a proper nonempty subset of the punctures.*

Example 28.2. Consider

$$\Sigma = \mathbb{R}^2/\mathbb{Z}^2 - \left\{ \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}, \begin{pmatrix} \frac{1}{2} \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ \frac{1}{2} \end{pmatrix} \right\},$$

where the action of f is specified via $\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$. Note that, working modulo $\mathbb{Z}^2$,

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ 0 \end{pmatrix}, \qquad \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} \frac{1}{2} \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ \frac{1}{2} \end{pmatrix}, \qquad \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} 0 \\ \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix}.$$

So f is well defined and f cyclically permutes the three punctures. Standard results imply f is pseudo-Anosov.

Notice that f and f^2 are irreducible, but f^3 is reducible. Therefore, f is irreducible, but not fully irreducible.

Definition 28.3. Φ is **atoroidal** if for all $p \geqslant 1$, Φ^p does not preserve any conjugate class in F_n except the identity.

Remark 28.4. A non-trivial result gives that Φ is atoroidal if and only if $F_n \rtimes_\Phi \mathbb{Z}$ is Gromov hyperbolic. For this reason, some of the literature uses "hyperbolic" as a synonym for "atoroidal". But this is not the same as saying that Φ acts hyperbolically on CV_n, as the following example shows.

Example 28.5. Consider Σ a surface with one puncture, and let $f : \Sigma \to \Sigma$ be pseudo-Anosov. f_* is irreducible, and hence acts hyperbolically on CV_n. But f_* preserves the conjugacy class of the loop around the puncture, so f is not atoroidal. (To rephrase:

this gives an example of a fully irreducible map for which there is a conjugacy class which is fixed and hence in particular does not grow under iteration.)

Example 28.6. Consider $f : \Sigma \to \Sigma$ which is pA, where Σ has two punctures, and each is preserved by f. Corollary 27.7 gives that f_* has exactly two periodic conjugacy classes of free factors. We can write

$$\begin{aligned} F_n &= \langle \text{loop around puncture } 1\rangle * \langle \text{some complement}\rangle \\ &= \langle \text{loop around puncture } 2\rangle * \langle \text{some different complement}\rangle \end{aligned}$$

Remark 28.7. In the example above, the two decompositions given are the only decompositions as in the criterion for reducibility, up to picking different complements. This is as if there was a matrix with two different Jordan Normal Forms, so you should think of it as quite weird. Each of the two decompositions are as good as the other; so one can't pick a decomposition at all canonically. (Additionally the complements cannot be chosen canonically or preserved, but that isn't surprising from the point of view of linear algebra.)

Exercise 41. *If d is a symmetric metric, and $\Phi \in \mathrm{Isom}(X, d)$. Then $\tau_\Phi = \tau_{\Phi^{-1}}$. Additionally, Φ is elliptic/hyperbolic/parabolic if and only if Φ^{-1} is.*

Exercise 42. *Suppose $\Phi \in \mathrm{Out}(F_2)$, and $\mathrm{Ab}(\Phi)$ has eigenvalues λ_1 and λ_2 s.t. $|\lambda_1| \geqslant |\lambda_2|$. Then Φ is parabolic if and only if $\lambda_1 = \pm 1, \lambda_2 = \pm 1$ and Φ is not finite order, and Φ is hyperbolic if and only if $|\lambda_1| > 1$. (As always, Φ is elliptic if and only if it is finite order. See Hint A.17.)*

Exercise 43. *In the situation above, show $\tau_\Phi = \log |\lambda_1|$. Deduce in particular that $\tau_\Phi = \tau_{\Phi^{-1}}$. (See Hint A.18.)*

Remark 28.8. Generally, in $\mathrm{Out}(F_n)$ for $n \geqslant 3$, $\tau_\Phi \neq \tau_{\Phi^{-1}}$. The following gives a concrete example where $\tau_\Phi \neq \tau_{\Phi^{-1}}$. Since translation length controls growth under iteration, this shows that the asymmetry in our metric is a natural and intrinsic feature of $Out(F_n)$, that would show up even if we didn't study CV_n.

Example 28.9. Consider the map

$$\Phi \in \mathrm{Out}(F_3), \quad \begin{cases} a \mapsto b \\ b \mapsto c \\ c \mapsto ab \end{cases}$$

If $\phi : R_3 \to R_3$ is the associated graph map, it is easy to check as follows that ϕ maps legal paths to legal paths. Label the gates as in Figure 96.

This gives mapping of directions as follows:

$$\phi_* : a \mapsto b \mapsto c \mapsto a, \qquad A \mapsto B \mapsto C \mapsto B$$

The only nontrivial gate is $\{A, C\}$. Legal turns map to legal turns, and each edge maps to a legal path, so this completes the verification that ϕ sends legal paths to legal paths.

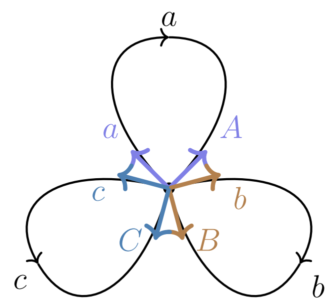

FIGURE 96

As in Example 23.1, there exists a metric that makes ϕ a train-track map with slope the largest eigenvalue of $\left(\begin{smallmatrix} 0&0&1\\ 1&0&1\\ 0&1&0 \end{smallmatrix}\right)$, which is about 1.32. Proposition 22.9 then gives $\tau_\Phi \approx \log(1.32)$.

Now consider the inverse:

$$\Phi^{-1} : \begin{cases} a \mapsto ca^{-1} \\ b \mapsto a \\ c \mapsto b \end{cases}$$

Let $\psi : R_3 \to R_3$ be the associated map. This gives

$$\psi_* : a \mapsto c \mapsto b \mapsto a, \qquad C \mapsto B \mapsto A \mapsto a.$$

We can give this a train-track structure with gates $\{b, A\}, \{c, B\}, \{C, a\}$, and we can check that ψ sends legal paths to legal paths. The slope is the largest eigenvalue of $\left(\begin{smallmatrix} 1&1&0\\ 0&0&1\\ 1&0&0 \end{smallmatrix}\right)$. So we have $\tau_{\Phi^{-1}} \approx \log(1.47) \neq \log(1.32) \approx \tau_\Phi$.

Remark 28.10. If $\phi : \Gamma \to \Gamma'$ with $\Gamma \neq \Gamma'$, the only train track structure we will ever use is $d_1 \sim d_2$ when $\phi_*(d_1) = \phi_*(d_2)$. If $\Gamma = \Gamma'$, one can alternatively use $d_1 \sim d_2$ when there exists $p > 0$ s.t. $\phi_*^p(d_1) = \phi_*^p(d_2)$. This is often better, as it automatically sends legal turns to legal turns, and all the results we proved using the first train track structure also work using the second.

Even using the second train track structure, it needs to be checked that for each vertex v there are at least 2 gates, and that each edge maps to a legal path.

Remark 28.11. We have given many examples of interesting $\Phi \in Out(F_n)$. Here are some that we have omitted:

(1) There exists Φ hyperbolic, with Φ^{-1} parabolic.
(2) There exists Φ reducible; but its restriction to invariant subgraphs always has smaller translation length.

29. PAUSE FOR REVIEW AND BIG PICTURE (11/06, RE, AB)

In this section we will review the various topics and results that we have studied and compiled in the course thus far. In Lecture 4, we proved an important result (Theorem

4.5) due to Stallings which stated that any morphism of graphs $\Gamma \to \Gamma'$ admits a factorization

$$\Gamma = \Gamma_0 \to \Gamma_1 \to \cdots \to \Gamma_k \to \Gamma' \tag{2}$$

where each $\Gamma_i \to \Gamma_{i+1}$ (for $i = 0, \ldots, k-1$) is a fold (cf. Definition 4.1) and the map $\Gamma_k \to \Gamma'$ is an immersion of graphs. Our intuition behind interpreting this result is that it should be thought of as being analogous to the procedure of row reduction from linear algebra. It should be noted that the factorization above is not unique, although it is unique in the sense that Γ_k is isomorphic to the core of the cover of Γ' corresponding to the image of $\pi_1(\Gamma)$ in $\pi_1(\Gamma')$ (see Corollary 7.6 for further details).

The factorization in (2) was an ingredient in the proof of Nielsen's Theorem (Theorem 9). There were three versions of this theorem—although most plainly Nielsen's theorem asserts that $\mathrm{Aut}(F_n)$ (for $n > 1$) is generated by the collection of signed permutations together with the automorphism f defined on generators a_i by

$$f(a_i) = \begin{cases} a_1 a_2 & i = 1, \\ a_i & i > 1. \end{cases}$$

An important consequence of Nielsen's result was stated in Corollary 9.5 which establishes an isomorphism $\mathrm{Out}(F_2) \cong \mathrm{GL}_2(\mathbb{Z})$.

In Lecture 10 and in successive lectures we investigated the structure of outer space, CV_n, which was the space of marked metric graphs. We also defined in (14.3) reduced outer space cv_n and proved that it was a particular deformation retraction of CV_n. There was a particularly nice topology that we imposed on CV_n which made it locally finite; moreover it could be realized as the complement $X \backslash X_0$ in a simplicial complex X of a subcomplex X_0 (see the discussion following Definition 12.9).

We also developed a fairly good understanding of CV_n when $n = 2$. In this case CV_2 was described in terms of the Farey graph wherein we established a homeomorphism

$$\mathrm{cv}_2 \cong \mathcal{F}_\Delta \backslash (\mathcal{F}_\Delta)_0,$$

where $\mathcal{F}_\Delta$ was the Farey complex and $(\mathcal{F}_\Delta)_0$ was its 0-skeleton (i.e. set of vertices). Another direction we pursued was understanding the structure of $\mathrm{Out}(F_n)$ by studying its action on CV_n. Indeed, in Lecture 15 we defined and studied an action of $\mathrm{Out}(F_n)$ on CV_n and proved that the stabilizer of a point in CV_n was finite (cf. Corollary 15.4). Furthermore, in Lemma 15.3, we gave an explicit description of these stabilizers.

Lectures 16, 17, and 18 were devoted to proving that CV_n was contractible (Theorem 15.7). The strategy of the proof was to exhibit a deformation retraction of CV_n onto the simplex of the rose R_n afforded by greedy folding paths. In the course of Lecture 18 we further defined the so-called "spine" K_n of CV_n—it is a particular simplicial complex of dimension $2n-3$. Like CV_n, K_n is a contractible space. The spine has the advantage that the space $\mathrm{Out}(F_n)$-orbits of K_n was shown to be compact. We used this fact to show that $\mathrm{Out}(F_n)$ was finitely presented in Corollary 18.4.

Starting in Lecture 19, we defined an "asymmetric" metric d on CV_n which was defined as

$$d(\Gamma, \Gamma') = \inf_\phi \log \sigma(\phi) \tag{3}$$

where ϕ ranged over all difference of marking maps $\phi\colon \Gamma \to \Gamma'$ and $\sigma(\phi)$ denoted the Lipschitz constant of ϕ. Via the notion of an optimal difference of marking map we observed that the infimum of (3) was realized. And in Corollary 19.15 we further characterized d by

$$d(\Gamma, \Gamma') = \log \max_{[\alpha]} \frac{\ell_\alpha(\Gamma')}{\ell_\alpha(\Gamma)}, \tag{4}$$

where $[\alpha]$ was the conjugacy class of an element $\alpha \in F_n$ and $\ell_\alpha(\Gamma)$ was the length of the unique immersed loop in the free homotopy class associated to α. The characterization above was useful as there was an algorithm that allowed us to compute the quantities on the right hand side of (4). The asymmetric metric established further rich geometric properties of CV_n. For example, using the metric we could define geodesics in CV_n which were particular paths obeying a certain isometry condition (Definition 21.3). Perhaps more markedly, we used the metric to qualitatively understand the behavior of the group action of $\mathrm{Out}(F_n)$ on CV_n through a program initiated in Lecture 22.

The abstract setup was that if (X, d) was an asymmetric metric space, $\Phi \in \mathrm{Isom}(X)$, then we could define the translation length $\tau(\Phi) = \inf\{d(x, \Phi(x))\colon x \in X\}$. Using the translation length, we could classify isometries into as one of three categories:

- (Elliptic) Isometries with fixed points;
- (Hyperbolic) Isometries for which the translation length τ was positive and is realized;
- (Parabolic) Isometries for which the translation length τ was not realized.

We spent a considerable amount of effort to understand and classify the elements of $\mathrm{Out}(F_n)$ when viewed as elements of $\mathrm{Isom}(\mathrm{CV}_n)$. Our primary results are tabulated below. Fix $\Phi \in \mathrm{Out}(F_n)$.

- Φ is elliptic if and only if it has finite order (Lemma 22.5).
- If Φ is parabolic then it is reducible (Theorem 24.1). The converse however is false.
- If Φ is hyperbolic then there is a map $\phi\colon \Gamma \to \Gamma$ representing Φ that is a train track map on its tension graph (Theorem 24.3).
- If Φ is irreducible then it is hyperbolic.
- Φ is reducible if and only if F_n admits a decomposition

$$F_n = A_1 * \cdots * A_n * C \tag{5}$$

  where the A_i's are nontrivial and Φ permutes the conjugacy classes $[A_i]$ (Proposition 27.1).

We also observed some "pathologies" as well. For example, we exhibited $\Phi \in \mathrm{Out}(F_n)$ such that $\tau_\Phi \neq \tau_{\Phi^{-1}}$ which in truth was not terribly surprising since we are working in an asymmetric space. Another example was an irreducible element Φ with the property that Φ^3 was reducible. Moreover, we also saw examples of reducible elements whose associated decomposition of F_n as in (5) was non-canonical.

Overall, we have made strides in understanding free groups, $Out(F_n)$ as a group, and individual elements of $Out(F_n)$. In terms of the big picture, our results have ratified the analogies to mapping class groups and lattices in Lie groups, and we have build up methods that more broadly are useful for importing ideas and results to $Out(F_n)$ via these analogies. Furthermore we are implementing the main insight of geometric group theory, which is that often groups can be studied via the metric spaces they act on. Future lectures will give other spaces (different from outer space) that $Out(F_n)$ acts on, allowing an even richer application of ideas from geometric group theory. Free factors will play a major role in this, so next we will spend some time investigating when a subgroup $A \subset F_n$ is a free factor.

30. Deciding if a subgroup is a free factor (11/08, AB, RE)

Our goal here is to find an algorithm to determine whether a given finitely generated subgroup $A < F_n$ is a free factor. We partially follow [HW19, Asc].

Recall that A is a free factor if there exists another subgroup $B < F_n$ such that $F_n = A * B$. Free factors are vastly better behaved than arbitrary finitely generated subgroups, and in the comparison to linear algebra they provide a reasonable analogue of a vector subspace. It may help to keep in mind that vector subspaces always have complements.

We say that a single element $g \in F_n$ is *primitive* if it generates a free factor, or equivalently if it is part of a basis for F_n. We begin by (re)introducing a few notions that will help us.

Definition 30.1. Let $A < F_n$ be a finitely generated subgroup. The *volume* of A is the number of edges in the core of the (unpointed) A-cover of R_n.

Exercise 44. *For $g \in F_n$, we have that $\mathrm{vol}(\langle g \rangle)$ is the length of the cyclic reduction of the word corresponding to g. (Recall that the cyclic reduction is obtained from a word by removing subwords of the form $a_i a_i^{-1}$ and $a_i^{-1} a_i$ for each i, and also by replacing words of the form $a_i w a_i^{-1}$ or $a_i^{-1} w a_i$ by w, i.e. it is a shortest word corresponding to an element in the conjugacy class of g. Up to cyclic permutation, it depends only on the conjugacy class.)*

Definition 30.2. An automorphism φ of F_n is *Whitehead of the second kind* if there is $m \in \{a_j^{\pm 1}\}_{j=1,\dots,n}$ such that

$$\varphi(a_i) \in \{a_i, m^{-1}a_i, a_i m, m^{-1} a_i m\}$$

for each i.

Next, we state the key theorem that will allow us to produce the algorithm.

Theorem 30.3. *If $A < F_n$, with $A \neq F_n$, and A is a free factor, but A is not conjugate to a subgroup of $\langle a_1, \dots, \hat{a}_i, \dots, a_n \rangle$ for any i (here this denotes the subgroup of F_n generated by the a_j for $j \neq i$), then there exists a Whitehead automorphism of the second kind φ such that $\mathrm{vol}(\varphi(A)) < \mathrm{vol}(A)$.*

Assuming this theorem, we can describe the algorithm. We begin with generators $g_1, \dots, g_p$ for A, which we do not assume form a basis:

(1) Fold the map from the graph in Figure 97 to R_n to compute the core Γ of the A cover of R_n. If $\Gamma = R_n$, then $A = F_n$, so A is trivially a free factor. If $\Gamma \to R_n$

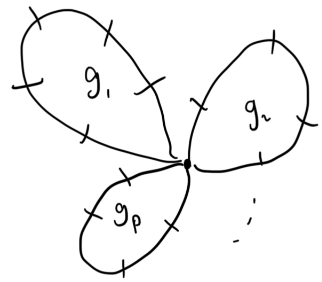

FIGURE 97

is not surjective, then A is conjugate into $\langle a_1, \ldots, \hat{a}_i, \ldots, a_n\rangle$ for some i. Recall that by Lemma 26.6, if $A < C < F_n$, and C is a free factor, then A is a free factor of F_n if and only if it is a free factor of C. Therefore, we may repeat the algorithm with F_n replaced by $\langle a_1, \ldots, \hat{a}_i, \ldots, a_n\rangle \cong F_{n-1}$, and A replaced by its conjugate lying in $\langle a_1, \ldots, \hat{a}_i, \ldots, a_n\rangle$, and we are done by induction on n. If $\Gamma \to R_n$ is surjective but not an isomorphism, move onto step 2.

(2) For each of the finitely many Whitehead automorphisms of the second kind φ, compute $\mathrm{vol}(\varphi(A))$. If $\mathrm{vol}(\varphi(A)) \geqslant \mathrm{vol}(A)$ for all φ, then A is not a free factor by Theorem 30.3. Otherwise, pick φ such that $\mathrm{vol}(\varphi A) < \mathrm{vol}(A)$. Since A is a free factor if and only if $\varphi(A)$ is, we may replace A by $\varphi(A)$, and return to step 1. Since the volume has decreased, we are making progress, and this algorithm will terminate in finite time.

Remark 30.4. The above is a conceptual proof that an algorithm exists. When implementing the algorithm, one does not however actually search over all φ as suggested above. Instead one uses the results of of Exercise 48 to quickly determine which φ, if any, will reduce the volume.

Now, we turn to the proof of Theorem 30.3. First, we introduce a helpful definition.

Definition 30.5. An *almost rose* is a directed graph Γ whose edges are labelled by symbols $\{a_i\}$ such that:

(1) There is a unique index $i \in \{1, \ldots, n\}$ such that there are exactly two edges labelled by a_i, and for $j \neq i$, there is exactly one edge labelled by a_j.
(2) Γ has exactly two vertices, which we call u and v.
(3) One of the two edges labelled a_i is a loop with vertex u, and the other connects u and v.
(4) Γ has no hanging tree, that is, there is a $j \neq i$ such that the edge labelled by a_j either connects u to v, or is a loop with vertex v.

Up to signed permutation, every almost rose looks like the graph in Figure 98 for some $1 \leqslant k \leqslant \ell \leqslant n$, with $k < n$.

Lemma 30.6. *If A is a free factor, and Γ is the core of the A-cover of R_n, there exists a label preserving map of Γ to an almost rose.*

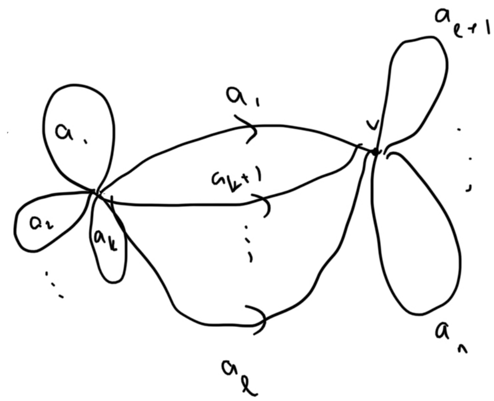

FIGURE 98. (The left vertex is meant to be labelled u)

Proof. Write $F_n = A * B$, and pick a basis $z_1, \dots, z_q$ for B. Let $\hat{\Gamma}$ be a graph obtained by adjoining q subdivided loops to Γ such that there is an induced morphism $\hat{\Gamma} \to R_n$ of graphs, where the morphism is defined on each of the q subdivided loops by the words corresponding to the z_j's. See Figure 99 for an example. Then, we fold the map $\hat{\Gamma} \to R_n$ to obtain a factorization:

$$\hat{\Gamma} = \hat{\Gamma}_0 \to \cdots \to \hat{\Gamma}_s \xrightarrow{\cong} R_n.$$

Note that the last map is an isomorphism by Exercise 45 below. Since the rank never goes down, each of these folds are either type I or type II. Therefore, $\hat{\Gamma}_{s-1}$ is a single rank preserving fold away from being a rose, so it has two vertices, and $\hat{\Gamma}_{s-1} \to \hat{\Gamma}_s$ is a type II fold (note that type I folds can only exist when the graph has at least three vertices). Finally, note that $\hat{\Gamma}_{s-1}$ has no hanging tree, since $\hat{\Gamma}$ has no hanging trees, and type I and II folds cannot introduce new hanging trees. Therefore, $\hat{\Gamma}_{s-1}$ is an almost rose, and $\Gamma \to \hat{\Gamma} \to \hat{\Gamma}_{s-1}$ is our desired map. □

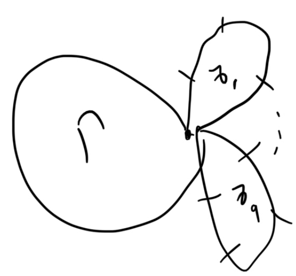

FIGURE 99. The graph $\hat{\Gamma}$

This proof used the exercise:

Exercise 45. *The map $\hat{\Gamma} \to R_n$ is an isomorphism on π_1.*

We also will need the following fact:

Exercise 46. *Let f be the map of Γ to an almost rose as in Lemma 30.6. Then, $f(\Gamma)$ has two vertices and no hanging tree. (See Hint A.19.)*

Now, we begin the proof of Theorem 30.3.

Proof of Theorem 30.3. Let f be the map of Γ to an almost rose as in Lemma 30.6. By applying a signed permutation if necessary, we may assume the almost rose is labelled as in Figure 98. Let e_0 be an edge of $f(\Gamma)$ connecting u, v, such that e_0 is labelled by a_1 whenever the unique edge labelled a_1 connecting u, v is in $f(\Gamma)$, and if this edge is not in $f(\Gamma)$, we can pick an arbitrary such e_0. In either case, let a_j be the label of e_0.

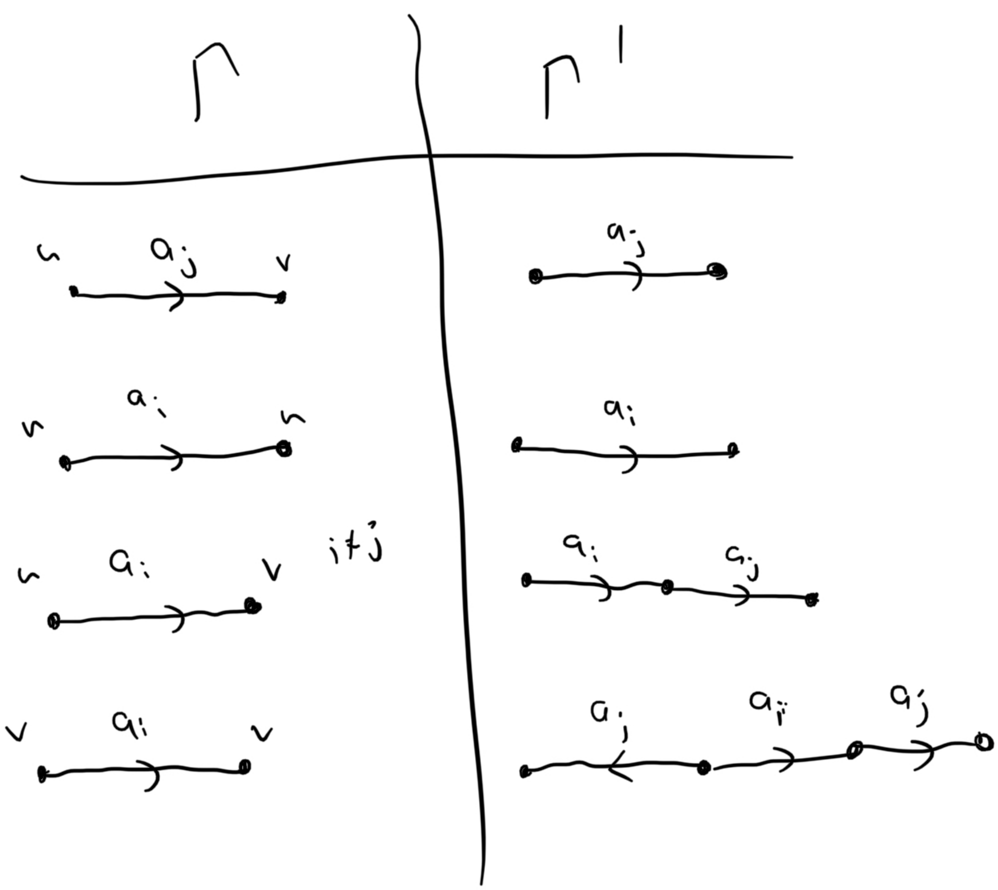

FIGURE 100. The labels of the vertices indicate whether the vertex maps to u or v in the almost rose. For each half edge entering or leaving a vertex labelled v, we are adding an edge labeled a_j pointing towards that vertex; the one exception is if the edge is already labelled a_j. So in a sense we're changing the graph "near the preimage of v" only.

Now, we produce a graph Γ' from Γ by applying the following replacement rules to each edge e of Γ:

(1) If $f(e) = e_0$, then e is unchanged in Γ'.
(2) If $f(e)$ is a loop with vertex u, then e is unchanged in Γ'.
(3) If $f(e) \neq e_0$ but connects u and v, then e is replaced by two edges labelled $a_i a_j$.
(4) If $f(e)$ is a loop with vertex v, e is replaced by three edges labelled $a_j^{-1} a_i a_j$.

This situation is summarized in Figure 100. The labels on Γ' give rise to a map $f' : \Gamma' \to R_n$. The rest of the proof will be explained next time. □

We will need the following easy exercise:

Exercise 47. *There is a Whitehead automorphism of the second kind* φ *such that* $\operatorname{im}(f'_*) = \varphi(A)$.

Our general strategy will be to fold f', and show that $\operatorname{vol}(\varphi(A)) < \operatorname{vol}(A)$.

31. Free factor algorithm, free factor complex (11/10, AB, JG)

Let's return to the proof of Theorem 30.3.

Proof of Theorem 30.3 (continued). By our choice of j (namely, requiring that $a_j = a_1$ if the edge labelled a_1 connecting u, v in the almost rose is in the image of f) we have $\operatorname{im}(f'_*) = \varphi(A)$. (If that edge was in the image but we used a $j \neq 1$, then the replacement rule above would effect the edges labelled a_1 that map to the edge from u to v, but not the edges labelled a_1 that map to a loop at u, and so we wouldn't have a well-defined φ with $\operatorname{im}(f'_*) = \varphi(A)$.)

Now, fold the morphism $f' : \Gamma' \to R_n$ and chop off hanging trees to get an immersion $f'' : \Gamma'' \to R_n$ with $\operatorname{im}(f''_*) = \varphi(A)$. By Lemma 7.4, the map $\Gamma'' \to R_n$ is the core of the $\varphi(A)$-cover of R_n. To conclude, it suffices to show that Γ'' has fewer edges than Γ (since φ is a Whitehead automorphism).

Let's think about the preimages of v in Γ. If $\hat{v}_0$ is such a vertex, note that $\hat{v}_0$ can have an incoming a_j edge but no outgoing a_j edge. Let $\hat{v}$ denote the vertex of Γ' corresponding to $\hat{v}_0 \in \Gamma$. Note that by construction, every edge adjacent to $\hat{v} \in \Gamma''$ is an incoming a_j edge. Since we may fold edges and remove hanging trees in any order we like, let's consider the graph G obtained from Γ' by only performing the following two operations:

(1) For each vertex $\hat{v}$ of Γ' corresponding to a vertex of Γ mapping to v, fold all the edges bordering $\hat{v}$ (which are incoming a_j edges) to produce a graph G_0.
(2) Chop off the (hanging) edge of G_0 corresponding to each of the $\hat{v}$, and call this graph G.

The situation is summarized in Figure 101. The graph Γ'' is obtained from G by performing further folds and removing hanging trees, so it has fewer edges than G. Therefore, it suffices to show that G has fewer edges than Γ.

Let's count precisely how many edges have been added and removed to produce G. The edges added to Γ' are in bijection with pairs $(\hat{v}, e)$, where $\hat{v}$ is a vertex of Γ mapping to v, and e is an (oriented) edge of Γ based at $\hat{v}$ which is not labelled by a_j. Let α be the number of these vertices. Let β be the number of edges labelled a_j which do not

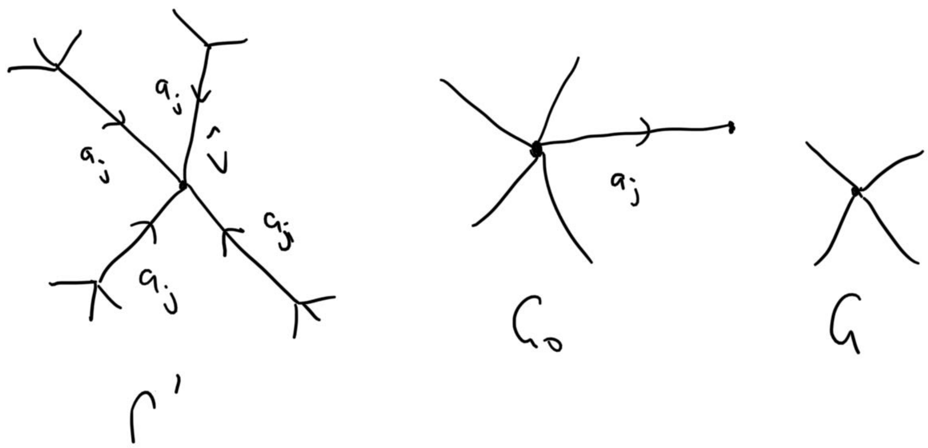

FIGURE 101

map to loops at u under f, and let γ be the number of edges added in the passage from Γ to Γ'. In summary

$$\begin{aligned} \alpha &= \text{number of } \hat{v} \\ \beta &= \text{number of } a_j \text{ edges mapping to an edge from } u \text{ to } v \\ \gamma &= \text{number of } (\hat{v}, e) \\ &= \text{number of edges added to } \Gamma \text{ to get } \Gamma'. \end{aligned}$$

So, to start we add γ edges to Γ to get Γ'. In step 1, we remove $\beta + \gamma - \alpha$ edges via the folding. In step 2, we remove α edges. Therefore, β is the difference between the number of edges of G and the number of edges of Γ. Since we chose a_j to label an edge in the image of f, $\beta > 0$, and so we are done. □

Remark 31.1. Here is come context on Theorem 30.3:

(1) This theorem is very closely related to Whitehead's "Cut Vertex Lemma." See Exercise 48.
(2) It is even more standard to prove this result using spheres in three manifolds. There is also a third approach using the boundary of the free group [Mar95, Appendix B].
(3) The algorithm based on this theorem is often called the "easy Whitehead algorithm." This also goes under the name "peak reduction".
(4) There is also a "hard Whitehead algorithm," which allows one to compute whether any given $w, w' \in F_n$ lie in the same orbit of $\mathrm{Aut}(F_n)$.
(5) There are also many other variations of these Whitehead algorithms. The moral is that "greedy" algorithms using Whitehead automorphisms are shockingly effective.

Exercise 48. *Let A be a finitely generated subgroup of F_n, and let Γ be the core of the A cover of the standard rose R_n. Define the Whitehead graph of A as follows:*

- *There are $2n$ vertices, labelled by $\{a_1, a_1^{-1}, \ldots, a_n, a_n^{-1}\}$.*
- *There is an edge from vertex v to vertex w if there a vertex of Γ with an outgoing edge labelled v and an outgoing edge labelled w. (There can also be more outgoing edges. The edges of Γ aren't really oriented, so what we really mean by "outgoing edge labelled v" is that there is an adjacent edge, such that when you oriented it away from v it maps to the rose according to its label $v \in \{a_1, a_1^{-1}, \ldots, a_n, a_n^{-1}\}$. If an edge is a loop based at v there are two ways to orient it away from v.)*
- *By definition of Γ, since $\Gamma \to R_n$ is an immersion, the Whitehead graph does not contain any loops based on a single vertex. We do not add multiple edges, so it is a graph in the strictest sense.*

Note that $\Gamma \to R_n$ is onto if and only if A is not conjugate into a subgroup defined by at most $n-1$ of the basis elements.

(1) *Show that $\Gamma \to R_n$ is not onto if and only if there is some i such that both a_i and a_i^{-1} are isolated vertices.*
(2) *Suppose that $\Gamma \to \Gamma'$ is a label preserving morphism of graphs. Show that the Whitehead graph of Γ is a subgraph of the Whitehead graph of Γ'.*
(3) *Show that the Whitehead graph of the almost rose illustrated in Figure 98 is as illustrated in Figure 102. Note in particular that if we remove the vertex labelled*

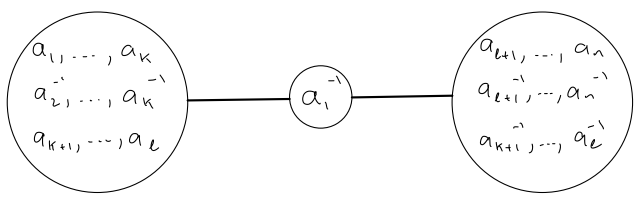

Figure 102. A schematic of the Whitehead graph of the almost rose illustrated in Figure 98. Two vertices should be joined by an edge if and only if they are either in the same bubble or their bubbles are joined by an edge in this schematic. Thus, in the schematic, each bubble represents a complete subgraph, and each edge represents a join. The Whitehead graph is, in this case, two complete subgraphs which overlap at a_1^{-1}. The vertex a_1^{-1} is a cut vertex.

a_1^{-1} from Whitehead graph of the almost rose, the Whitehead graph becomes disconnected. That is, a_1^{-1} is a cut vertex.

(4) *Deduce that if Γ maps to an almost rose, then its Whitehead graph has a cut vertex.*
(5) *Conversely, show that if the Whitehead graph of Γ has a cut vertex, then Γ has a map to an almost rose.*

In summary, Γ has a map to an almost rose if and only if its Whitehead graph has a cut vertex.

Exercise 49. *Suppose that A is a finitely generated subgroup of F_n that is contained in some free factor. Show that the conclusion of Theorem 30.3 holds for A. (See Hint A.20.)*

Our next goal is to introduce a geometric space parameterizing all conjugacy classes of free factors.

Definition 31.2. A free factor is *proper* if it is neither trivial nor all of F_n.

Definition 31.3. Given two conjugacy classes of free factors $[A], [B]$, say $[A] < [B]$ if a representative of $[A]$ is contained in a representative of $[B]$. (Equivalently, if there exists g with $gAg^{-1} \subset B$.)

Now we are ready to define the free factor complex.

Definition 31.4. If $n > 2$, the *complex of free factors* FF_n is the simplicial complex whose vertices are parameterized by all conjugacy classes of proper free factors of F_n, and $([A_1], \ldots, [A_k])$ span a $(k-1)$-simplex if, up to reordering, we have:

$$[A_1] < \cdots < [A_k].$$

When $n = 2$, we define FF_2 to be a graph whose vertices are conjugacy classes of free factors of rank 1, and $[A], [B]$ span an edge if $F_2 = A * B$ for certain representatives of these conjugacy classes.

In general, we define the *free factor graph* $FF_n^{(1)}$ to be the 1-skeleton of FF_n.

Remark 31.5. We already proved that FF_2 is the Farey graph; see Remark 14.14.

Exercise 50. *Describe the subset of FF_n given by all free factors generated by a subset of the standard basis. (See Hint A.21.)*

Exercise 51. *Let $\mathcal{P}_n$ be the graph which has as a vertex for each conjugacy class of pair $\{x, x^{-1}\}$ with x primitive, and an edge from $[\{x, x^{-1}\}]$ to $[\{y, y^{-1}\}]$ if there is a basis for F_n which contains x and a conjugate of y. Show $\mathcal{P}_n$ is quasi-isometric to FF_n. (See Hint A.22.)*

We define a metric on $FF_n^{(1)}$ by declaring each edge to have length 1.

Remark 31.6. $\mathrm{Out}(F_n)$ acts simplicially on FF_n, hence it acts on $FF_n^{(1)}$.

Let $\mathcal{S}$ be the set of finite sets of vertices of FF_n. There is a map $\pi : CV_n \to \mathcal{S}$ given by

$$\pi(\Gamma) = \{\pi_1(\Gamma_0) : \Gamma_0 \subset \Gamma \text{ a connected subgraph with } \pi_1(\Gamma_0) \text{ proper}\}.$$

Let ∞, D, Θ denote the figure-eight, dumbbell, and theta graphs, respectively. These represent the only homeomorphism classes of connected graphs whose fundamental groups have rank 2 and no hanging trees. For $X \in \{\infty, D, \Theta\}$, and a, b a basis of F_2, let $X_{a,b}$ be a copy of X, with a marking given by labelling two oriented edges by a and b. Then, we can see the following:

- $\pi(\Theta_{a,b}) = \{[\langle a\rangle], [\langle b\rangle], [\langle ab^{-1}\rangle]\}$
- $\pi(D_{a,b}) = \pi(\infty_{a,b}) = \{[\langle a\rangle], [\langle b\rangle]\}$.

This is most of the proof of the following:

Exercise 52. *If $\Gamma \in CV_2$, $\pi(\Gamma)$ is either:*

- *two vertices joined by an edge*
- *three vertices, where each pair are joined by an edge.*

In particular, $\mathrm{diam}(\pi(\Gamma)) = 1$.

32. Bounding distance in FF_n (11/13, KS, SL)

Lemma 32.1. $\mathrm{diam}(\pi(\Gamma)) \leqslant 4$ *for all* $\Gamma \in CV_n$.

Proof. Assume $n \geqslant 3$. Say Γ_A, Γ_B are subgraphs of Γ with $\pi_1(\Gamma_A) = A, \pi_1(\Gamma_B) = B$. Pick non-separating edges a, b with $a \notin \Gamma_A, b \notin \Gamma_B$. Since $n \geqslant 3$, $\Gamma - (a \cup b)$ is not a forest. Let Γ_C be a non-tree component of $\Gamma - (a \cup b)$ and let $C = \pi_1(\Gamma_C)$.

Since $\Gamma_A \subset \Gamma - a$, we have $[A] < [\pi_1(\Gamma - a)]$, so there is an edge in the complex of free factors, hence $d(A, \pi_1(\Gamma - a)) \leqslant 1$. (This distance is 0 if $\Gamma_A = \Gamma - a$.) So

$$\begin{aligned} d(A,B) &\leqslant d(A, \pi_1(\Gamma - a)) + d(\pi_1(\Gamma - a), C) + d(C, \pi_1(\Gamma - b)) + d(\pi_1(\Gamma - b), B)) \\ &\leqslant 1 + 1 + 1 + 1 = 4, \end{aligned}$$

as desired. □

Definition 32.2. $x \in F_n$ is *simple* if it is contained in a proper free factor.

Example 32.3. $a_1 a_2 a_1^{-1} a_2^{-1} \in F_2$ is not simple.

Remark 32.4. Primitive implies simple.

Example 32.5. $a_1^3 a_2^3 \in F_3$ is simple but not primitive.

Exercise 53. *Show every simple element of F_n can be written as a product of two primitive elements. (See Hint A.23.)*

Exercise 54. *Show that if $w \in F_n$ is cyclically reduced and has all possible $2n(2n-1)$ possible subwords of length 2, then w is not simple. (See Hint A.24.)*

Remark 32.6. A slight variation on Lemma 26.6 shows that the intersection of free factors is a free factor. See the discussion after that lemma for a reference.

Definition 32.7. If x is simple, let $\dot{x} \in FF_n^{(0)}$ be the conjugacy class of the intersection of free factors containing x.

Remark 32.8. In general $\dot{\alpha} \neq \langle \alpha \rangle$. For example, $\alpha = a_1^3 a_2^3 \in F_3$ but $\dot{\alpha} = \langle a_1 a_2 \rangle$.

Definition 32.9. If $x \in F_n$, $\Gamma \in CV_n$, let $x|\Gamma$ be the (unbased) immersed loop in Γ with conjugacy class x.

The following is an improved version of [BF14a, Lemma 3.2] with a new proof.

Proposition 32.10. *Take $\Gamma \in CV_n$, e an edge of Γ, and let $\alpha \in F_n$ be simple such that $\alpha|\Gamma$ crosses e at most k times. Then there exists $A \in \pi(\Gamma)$ with*

$$d(A, \dot{\alpha}) \leqslant \max\left(2 \lfloor \log_2 k \rfloor + 5, 1\right).$$

Remark 32.11. The proposition is already interesting and useful if α is primitive, in which case $\dot{\alpha} = \langle \alpha \rangle$. See Remark 32.12 for how the proof can be made slightly simpler in this case.

Proof. If $k = 0$, take $A = \pi_1(\Gamma - e)$. In this case $d(A, \dot{\alpha}) \leqslant 1$, since $\alpha \in \pi_1(\Gamma - e)$.

If $k \geqslant 1$ we will use the following statement.

Exercise 55. *There exists a marked graph Γ' and a morphism $f : \Gamma' \to \Gamma$ that is an isomorphism on π_1, such that $\alpha|\Gamma'$ is contained in a proper subgraph and intersects $f^{-1}(e)$ at most k times. (See Hint A.25.)*

Let Γ'_0 be a proper subgraph of Γ' containing $\alpha|\Gamma'$. If α is not embedded, we view (the image of) its immersed representative $\alpha|\Gamma'$ as a subgraph of Γ', and let β be embedded loop in $\alpha|\Gamma' \subset \Gamma'$, as in Figure 103. So $\beta \in F_n$ is primitive and

$$d(\langle\beta\rangle, \dot{\alpha}) \leqslant d(\dot{\alpha}, \pi_1(\Gamma'_0)) + d(\pi_1(\Gamma'_0), \langle\beta\rangle) \leqslant 1 + 1 = 2.$$

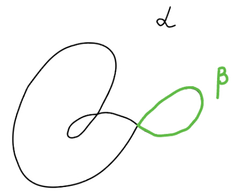

FIGURE 103

Note that $\beta|\Gamma'$ intersects $f^{-1}(e)$ at most k times, since $\beta|\Gamma' \subset \alpha|\Gamma'$ and $\alpha|\Gamma'$ intersects $f^{-1}(e)$ at most k times.

Remark 32.12. If α was primitive, we could slightly simplify the above, for example by taking Γ' to be a subdivided rose labelled by a basis containing α, and taking $\beta = \alpha$.

Let us fold to get

$$\Gamma' = \Gamma'_0 \to \Gamma'_1 \to \cdots \to \Gamma'_k = \Gamma,$$

and let $f_i : \Gamma'_i \to \Gamma$ be the induced maps.

Let $i_0 = 0$ and i_1 be the first time when β stops being embedded (that is, the smallest i such that $\beta|\Gamma'_i$ is not embedded). The only situation when that occurs is when $\beta|\Gamma'_{i_1}$ is figure eight loop, since $\beta|\Gamma'_{i_1-1}$ is embedded and a fold identifies just one pair of vertices. Let β_1 be one of the two embedded subloops of $\beta|\Gamma'_{i_1}$ whichever has at most $k/2$ preimages of e ($\beta|\Gamma'_{i_1}$ has at most k preimages).

Continue folding until β_1 stops being embedded and repeat. Stop at t_l with β_l disjoint from $f_{t_l}^{-1}(e)$, so $l \leqslant \lfloor \log_2 k \rfloor + 1$. Note that $d(\langle\beta_j\rangle, \langle\beta_{j+1}\rangle) \leqslant 2$, since in $\Gamma'_{t_{j+1}}$ both β_j and β_{j+1} are contained in the figure eight subgraph defined by β_j. Let A be the free factor defined by the component of $\Gamma - e$ containing $\beta_\ell|\Gamma$, so $d(A, \langle\beta_\ell\rangle) \leqslant 1$. Overall,

$$\begin{aligned} d(A, \dot{\alpha}) &\leqslant d(\dot{\alpha}, \langle\beta\rangle) \\ &\quad + d(\langle\beta\rangle, \langle\beta_1\rangle) + \cdots + d(\langle\beta_{l-1}\rangle, \langle\beta_l\rangle) \\ &\quad + d(\langle\beta_l\rangle, A) \\ &\leqslant 2 + 2(\lfloor \log_2 k \rfloor + 1) + 1 \\ &\leqslant 2\lfloor \log_2 k \rfloor + 5. \end{aligned}$$

□

33. $CV_n \to FF_n$, Gromov hyperbolicity (11/15, UP, NL)

Remark 33.1. For any $\alpha \in F_n$ and $\Gamma \in CV_n$, if $\alpha|\Gamma$ crosses each edge of Γ at least k times, then

$$l_\alpha(\Gamma) \geqslant k \cdot \mathrm{vol}(\Gamma) \geqslant k$$

Hence some edge is crossed at most $l_\alpha(\Gamma)$ times.

Corollary 33.2. *If* $\Gamma, \Gamma' \in CV_n$, $d(\Gamma, \Gamma') \leqslant \log \sigma$, *then* $\mathrm{diam}(\pi(\Gamma) \cup \pi(\Gamma')) \stackrel{\sim}{\leqslant} \log \sigma$.

The $\stackrel{\sim}{\leqslant}$ sign indicates less than or equal up to additive and multiplicative error. (We could easily be precise but choose not to bother.)

Proof. Pick $f : \Gamma \to \Gamma'$ optimal so f has slope $\leqslant \sigma$. Let α be a witness of the form of an embedded loop, a figure 8, or a dumbbell (See Figure 104). So α is simple and

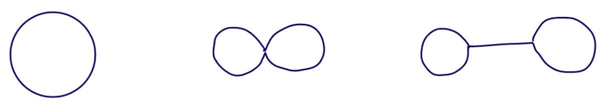

Figure 104. A witness can always be chosen to be take of these forms

$l_\alpha(\Gamma) \leqslant 2$ and so $l_\alpha(\Gamma') \leqslant 2\sigma$. There is an edge e' of Γ' such that $\alpha|\Gamma'$ crosses e' at most 2σ times, so Proposition 32.10 gives a $A \in \pi(\Gamma)$ with

$$d(A, \dot{\alpha}) \stackrel{\sim}{\leqslant} \log(2\sigma).$$

□

One often defines a "distance" between subsets of a metric space by $d(P, Q) = \mathrm{diam}(P \cup Q)$. This definition of distance is symmetric and satisfies the triangle inequality, but the distance from a point to itself may be greater than 0. Thus the corollary says that

$$d(\pi(\Gamma), \pi(\Gamma')) \leqslant k d(\Gamma, \Gamma') + k$$

for some k. We say π is "coarsely Lipschitz".

Definition 33.3. A metric space is *geodesic* if there is at least one geodesic between every pair of points.

Definition 33.4. Let $\delta \geqslant 0$. A geodesic metric space is δ-*hyperbolic* if for any triple of points x, y, z and a triple of geodesics $\gamma_{xy}, \gamma_{yz}, \gamma_{zx}$ joining them, we have

$$\gamma_{xy} \subset N_\delta(\gamma_{yz} \cup \gamma_{zx}).$$

Here, $N_\delta(S) = \{p : \exists q \in S, d(p, q) \leqslant \delta\}$. Figure 105 shows a picture of what this condition looks like. We say triangles are "δ-thin".

Definition 33.5. A space is *Gromov hyperbolic* if it is δ-hyperbolic for some δ.

Example 33.6. $\mathbb{R}^2$ is not Gromov-hyperbolic.

Example 33.7. Trees are 0-hyperbolic. Figure 106 shows a 0-thin triangle in a tree.

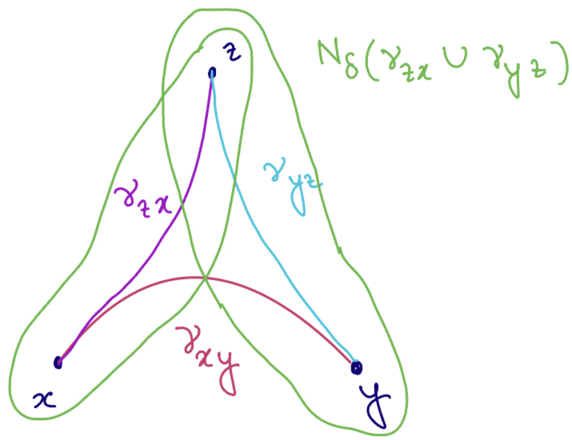

FIGURE 105

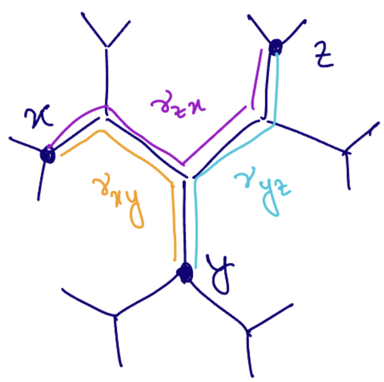

FIGURE 106

Morally speaking, "Gromov-hyperbolic = coarsely tree-like in many ways."

Example 33.8. $\mathbb{H}^2$, and in general the universal cover of a compact, negatively curved manifold, is Gromov hyperbolic.

Exercise 56. FF_2 *(which is also the Farey graph) is 1-hyperbolic.*

The general case is vastly harder and we'll only mention the result without giving the proof [BF14a]. See [Vog15, Section 8] for a nice and short introduction, and see also the discussion after Theorem 36.9.

Theorem 33.9 (Bestvina, Feighn)**.** $FF_n^{(1)}$ *is Gromov hyperbolic.*

Let us now try to very briefly indicate one way such a result can be used.

Remark 33.10. In the world of coarse geometry, the definitions of elliptic, parabolic, and hyperbolic for an isometry $\Phi \in \mathrm{Isom}(X)$ get changed:

- Elliptic $\iff$ there is a bounded orbit
- Parabolic $\iff$ not elliptic and $\widehat{\tau}_\Phi = 0$
- Hyperbolic $\iff$ $\widehat{\tau}_\Phi > 0$.

This is not consistent with the old definition. To mitigate this, we'll use "loxodromic" to say $\widehat{\tau}_\Phi > 0$.

Just like random walks on trees drift, one can get the following result [MT18]:

Theorem 33.11 (Maher-Tiozzo)**.** *Let X be a Gromov hyperbolic (separable) metric space. Let μ be a probability measure on Isom(X) with countable support that generates a non-elementary subgroup. Let $w_n \in$ Isom(X) be sampled after n steps. Then $\exists L$ such that*

$$\mathbb{P}(\widehat{\tau}_{w_n} \leqslant Ln) \to 0 \text{ as } n \to \infty$$

In particular,

$$\mathbb{P}(w_n \text{ is loxodromic }) \to 1$$

Corollary 33.12. *Fix a finite generating set for* $\mathrm{Out}(F_n)$*. Let w_n be the result of composing n randomly chosen generators. Then*

$$\mathbb{P}(w_n \text{ is fully irreducible}) \to 1$$

as $n \to \infty$.

Proof sketch ignoring "non-elementary". If w_n is not fully irreducible, then $\exists p$ such that w_n^p is reducible. So there exist $[A_1], \ldots, [A_k] \in FF_n^{(0)}$ permuted by w_n^p. So $\widehat{\tau}(w_n^p) = 0$, so $\widehat{\tau}(w_n) = 0$. □

We haven't defined what non-elementary means, but let us at least mention that to have a non-elementary subgroup of isometries the space must have infinite diameter. We haven't shown that FF_n has infinite diameter if $n > 2$.

Exercise 57. *A map $f : X \to Y$ between metric spaces is coarsely Lipschitz if there is a $C > 0$ such that*

$$d(f(x), f(x')) \leqslant Cd(x, x') + C$$

for all x, x'. A map $g : Y \to X$ is a coarse inverse to f if there is a $C > 0$ such that $d(x, g(f(x))) \leqslant C$ for all x and $d(y, f(g(y))) \leqslant C$ for all y.

Show that if f and g are coarsely Lipschitz and are coarse inverses to each other, then both are quasi-isometries.

Exercise 58. *Let X be a (possibly asymmetric) metric space and let Γ be a graph. Suppose there is a map π from X to finite subsets of vertices of Γ, such that $\pi(x)$ always has diameter at most C for some $C > 0$. (Optional: Start off with the special case when $\pi(x)$ is always a singleton.) For v a vertex of Γ, define $\pi^{-1}(v)$ to be the set of x with $v \in \pi(x)$.*

Let E be defined by starting with X, and for each vertex v of Γ, adding a "cone point" c_v. For each point $x \in \pi^{-1}(v)$ also add a unit length segment joining x to c_v. (The space E is sometimes called an "electrification" of X. It has the feature that it contains X, and each $\pi^{-1}(v)$ now has diameter at most 2.)

Suppose that, for all pairs of adjacent vertices v, w, the intersection of $\pi^{-1}(v)$ and $\pi^{-1}(w)$ is non-empty. Show that the induced map $E \to \Gamma$ is a quasi-isometry. (If you don't like that $\pi(x)$ is a finite set, define a $\hat{\pi}$ by setting $\hat{\pi}(x)$ to be an arbitrarily chosen element of $\pi(x)$, and show that $\hat{\pi}$ is a quasi-isometry.)

Apply this to $X = CV_n$ and $\Gamma = FF_n$ to conclude that the free factor graph is quasi-isometric to an electrification of CV_n. (See Hint A.26.)

34. Groups acting on trees (11/17, SY, JG)

Definition 34.1. A group action on a tree is called **minimal** if there is no proper invariant subtree.
The action has no **edge inversions** if whenever $g(e) = e$, g fixes both vertices of e.

Remark 34.2. If $G \curvearrowright T$ is arbitrary, and T' is the result of subdividing all edges, then $G \curvearrowright T'$ has no edge inversions. So if you have edge inversions, it is trivial to get rid of them by adding vertices.

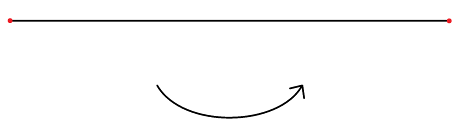

Figure 107. An edge inversion

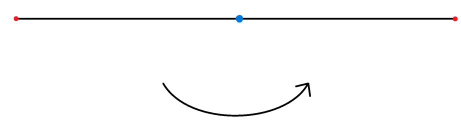

Figure 108. No edge inversions after subdividing

Convention: Unless otherwise specified, all our actions will be without inversions. (This implies $T \to T/G$ is a graph morphism.)

Lemma 34.3. *$G \curvearrowright T$ is minimal if and only if $core(T/G) = T/G$.*

Proof. If $core(T/G) \neq T/G$, its preimage is an invariant subtree T'.
Conversely, If T has an invariant subtree T',

$$T'/G \subsetneq T/G \qquad \pi_1(T'/G) = G = \pi_1(T/G),$$

So $core(T/G) \neq T/G$ by Lemma 7.3. □

The picture of general action $G \curvearrowright T$ is as follows:

- preimage of $core(T/G) \neq T/G$, where action is minimal;
- extra hanging trees that are permuted. (See Figure 109.)

Definition 34.4. A **free splitting** is a minimal action of F_n on a tree with trivial edge stabilizers.
(It is a **k-edges free splitting** if there are k orbits of edge.)

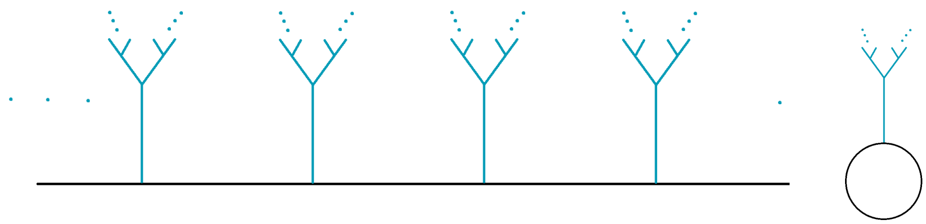

FIGURE 109. An example of T and T/G when $G = \mathbb{Z}$. The generator of $\mathbb{Z}$ translates to the right.

Definition 34.5. Two splittings are conjugate if there is an equivariant homeomorphism between them. (We think of them as "the same".)

Definition 34.6. A free splitting T_1 **collapses to** a free splitting T_2 if there is an invariant collection of edges of T_1, s.t. collapsing them gives (something conjugate to) T_2.

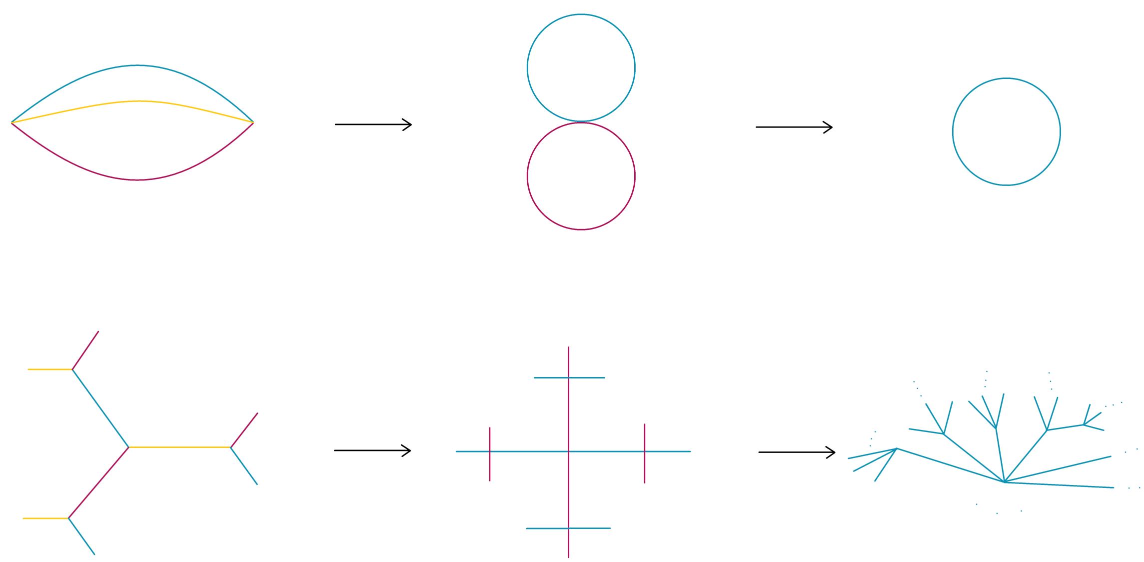

FIGURE 110. Collapsings of free splittings

Remark 34.7. Collapsing a proper invariant subset of edges always gives a new free splitting (still minimal, edge stabilizer still trivial).

Example 34.8. Let Σ be a surface with one boundary circle. Consider a collection of k disjoint arcs that go from the boundary to itself. We assume they are not isotopic to each other and cannot be isotoped rel boundary into $\partial\Sigma$.

Lift to $(\widetilde{\Sigma - \partial\Sigma}) = \mathbb{D}$, get an associated k-edge splitting of $\pi_1(\Sigma) = F_n$, with a vertex for each component of $\mathbb{D} - preimage\ of\ \bigcup\ arcs$.

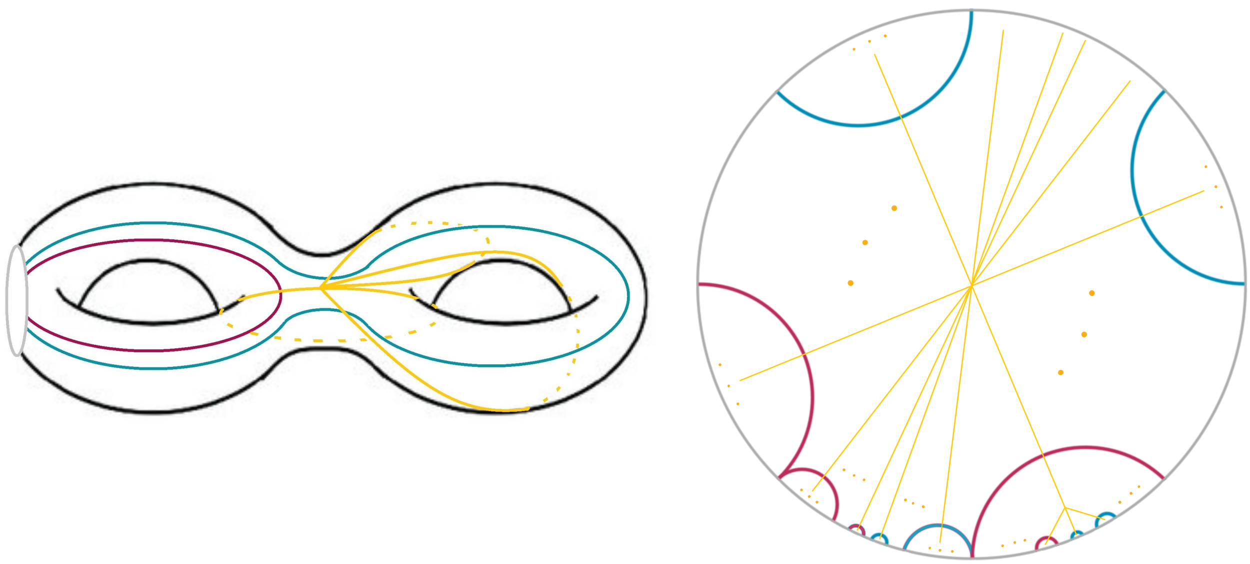

FIGURE 111. Σ with $\partial\Sigma$ a circle and universal cover of $\Sigma - \partial\Sigma$. (The lifted purple loops should not actually touch at infinity.)

We'll now take a slight digression, to address the question: What algebraic information is encoded in a free splitting? The answer, given by Bass-Serre Theory, is that it produces a collection of presentations of the group. We'll only be able to give one result from this theory, but you can read more in sources such as [Ser03, SW79]. We'll roughly follow part of [DD89], because its presentation isolates and gives a short proof of one of the main results.

Definition 34.9. If $v \in V(T)$, let $G_v = \{g \in G | gv = v\}$, similarly for $e \in E(T)$.

Let S be the lift to T of a maximal tree of T/G.

Let B be the subtree of T defined by S union, for each edge e of $G/T -$(image of S), a lift of e adjacent to S. See Figure 112.

Remark 34.10. The letters "S" and "B" are chosen to indicate that these are small and big "almost" fundamental domains. For example, S contains exactly one vertex from each orbit of vertices, and B contains exactly one edge from each orbit of edges. (Often when one talks about fundamental domains one ask that the group element that moves a point into the fundamental domain be typically unique. This won't hold in our situation because the vertex and edges stabilizers can be non-trivial.)

Remark 34.11. Note that the construction of either domain is not canonical, as we have made many choices. We are trying to find a group presentation, and group presentations

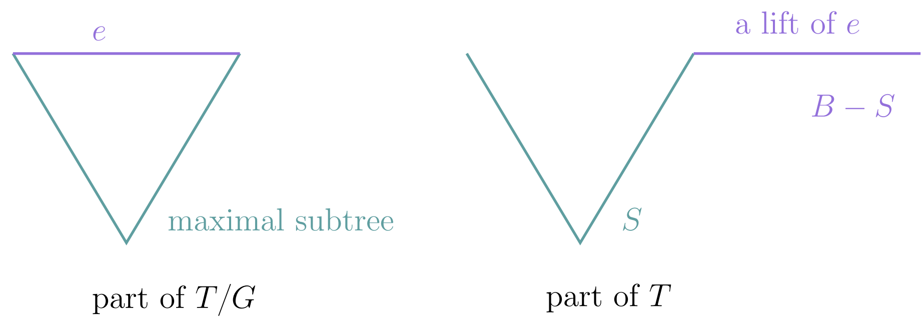

FIGURE 112. Small and big fundamental domain

are typically very non-unique and non-canonical, so it's expected that we have to make a lot of choices.

Theorem 34.12 (Part of Bass-Serre Theory)**.** *Say $G \curvearrowright T$ as above. For each $e \in B - S$, pick $s_e \in G$ that takes the other vertex of e into S. If $e \in S$, set $s_e = \mathrm{id}$.*

Then if P is the group

$$\left\langle \bigcup_{v \in V(S)} G_v \bigcup \left\{ t_e : e \in E(B) \right\} \middle| \text{ all relations from each } G_v; \right.$$

$$\left. \forall e \in E(B),\ g \in G_e,\ t_e g {t_e}^{-1} = s_e g {s_e}^{-1};\ t_e = 1 \text{ if } e \in E(S) \right\rangle$$

the natural map $P \to G$ is an isomorphism.

The unions describing the generating set are formal disjoint unions; it is somewhat of an abuse of notation to say that the generators come from G_v, and it would be better to instead say we are taking one new, formal symbol for each element of each G_v. If some element of G lies in two different $G_v, v \in V(S)$, it formally gives two distinct generators of P. These two distinct generators end up being equal in P, so the abuse of notation is only very mild.

35. Presentations for groups acting on trees (11/20, YL, ZH)

We begin with some remarks on Theorem 34.12.

Remark 35.1. Suppose $e \in E(S)$ is a directed edge from v to w. If $g \in G_e$, then we get two generators from g: $(g \in G_v)$ and $(g \in G_w)$. The relation in this case is

$$t_e(g \in G_v)t_e^{-1} = s_e(g \in G_w)s_e^{-1}$$

but since $t_e = id, s_e = id$, we have $(g \in G_v) = (g \in G_w)$ in P.

Exercise 59. *Use the previous remark to fully justify the claim from the end of last lecture: If $g \in G$ is contained in G_v and in G_w, with v, w in $V(S)$, then the two generators $(g \in G_v)$ and $(g \in G_w)$ are equal in P. (See Hint A.27.)*

Remark 35.2. If $e \in E(B) - E(S)$ is an edge from $v \in V(S)$ to $w \in V(B) - V(S)$, the relation

$$t_e g {t_e}^{-1} = s_e g {s_e}^{-1}$$

of P requires some interpretation. Namely, here g denotes the generator that might be more formally denoted $(g \in G_v)$. The element s_e by definition takes w to some vertex $u \in V(S)$. Note that $s_e g s_e{}^{-1}$ fixes u. In the relation, $s_e g s_e{}^{-1}$ denotes the generator that might be more formally denoted $(s_e g s_e{}^{-1} \in G_u)$.

Remark 35.3. The map $P \to G$ is defined by sending $(g \in G_v)$ to g and t_e to s_e.

Remark 35.4. Fix a presentation for each G_v. The theorem is equivalent to G having a presentation consisting of the union of the generating sets of each $G_v, v \in V(S)$, together with the indicated new generators t_e, and all the relations from each of the presentations of each G_v together with the indicated new relations involving the t_e. So the theorem says what *new* generators and relations are needed when passing from the collection of G_v to the whole group G.

Example 35.5. If T/G is an edge with two vertices, we have that $S = B$ consists of a single edge e with two vertices v and w, and the theorem tells us $P = G_v *_{G_e} G_w$.

Example 35.6. If T/G is a self loop at one vertex, we have S is just one vertex v and B is an edge e with two vertices v and w. In this situation, we can drop some subscripts, and use t, s instead of t_e, s_e. Define

$$\begin{aligned} \phi : (G_e \subset G_v) &\to G_v \\ h &\mapsto shs^{-1}. \end{aligned}$$

Then $G = \langle G_v, t \,|\, \forall h \in G_e,\ tht^{-1} = \phi(h) \rangle$. This description of G is exactly the definition of the HNN extension $G_v *_\phi$.

Example 35.7. If all edge stabilizers are trivial, then

$$G = \mathop{\ast}_{v \in V(B)} G_v \mathop{\ast}_{e \in E(B \setminus S)} \langle t_e \rangle.$$

In particular, note each G_v is a free factor of G.

Remark 35.8. What is the difference between the HNN extension $F_n = F_{n-1} *_\phi$, where $\phi : \{id\} \to \{id\}$ is the identity, and the free product $F_n = F_{n-1} * \langle t \rangle$? In a sense nothing, but in practice usually if you get $F_{n-1} * \langle t \rangle$, then $\langle t \rangle$ is canonical, and if you get $F_{n-1} *_\phi$, then the last generator comes from a non-canonical choice.

Example 35.9. If α is an arc on a surface with one boundary component, we get an associated free splitting of $\pi_1 = F_n$, which has quotient in the two cases below:

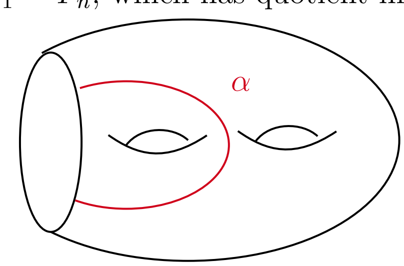

quotient is 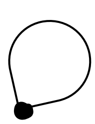 when α is non-separating

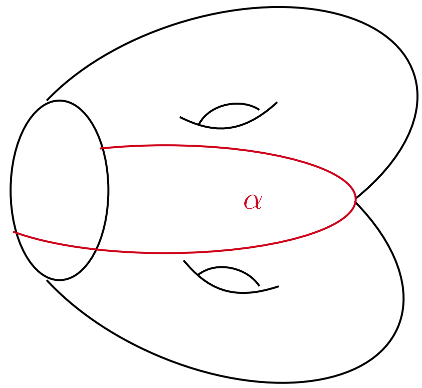

quotient is 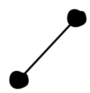 when α is separating

Outline of Proof for Theorem 34.12. We proceed in the case where edge stabilizers are trivial, since this is sufficient for our purposes and simplifies the presentation. (We only use this extra assumption in the last lemma.)

Lemma 35.10. *We have* $T = \bigcup_{g \in P} gB$ *(and* P *acts on* T *via the map* $P \to G$*).*

Proof. We will show that $\bigcup_{g\in P} gB$ has all edges adjacent to B. The result will then follow by connectivity, since $\bigcup_{g\in P} gB$ has every adjacent edge to itself.

Let f be an edge adjacent to B; it has one vertex p in $V(B)$, and the other vertex q not in $V(B)$. Our goal is to show f is in $\bigcup_{g\in P} gB$. By definition of B, there is some $h \in G$ such that $hf \in B$.

First assume $p \in V(S)$ and proceed as follows.

- If $hp \in V(S)$, then $hp = p$, by definition of S. In that case $h \in G_p$, so the fact that $f \in h^{-1}B$ proves the result, keeping in mind that $G_p \subset P$.
- If $hp \notin V(S)$, then $hf \in E(B) - E(S)$, and $s_{hf}h(p) = p$. So

  $$f = (s_{hf}h)^{-1} t_{hf}(hf),$$

  which proves the result since $hf \in E(B), t_{hf} \in P$, and $(s_{hf}h)^{-1} \in G_p \subset P$.

Next assume that $p \notin V(S)$. In that case, p is an endpoint of an edge $e \in E(B) - E(S)$ and we have $s_e(p) \in V(S)$. Since $s_e(f)$ is in $\bigcup_{g\in P} gB$ by the previous case, applying t_e^{-1} to $s_e(f)$ gives the result. □

Lemma 35.11. *The map* $P \to G$ *is surjective.*

Proof. We will denote this map by I and write a bit more formally than usual. Pick $g \in G$; our goal is to show g is in the image of I.

Pick $v \in V(S)$. By Lemma 35.10, there exists a $p \in P$ such that $gv \in I(p)B$, and so $(I(p)^{-1}g) \cdot v \in V(B)$.

If $(I(p)^{-1}g)\cdot v \in V(S)$, then by definition of S, $I(p)^{-1}gv = v$. Thus, $I(p)^{-1}g \in G_v \subset P$, and

$$g = I(p \cdot I(p)^{-1}g)$$

is in the image as desired, since p maps to $I(p)$ via I and $I(p)^{-1}g$ viewed as a generator of P maps to $I(p)^{-1}g$ in G.

If $I(p)^{-1}gv \in V(B)\backslash V(S)$, let e be the edge in B, with $I(p)^{-1}gv$ as an endpoint. Thus $s_e(I(p)^{-1}gv) \in V(S)$, and hence the definition of S gives $s_eI(p)^{-1}gv = v$. Hence $s_eI(p)^{-1}g \in G_v \subset P$, and

$$g = I(p \cdot t_e^{-1} \cdot s_eI(p)^{-1}g)$$

is in the image as desired. □

Lemma 35.12. *Let G be a group, $g_1, \dots, g_n \in G$ and let*

$$\widehat{g}_n = (g_1 \cdots g_{n-1})g_n(g_1 \cdots g_{n-1})^{-1}.$$

Then $\widehat{g}_n\widehat{g}_{n-1} \cdots \widehat{g}_1 = g_1g_2 \cdots g_n$.

Remark 35.13. If B is a fundamental domain for a group action, then the translates gB tile the space. If g_iB is adjacent to B, we can think of g_i as a "local move" at B, in that it moves B to an adjacent "copy of B" in the tiling. See Figure 113.

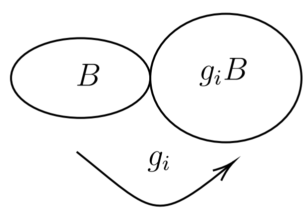

FIGURE 113

If g_n is a local move at B, then $\hat{g}_n$ is a local move at $g_1...g_{n-1}B$; see Figure 114. The conclusion of this observation and Lemma 35.12, as illustrated in Figure 115, is that if all the g_i are local moves at B, then by "adding hats and reversing the order of multiplication" we can get $g_1 \cdots g_nB$ via a chain of local loves starting at B.

Lemma 35.14. *$P \to G$ is injective.*

Proof. Recall that we are assuming that all edge stabilizers are trivial. This implies that for each edge in T there is a unique group element moving it into B. In particular, we have that if $g \in G$ has $g(B) = B$, then g is the identity.

Assume in order to find a contradiction that the lemma is false. Let $g_1 \cdots g_n$ be a minimal length word in P that is non-trivial in P but trivial in G. Here all the g_i are generators of P, and minimality in particular implies that

- none of the g_i are the identity element of any $G_v, v \in V(S)$,
- no consecutive pair g_i, g_{i+1} lies in the same $G_v, v \in S$, and
- none of the g_i are $t_e^{\pm}, e \in E(S)$.

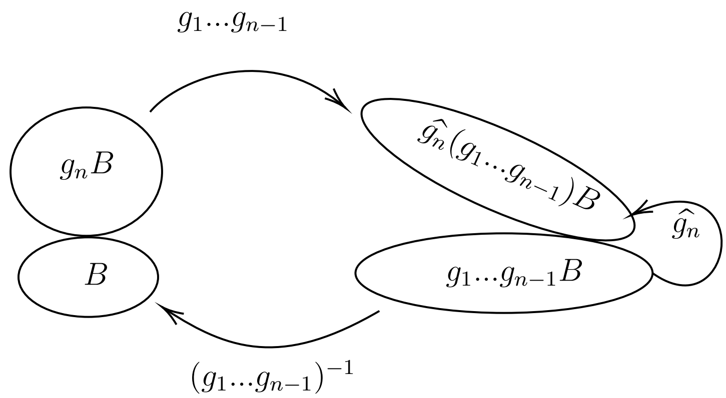

FIGURE 114

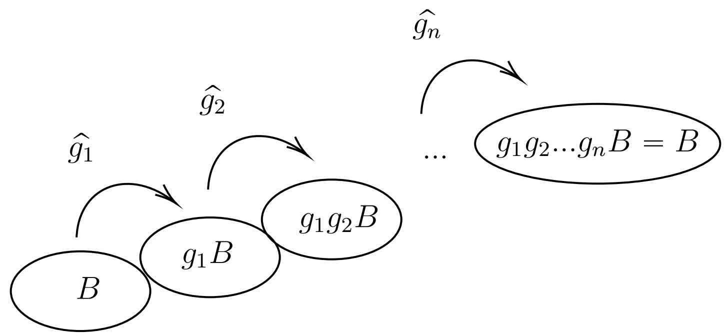

FIGURE 115

Using the action on T, one can check that all non-trivial generators of P are non-trivial in G, so we have $n \geqslant 2$.

Since T is a tree, $B = g_1 \cdots g_n B$, the path of translates of B in Figure 115, must backtrack. Thus, there exists an i such that

$$g_1 \cdots g_{i-1} B = g_1 \cdots g_{i+1} B.$$

Thus, $g_i g_{i+1} B = B$. Hence, by the first paragraph of this proof, $g_i g_{i+1}$ is the identity. Since we chose our word to be as small as possible, this implies $n = 2$.

We now make the following observations:

- Since edge stabilizers are trivial, a non-identity element of G either has fixed point set which is either a single vertex or is empty.
- Thus if $g_1 \in G_v, g_2 \in G_w$, with v, w distinct elements of $V(S)$, then

$$g_1 g_2(w) = g_1(w) \neq w.$$

  In particular, $g_1 g_2$ is not the identity.
- If $e \in E(B) - E(S)$, then $s_e(S)$ is disjoint from S. Similarly $s_e^{-1}(S)$ is disjoint from S. See Figure 116.

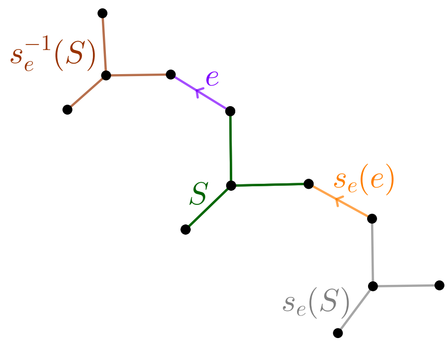

FIGURE 116

- Hence if $g \in G_v, v \in V(S)$, and $e \in E(B) - E(S)$, then $s_e^{\pm 1} g$ sends v outside of $V(S)$. (Keep in mind that $s_e^{\pm 1} g(v) = s_e^{\pm 1}(v)$, and apply the previous point.) So $s_e^{\pm 1} g$ is not the identity in G. The same of course applies for the inverse.
- A similar argument shows that if $e, f \in E(B) - E(S)$ are distinct, then non-trivial words of length two in s_e and s_f are not the identity.

It follows that $n = 2$ is not possible, giving a contradiction. □

Since we have shown $P \to G$ is surjective and injective, it must be an isomorphism. We never used the relation

$$\forall e \in \ E(B),\ g \in \ G_e,\ t_e g {t_e}^{-1} = s_e g {s_e}^{-1}$$

because this relation does not appear under our simplifying assumption that all G_e are trivial.

36. THE FREE SPLITTING COMPLEX (11/27, SK, NL)

The following is a corollary of the theorem proved in the last class:

Corollary 36.1. $\{[G_v] : v \in V(T)\}$ *is a set of vertices of bounded diameter in* FF_n.

Definition 36.2. The **free splitting complex** FS_n is the simplicial complex with a vertex for each (conjugacy class of) free splitting and a k-simplex for each tuple of splittings

$$F_n \curvearrowright T_0,\ F_n \curvearrowright T_1,\ \ldots,\ F_n \curvearrowright T_k,$$

where T_{i+1} collapses to T_i.

Remark 36.3. Each point $\Gamma \in CV_n$ with k-edges determines 1 k-edge free splitting, namely $F_n = \pi_1(\Gamma) \curvearrowright \tilde{\Gamma}$, and $\binom{k}{i}$ i-edge free splitting for $1 \leqslant i \leqslant k$ (collapse all but i orbits of edges in $\tilde{\Gamma}$). See Figure 117.

Exercise 60. *These are all non-conjugate. (See Hint A.28.)*

Example 36.4. Figure 117 is an example for $k = 3$.

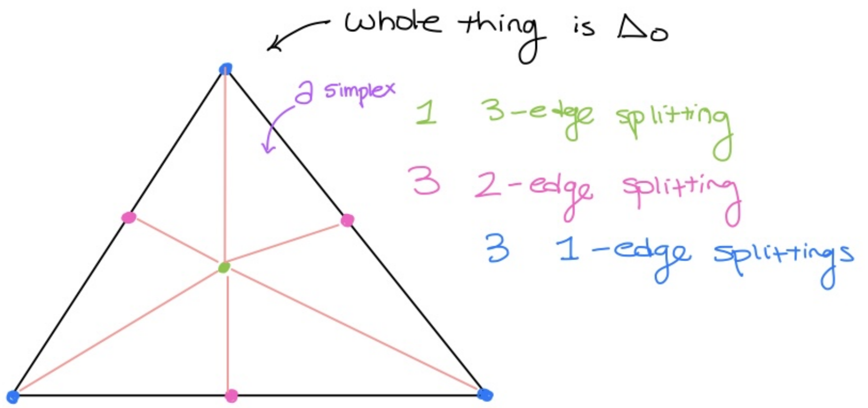

FIGURE 117. Δ_0, with how many k-edge splittings it has.

Thus Γ determines a subset of CV_n equal to the barycentric subdivision of a $(k-1)$-simplex Δ_0 whose vertices can be labeled by edges of Γ (collapse all but the preimage of that edge).

Thus Δ_0 is parametrized by $\{(\ell_e)_{e \in E(\Gamma)} : \sum \ell_e = 1, \ell_e \geqslant 0\}$. Mapping Γ to $(\text{length}(e))_{e \in E(\Gamma)}$ gives an inclusion $CV_n \hookrightarrow FS_n$. CV_n is a union of open simplices of FS_n.

Exercise 61. *Show that CV_n is dense in FS_n. (See Hint A.29.)*

Remark 36.5. To get FS_n, barycentrically subdivide CV_n and add in any missing faces. Actually, we have presented a slightly non-standard variant of the free splitting complex; usually one defines it so that barycentric subdivision is not required, as we later discuss in Definition 38.7. The hard part here is to understand the standard version of FS_n: One needs to first show that for every k-tuple of 1-edge free splittings, any two of which are collapses of a 2-edge free splitting, there is a unique k-edge free splitting that collapses to all of them. This additionally shows that the standard version of FS_n is a flag complex. One approach to that is via 3-manifolds; another is sketched in [Besc, Hint to Exercise 4].

Corollary 36.6. *FS_n is connected.*

Remark 36.7. On FS_n, we use the simplicial metric, i.e. where each edge has size 1 (or use graph metric on $FS_n^{(1)}$). This is very different from the asymmetric metric on CV_n.

Exercise 62. *Show that the map from cv_2 to the Farey complex extends to an isomorphism from the closure of cv_2 in FS_2 to the Farey complex. Show that this closure is dense, and that the Farey complex is quasi-isometric to FS_2. (This shows $FS_2 \to FF_2$ is a quasi-isometry. This is not true in higher rank. See Hint A.30.)*

It is combinatorially easier to use

$$\rho : CV_n \to FS_n, \qquad \Gamma \mapsto F_n \curvearrowright \tilde{\Gamma},$$

so we will use this. This map ρ is not continuous, but $d(\Gamma, \rho(\Gamma)) \leqslant 1$ in the simplicial matrix; see Figure 118.

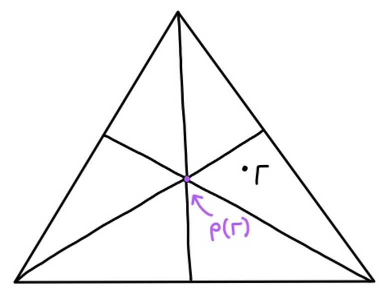

FIGURE 118. In this case, $\rho(\Gamma)$ is the barycentric center.

Proposition 36.8. *ρ is coarsely Lipschitz.*

Proof. Consider $\Gamma, \Gamma' \in CV_n$. There is an optimal map

$$\phi : \Gamma \to \Gamma'$$

of slope $\sigma = e^{d(\Gamma,\Gamma')}$. As in Corollary 21.7, we can change edge lengths so the tension graph Δ becomes all of Γ. This doesn't change $\rho(\Gamma)$ and moves along the start of a geodesic from Γ to Γ', so without loss of generality we assume $\Delta = \Gamma$.

Scale Γ by σ, so $vol(\Gamma) = \sigma$ and $\Gamma \to \Gamma'$ is a local isometry on edges (slope 1). Subdivide so $\Gamma \to \Gamma'$ is a morphism. Continuously fold (1 edge at a time, unit speed, arbitrary order) to get a path $\Gamma_s, s \in [0, t_0]$ with

(1) $\Gamma_0 = \Gamma, \Gamma_{t_0} = \Gamma'$,
(2) $vol(\Gamma_s) = vol(\Gamma) - s = \sigma - s$,
(3) $t_0 = \sigma - 1$.

We claim that for $s < \sigma/2$, there exists a point Γ_s on the interior of an edge with only one preimage on $\Gamma_0 = \Gamma$. We now prove this claim. Subdivide so that $\Gamma_0 \to \Gamma_s$ is a morphism. If each edge had $\geqslant 2$ preimages, we would have $vol(\Gamma_s) \leqslant \frac{vol(\Gamma_0)}{2}$, which proves the claim. See the diagram below to see this exemplified.

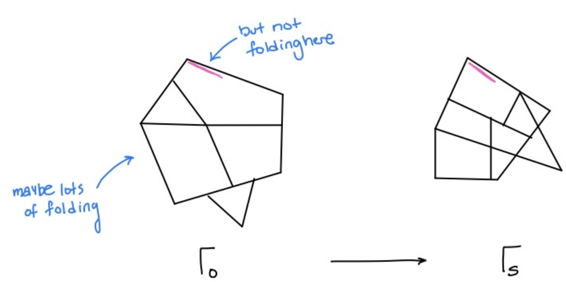

FIGURE 119. An example of a map from $\Gamma_0 \to \Gamma_s$

Note that $\{\rho(\Gamma_s) : 0 \leqslant s \leqslant \frac{\sigma}{2} - \epsilon\}$ has diameter $\leqslant 2$ since it all collapses to a 1 edge free splitting given by a segment of Γ_s with one preimage in Γ_0.

Iterating gives a linear bound in $\log_2 \sigma = \frac{d(\Gamma,\Gamma')}{\log 2}$. □

Note that the map $CV_n \to 2^{FF_n^{(0)}}$ can be extended to $FS_n \to 2^{FF_n^{(0)}}$ via $F_n \curvearrowright T \to$ {all vertex stabilizers of all collapses}. This is a finite set.

Exercise 63. *This agrees with the definition for CV_n (using $CV_n \subset FS_n$).*

Exercise 64. *The image of each point is a bounded diameter subset of FF_n and $FS_n \to FF_n$ is coarsely Lipschitz. (See Hint A.31.)*

The main result on FS_n is:

Theorem 36.9. *[Handel-Mosher] FS_n is Gromov hyperbolic.*

This was proven in [HM13], but a simplification is available in [BF14b]. It is known that hyperbolicity of FS_n implies hyperbolicity of FF_n [KR14].

Exercise 65. *This exercise builds on Exercises 58 and 61. For any 1-edge free splitting, consider the subset of CV_n where the universal cover collapses to this 1-edge free splitting. (On this subset, there is a well defined edge that doesn't disappear, and in the universal cover collapsing onto this edge gives the 1-edge free splitting.) Show that FS_n is quasi-isometric to the electrification of CV_n along the collection of these subsets corresponding to all possible 1-edge free splittings. (See Hint A.32.)*

37. Doubled handlebodies, Dehn twists (11/29, KS, LS)

One reference for some of the material in this section and the next is [BBP23].

Definition 37.1. The *handlebody* H_n is the 3-manifold with boundary obtained as the inside of the standard embedding of genus n surface $\Sigma_n \hookrightarrow \mathbb{R}^3$.

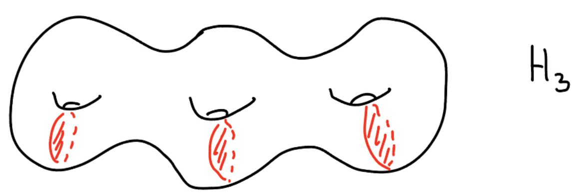

Figure 120. handlebody H_3.

Definition 37.2. The *doubled handlebody* M_n is two copies of H_n glued by the identity map on their boundary.

Example 37.3. $M_1 = S^1 \times S^2$ (see Figure 121).

Exercise 66. *M_n is homeomorphic to a connected sum of n copies of $S^1 \times S^2$.*

An explicit bijection $\pi_1(M_n) \to F_n$ arises as follows: Label n spheres S^2 by $a_1, \ldots, a_n$ as in Figure 123 and choose a positive direction at each. For each loop γ in $\pi_1(M_n)$ we determine its image in F_n just by reading of which spheres it intersects, taking

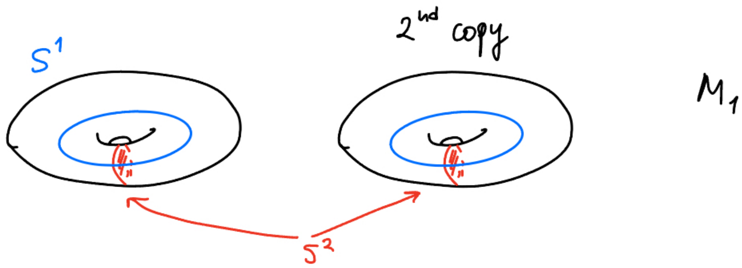

FIGURE 121. Doubled handlebody M_1.

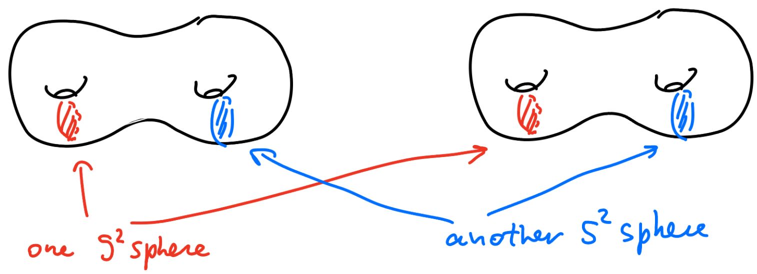

FIGURE 122. Doubled handlebody M_2.

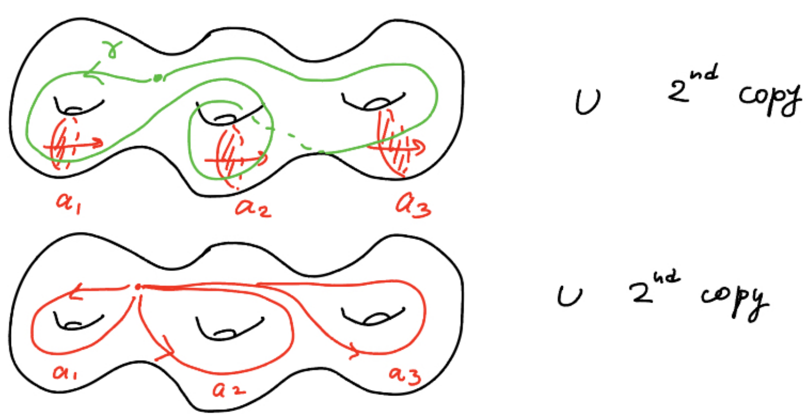

FIGURE 123

an inverse if the sphere is crossed in the negative direction. For example, for the γ illustrated, we get $\gamma = a_1 a_2^{-1} a_3$.

If we cut a handlebody along n discs as shown in the top of Figure 124, what we get is homeomorphic to the 3-ball. There are $2n$ discs on the boundary of this three ball, which arise from the n discs of the handlebody. If we glue two 3-balls along their boundary we get S^3. When we glue along just the part of the boundary not contained in the $2n$ discs, what we get is S^3 minus $2n$ 3-balls.

Remark 37.4. Thus we can draw M_n as $S^3 - 2n$ 3-dimensional balls with boundary spheres S^2 identified in pairs (see Figure 124). Note that the complement of the n spheres is simply connected.

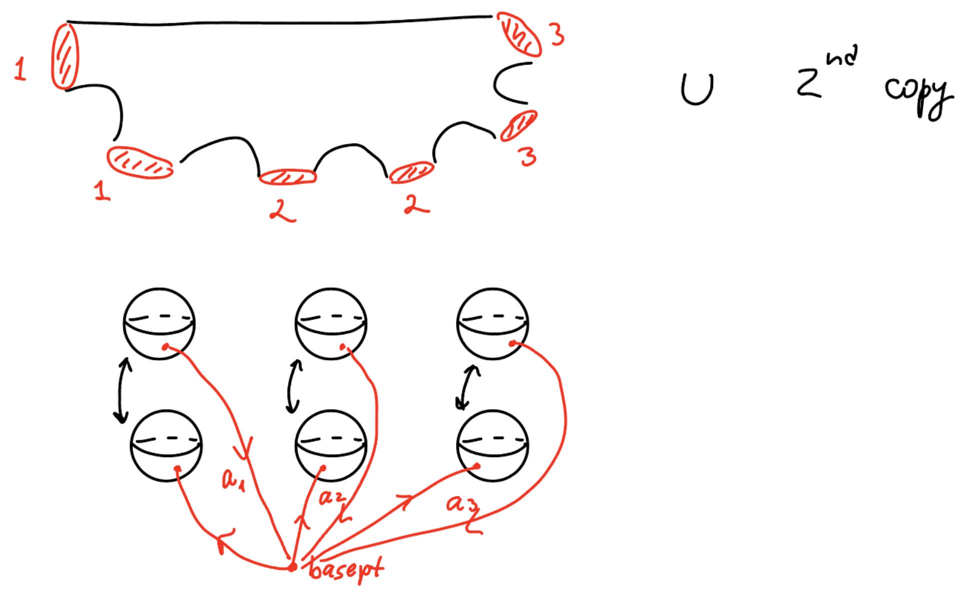

FIGURE 124

Example 37.5. The green region A corresponds to a component of the complement of 3 spheres (see Figure 125). We can set the middle red sphere centered at infinity, so A will be identified with interior of the middle sphere cutting out blue and orange spheres. Note that connected sum of blue and orange spheres along the tube is topologically equal to the red sphere.

Definition 37.6. $\mathrm{MCG}(M_n) = \mathrm{Diffeo}^+(M_n)/\mathrm{Diffeo}_0^+(M_n)$, where $+$ means orientation preserving and 0 means diffeomorphisms isotopic to identity.

Remark 37.7. The map $\mathrm{Diffeo}^+(M_n)/\mathrm{Diffeo}_0^+(M_n) \to \mathrm{Homeo}^+(M_n)/\mathrm{Homeo}_0^+(M_n)$ is an isomorphism [Sta].

In the following, we always assume that a sphere S does not bound a ball.

Definition 37.8. Let S be a sphere in M_n with a regular neighborhood $S \times [0,1]$. Fix an identification

$$S = \{(x_1, x_2, x_3) \in \mathbb{R}^3 : x_1^2 + x_2^2 + x_3^2 = 1\}.$$

Let $g_t, t \in [0,1]$ be a loop generating $\pi_1(SO(3)) = \mathbb{Z}/2\mathbb{Z}$. To be specific, g_t is a rotation by $2\pi t$ about z-axis. Let $T_S : M_n \to M_n$ be the diffeomorphism that is the identity off $S \times [0,1]$ and on $S \times [0,1]$ is defined by $(p,t) \mapsto (g_t p, t)$. This diffeomorphism T_S is called *Dehn twist*.

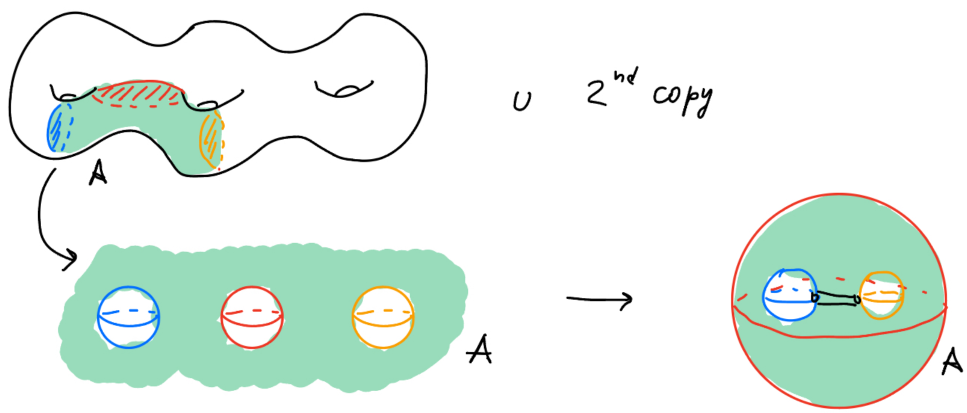

FIGURE 125

Exercise 67. $[T_S]^2 = \text{id}$ *in* $\text{MCG}(M_n)$.

Fact 37.9. $[T_S] \in \text{MCG}(M_n)$ *depends only on the homotopy class of* S.

Fact 37.10. *If* $i_1, i_2 : S \hookrightarrow M_n$ *are embedded spheres and* i_1 *is homotopic to* i_2*, then they are (ambient) isotopic.*

Example 37.11. Imagine a manifold (e.g. sphere) N intersects S in a circle (see Figure 126). In the picture, the inside sphere of the regular neighborhood is $S \times \{0\}$, and the outside outside sphere is $S \times \{1\}$. Note that N is homotopic to $T_S(N)$. In fact, T_S acts trivially on the homotopy (isotopy) class of all embedded 1 and 2-manifolds.

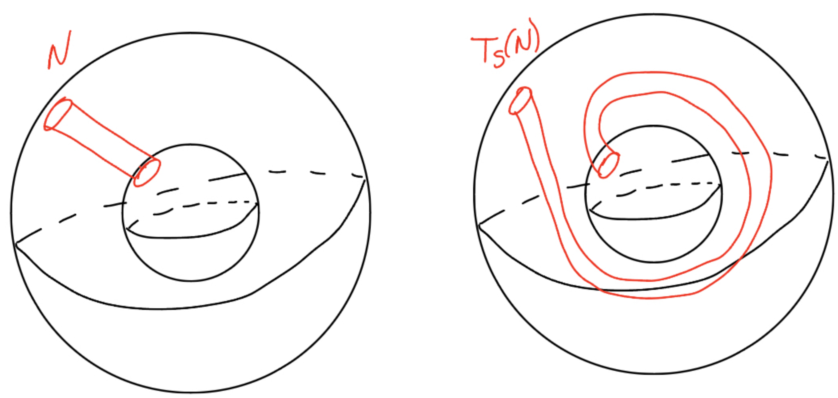

FIGURE 126

Remark 37.12. It is not easy to show that $[T_S] \neq \text{id}$.

Remark 37.13. $[T_{S'}]^{-1}[T_S][T_{S'}] = [T_{T_{S'}(S)}] = [T_S]$, since $T_{S'}(S)$ is homotopic to S.

Definition 37.14. Let DT be the subgroup of $\text{MCG}(M_n)$ generated by all $[T_S]$.

Corollary 37.15. DT *is abelian.*

38. Laudenbach's Theorem, sphere systems (12/01, LS, UP)

Theorem 38.1 (Laudenbach)**.** *The morphism*

$$MCG(M_n) \to Out(F_n)$$

is surjective with kernel DT*, which is the subgroup generated by all Dehn-twists. Furthermore*

$$\mathrm{DT} \cong (\mathbb{Z}/2)^n$$

and is generated by the twists in n disjoint "standard" spheres.

Remark 38.2. If M_n happened to be a $K(\pi, 1)$, then every homotopy equivalence would be determined by its action on π_1. Therefore it makes philosophical sense that the kernel above is related to 2-spheres in M_n, as these are the obstructions to M_n being a $K(\pi, 1)$.

In fact there is a recent result that improves Laudenbach's Theorem.

Theorem 38.3 (Brendle-Broaddus-Putman)**.** *The short exact sequence*

$$1 \to \mathrm{DT} \to MCG(M_n) \to Out(F_n) \to 1$$

splits.

The upshot of these theorems is that $MCG(M_n)$ is "really close" to $Out(F_n)$: $Out(F_n)$ is both a subgroup and a quotient of $MCG(M_n)$, and the "difference" between these two groups is a tiny group DT that acts trivially on just about everything.

The goal of today's lecture is to explain the proof of the surjectivity of $MCG(M_n) \to Out(F_n)$.

Proposition 38.4. *The morphism $MCG(M_n) \to Out(F_n)$ is surjective.*

The proof follows a standard strategy: we prove that a generating set of $Out(F_n)$ is contained in the image. We choose to show that all signed permutations are in the image and so is one Whitehead morphism, at which point we can invoke Nielsen's Theorem (Theorem 8.2).

Lemma 38.5. *All signed permutations are contained in the image of $MCG(M_n) \to Out(F_n)$.*

Proof. We starting by indicating a diffeomorphism whose mapping class maps to the signed permutation that inverts a single element. As in Figures 127 and 128, we define this diffeomorphism by twisting a single handle (and identically twisting its mirror copy). This diffeomorphism induces an automorphism of π_1 which sends $a_1 \mapsto a_1^{-1}$.

Next we need to give a diffeomorphism which swaps a_i, a_{i+1}, up to inverses. We can apply a similar idea, this time rotating the handles corresponding to a_i and a_{i+1}, so they get swapped, as indicated in Figure 129. □

It is now enough to show that one specific additional automorphism is in the image.

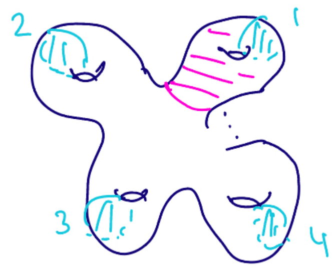

FIGURE 127. Drawing of M_n

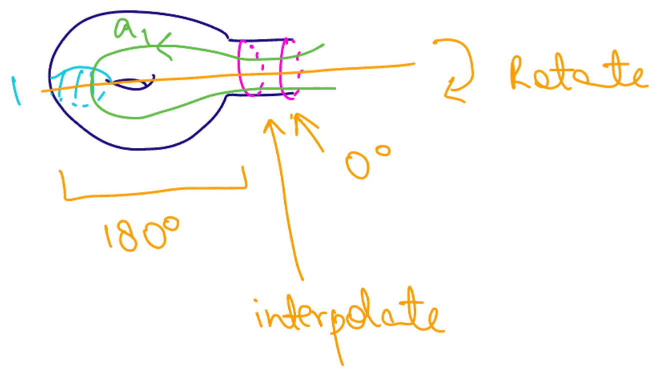

FIGURE 128. Rotation inducing $a_i \mapsto a_i^{-1}$

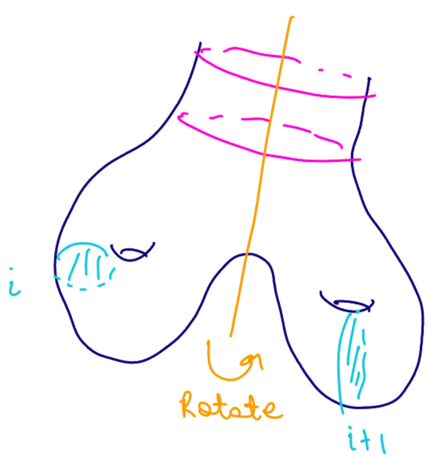

FIGURE 129. Rotation swapping a_i, a_{i+1}

Lemma 38.6. *The outer automorphism*

$$\begin{array}{rcl} a_1 & \mapsto & a_1 a_2 \\ a_2 & \mapsto & a_2^{-1} \\ a_i & \mapsto & a_i, \quad i > 2 \end{array}$$

is in the image.

Proof. To construct this automorphism, we can focus on two handles of M_n, see Figure 130. Then we can cut along the pink sphere and the blue sphere labeled "1", which yields two pairs of pants. These pants fit together as in Figure 131. We can then

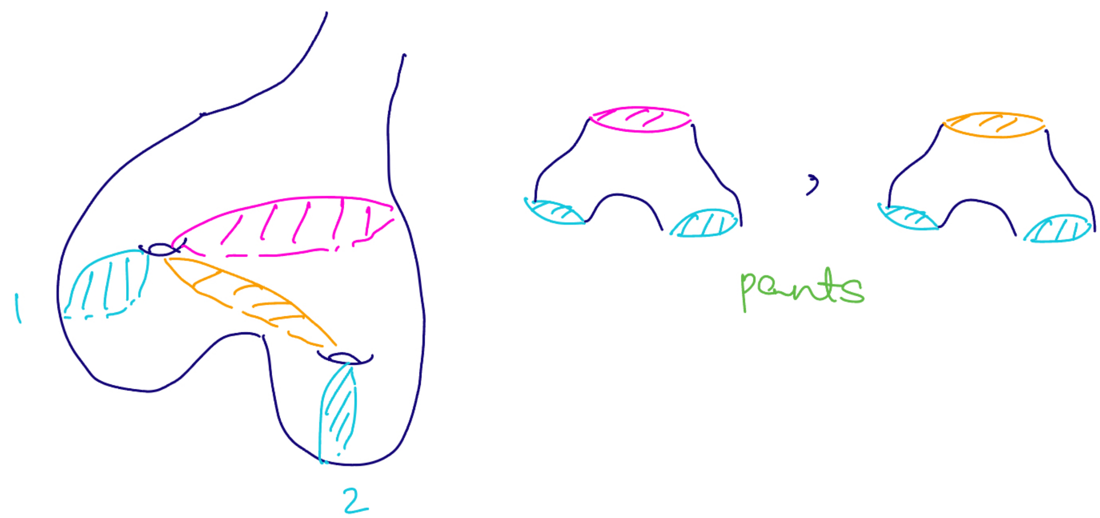

FIGURE 130. One of the boundaries on the middle part of the figure should be orange.

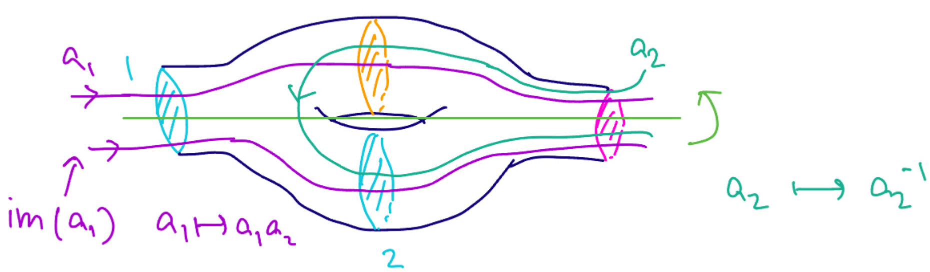

FIGURE 131

rotate by 180 degrees in the middle of this drawing, staying fixed on the boundary (interpolating near the boundary to make this a diffeomorphism). By tracing the arrows and seeing which spheres we intersect, one can check that this induces the desired Whitehead automorphism. In Figure 131, the loop corresponding to a_1, and its image under this map are denoted by purple, and the loop a_2 in green. □

Going back to Proposition 38.4, the statement now follows using Nielsen's Theorem (Theorem 8.2), since $Out(F_n)$ is generated by elements contained in the image by Lemma 38.5 and Lemma 38.6 respectively.

We now give a slightly different description of the free splitting complex.

Definition 38.7. FS_n^{std} is the delta-complex with a k-simplex with a k-face for every $(k+1)$edge free splitting, with glueing given by collapse maps.

Remark 38.8. FS_n as we first defined it is the barycentric subdivision of FS_n^{std}. As discussed briefly in Remark 36.5, it turns out that FS_n^{std} is a flag simplicial complex.

Definition 38.9. The sphere complex $\mathcal{S}_n$ is a simplicial complex with vertices for each isotopy class of spheres (not bordering a ball) in M_n where $[S_0], \ldots, [S_k]$ span a k-simplex if they can be realized disjointly.

Actually a priori this is a delta complex, but it turns out to be a simplicial complex.

There exists a morphism
$$\mathcal{S}_n \to FS_n^{std}$$
which is induced by sending $\{[S_0], \ldots, [S_n]\}$ to the $\pi_1(M_n) = F_n$ action on the dual tree of the preimage of the spheres in the universal cover of M_n.

Proposition 38.10. *This is an isomorphism.*

A proof sketch can be found in [AS11, Lemma 2]; note that the arxiv version has more discussion, but the published version has precise references that more clearly add up to a complete proof.

Remark 38.11. A few remarks:

(1) Sphere systems were the original perspective on the Whitehead algorithm.
(2) Sphere systems can be used to show that FS_n and CV_n are contractible.
(3) They give a useful perspective to show that FS_n is Gromov hyperbolic [HH17].

Remark 38.12. We end with a remark on why we used doubled handlebodies, instead of the handlebodies themselves. The mapping class group of a handlebody is sometimes called a handlebody group, and is a topic of study it its own right [Hen20]. It can be viewed as a subgroup of the mapping class group of its boundary. It does have a surjection to $Out(F_n)$, but unlike the case for doubled handlebodies, its kernel is huge and complicated. The kernel is even infinitely generated! Infinite generation isn't obvious, but it is obvious it has infinite kernel: The handlebody group, viewed as a subgroup of the mapping class group of its boundary, contains the Dehn twist about any simple closed curve in the boundary that is the boundary of a disc in the handlebody. These are all infinite order elements of the kernel of the map to $Out(F_n)$. The benefit of doubling the handlebody is now clear, since as soon as you double these all become order at most 2!

39. Survey of additional results (12/04, HT, SZ)

39.1. **Individual Automorphisms.**

- Each $\Phi \in \mathrm{Out}(F_n)$ is either exponentially growing (EG) or polynomially growing (PG). In either case, there is a "relative train track map" $\phi : \Gamma \to \Gamma$ realizing Φ. A relative train track map means that there is a filtration of invariant subgraphs
$$\Gamma^{(0)} \subset \cdots \subset \Gamma^{(k)} = \Gamma$$
such that the failure of ϕ on $\Gamma^{(i)}$ to be a train track map is entirely contained in $\Gamma^{(i-1)}$. The closure of $\Gamma^{(i+1)} - \Gamma^{(i)}$ is called a stratum, and each stratum can be either exponentially or polynomially growing.
- There is a theory of currents/laminations of F_n, which are useful for understanding the iterative dynamics of individual automorphisms.
- If $\Phi \in \mathrm{Aut}(F_n)$, consider the subgroup $\mathrm{Fix}(\Phi)$ of F_n containing all elements fixed by Φ. This subgroup is always finitely generated, and in fact always has rank at most n. This is called the Scott conjecture, and the proof used relative train track maps [BH92].
- Φ acts on FF_n loxodromically if and only if Φ is fully irreducible. A classification is known about when the action of Φ on FS_n is loxodromic; the answer is in terms of laminations [HM19].
- Much subsequent work relies on finding better and better relative train track maps.

39.2. **Mapping Torus.**

- If $\Phi \in \mathrm{Aut}(F_n)$, one can consider the 'mapping torus' group $F_n \rtimes_\Phi \mathbb{Z}$. This group is Gromov hyperbolic if and only if Φ is atoroidal [Bri00].
- One can ask when
$$F_n \rtimes_\Phi \mathbb{Z} \cong F_{n'} \rtimes_{\Phi'} \mathbb{Z}$$
Remarkably, this isomorphism can hold when $n \neq n'$.
- There are many open problems on mapping tori [Aim]. One can also ask when $F_n \rtimes_\Phi G$ is hyperbolic for some group G.

39.3. **General Properties.**

- A group is *linear* if it has a finite-dimensional faithful representation over some field. $\mathrm{Aut}(F_n)$ is not linear for $n \geqslant 3$ [FP92], and since $\mathrm{Aut}(F_{n-1})$ is a subgroup of $\mathrm{Out}(F_n)$, $\mathrm{Out}(F_n)$ is also not linear for $n \geqslant 4$.

  Exercise 68. *Show* $\mathrm{Aut}(F_{n-1})$ *is a subgroup of* $\mathrm{Out}(F_n)$*, by showing that the natural map* $\mathrm{Aut}(F_{n-1}) \to \mathrm{Aut}(F_n)$ *obtained by fixing the n-th generator stays injective after composing with* $\mathrm{Aut}(F_n) \to \mathrm{Out}(F_n)$*.*

  Clearly $\mathrm{Out}(F_2) = GL_2(\mathbb{Z})$ is linear. Whether $\mathrm{Out}(F_3)$ and $\mathrm{Aut}(F_2)$ are linear is unknown.
- Since there is a homomorphism from a finite-index subgroup of $\mathrm{Out}(F_3)$ to $GL_2(\mathbb{Z})$, $\mathrm{Out}(F_3)$ does not have property (T).

Remark 39.1. This homomorphism to $GL_2(\mathbb{Z})$ can be constructed as follows. Let R_3 be the rose, and consider a degree 2 cover $\Gamma \to R_3$. We can write

$$H_1(\Gamma) = H_1^+(\Gamma) \oplus H_1^-(\Gamma)$$

as the sum of 1 and -1 eigenspaces for the involution whose quotient is R_3. There is a finite index subgroup G of $\mathrm{Out}(F_3)$ that lifts to homotopy equivalences of Γ. The desired homomorphism can be obtained by considering the action of G on H_1^-.

However, $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$ have property (T) for $n \geqslant 4$ [KKN21, Nit]. Note that since $\mathrm{Out}(F_n)$ is a quotient of $\mathrm{Aut}(F_n)$, it inherits (T) from $\mathrm{Aut}(F_n)$ (a quotient of a group with (T) has (T)).

39.4. **Subgroups.**

- $\mathrm{Out}(F_n)$ satisfies the Tits alternative: every subgroup H of $\mathrm{Out}(F_n)$ either contains a sub-subgroup isomorphic to F_2 or is virtually abelian [BFH00].
- Every subgroup H of $\mathrm{Out}(F_n)$ either has a fully irreducible element, or virtually fixes some free factor [Hor16].

39.5. **Rigidity.**

- All automorphisms of $\mathrm{Aut}(F_n)$ and $\mathrm{Out}(F_n)$ are inner automorphisms for $n \geqslant 3$. In particular, $\mathrm{Out}(\mathrm{Aut}(F_n))$ and $\mathrm{Out}(\mathrm{Out}(F_n))$ are both trivial for $n \geqslant 3$ [BV00].
- $\mathrm{Aut}(FF_n) = \mathrm{Out}(F_n)$ [BB].
- Recalling that K_n is the spine of CV_n and considering its simplicial automorphisms, $\mathrm{Aut}(K_n) = \mathrm{Out}(F_n)$ [BV01].
- The isometry group of CV_n is $\mathrm{Out}(F_n)$ [FM12b].
- These are all instances of a metaconjecture that any sufficiently rich structure associated to $\mathrm{Out}(F_n)$ will have its group of automorphisms equal to $\mathrm{Out}(F_n)$. This is also studied in the context of mapping class groups, where it is called Ivanov's metaconjecture.

39.6. **Algebraic Topology.**

- FF_n is homotopic to a wedge of spheres [HV].
- FS_n is contractible [HV].
- $\mathrm{Out}(F_n)$ is $2n-5$-connected at ∞ and is a virtual duality group of dimension $2n-5$ [BF00].
- The rational Euler characteristic of $\mathrm{Out}(F_n)$ is negative, and grows super-exponentially in n [BV23].
- There is a linear function $l(i)$ so that

  $$H_i(\mathrm{Aut}(F_n)) \cong H_i(\mathrm{Out}(F_n))$$

  for $n > l(i)$. For this and the next results an expository reference is [Vog02].
- There is another linear function $l'(i)$ so that

  $$H_i(\mathrm{Aut}(F_n)) \cong H_i(\mathrm{Aut}(F_{n+1}))$$

  for $n > l'(i)$.

- These facts allow us to consider the stable homology of $\mathrm{Aut}(F_n)$. This stable homology is the same as for the symmetric group S_n, and in particular is all torsion.

39.7. **Boundaries.**

- Let $\mathcal{C}$ be the set of conj. classes of non-identity elements of F_n. Then there is a map $CV_n \to (0,\infty)^{\mathcal{C}}$ sending each Γ to the map $(\alpha \to l_\alpha(\Gamma))$. This map is injective, and is in fact injective to the projectivization $\mathbb{P}(0,\infty)^{\mathcal{C}}$.

  Exercise 69. *Show the map $CV_n \to \mathbb{P}(0,\infty)^{\mathcal{C}}$ is injective. (See Hint A.33.)*

  Moreover, the closure of the image of CV_n is compact; see [Besc, Section 1.7] for a sketch of this. It can be described explicitely as follows. A $\mathbb{R}$-tree is a 0-hyperbolic metric space, or equivalently a geodesic metric space where there is a unique embedded path between any two points. If T is a $\mathbb{R}$-tree with F_n acting on T, one can define a variant of the stable translation length function $g \to \hat{\tau}_g$ that only depends on the conjugacy class of g. $\overline{CV_n}$ can then be described explicitly in terms of length functions of very small actions on $\mathbb{R}$-trees [Besd].
- CV_n has a 'Borel-Serre' bordification, and the boundary of FF_n can be computed as a quotient of a subset of the boundary of CV_n [BSV18, BR15].

39.8. **RAAGs.**

- If Γ is a graph with vertex set $1, ..., n$, the right angled Artin group for Γ is
$$A_\Gamma = \langle a_1, ..., a_n : a_j a_i = a_i a_j \text{ if } (i,j) \in E(\Gamma)\rangle$$
  For example, if Γ has no edges, $A_\Gamma = F_n$, and if Γ is complete, $A_\Gamma = \mathbb{Z}^n$. Therefore, RAAGs allow interpolation between $\mathbb{Z}^n$ and F_n.
- RAAGs A_Γ and the corresponding $\mathrm{Out}(A_\Gamma)$ are closely connected to CAT(0) cube complexes. One can also define a contractible 'outer space' for A_Γ with an $\mathrm{Out}(A_\Gamma)$ action [BCV23]. For $A_\Gamma = \mathbb{Z}^n$, the corresponding outer space is the symmetric space for $SL_n(\mathbb{Z})$.
- There are many other open problems in this area. See [Cha07] for a survey.

40. Open problems (12/06, AW, YL, guest lecture by Bestvina)

We won't focus on famous open problems (like the conjugacy problem, which in any case may be approaching a solution) but rather on problems which are less studied despite their importance.

The most important question is: How similar is $\mathrm{Out}(F_n)$ to mapping class groups? One wants the $\mathrm{Out}(F_n)$ theory to catch up with the mapping class group theory.

40.1. **Actions on hyperbolic spaces.**

The following is an example of a statement known for mapping class groups but not $\mathrm{Out}(F_n)$. By "hyperbolic" we mean "Gromov hyperbolic".

Theorem 40.1. *Let $G = MCG(\Sigma)$, and let $g \in G$ have infinite order. Then there exists a finite index subgroup H of G and an action of H on a hyperbolic space such that there is an $n > 0$ with $g^n \in H$ and such that g^n acts loxodromically.*

This is already not known for $G = \mathrm{Out}(F_3)$ and g described by

$$a \mapsto b, \quad b \mapsto ab, \quad c \mapsto c.$$

It is known for $G = \mathrm{Out}(F_n)$ and g of exponential growth, but it is not known for any g of polynomial growth.

Let Γ be graph endowed with a total order on its edges. For example, $\Gamma = R_3$ with edges labeled by a, b, c with order $a < b < c$. Continuing with this example, consider the automorphisms of the form

$$\begin{array}{rcl} a & \mapsto & a \\ b & \mapsto & a^i b \\ c & \mapsto & w(a,b)cv(a,b) \end{array}$$

with $i \in \mathbb{Z}$ and $w(a,b), v(a,b) \in \langle a, b\rangle$. The set K of all these automorphisms is an example of a Kolchin subgroup of $\mathrm{Out}(F_3)$. Kolchin subgroups are defined similarly for arbitrary Γ [BFH05].

There is a homomorphism $K \to \mathbb{Z}$ which in our example send the automorphism above to i, and in fact in our example there is a short exact sequence

$$1 \to F_2 \times F_2 \to K \to \mathbb{Z} \to 1.$$

In general one can understand Kolchin subgroups inductively via such short exact sequences; they are not mysterious groups.

Problem 40.2. If K is a Kolchin subgroup of $\mathrm{Out}(F_n)$, $g \in K, g \neq 1$, is there an action of a finite index subgroup of K on a hyperbolic space such that a power of g acts loxodromically?

The expectation is that the answer will in general be no, although it may be yes for the specific K above. Note that the K above can be expressed as

$$K = (F_2 \times F_2) \rtimes \mathbb{Z}$$

via the automorphism $a \mapsto a, b \mapsto ab$. Although there are many actions of $F_2 \times F_2$ on hyperbolic spaces, one has to worry about extending them to K. A related group is

$$K' = F_2 \rtimes \mathbb{Z}$$

using the same automorphism, which is the fundamental group of a 3-manifold obtained as the mapping torus of a Dehn twist map on a punctured torus; see Figure 132. The question is not known even for K'; it is related to hierarchical hyperbolicity. (Note that this K' is famous for not being LERF.)

Remark 40.3. The Kolchin subgroup of the Γ illustrated in Figure 133 is not contained in the Kolchin subgroup of a rose.

The answer to the problem above is not expected to depend much on the graph or the ordering of its edges.

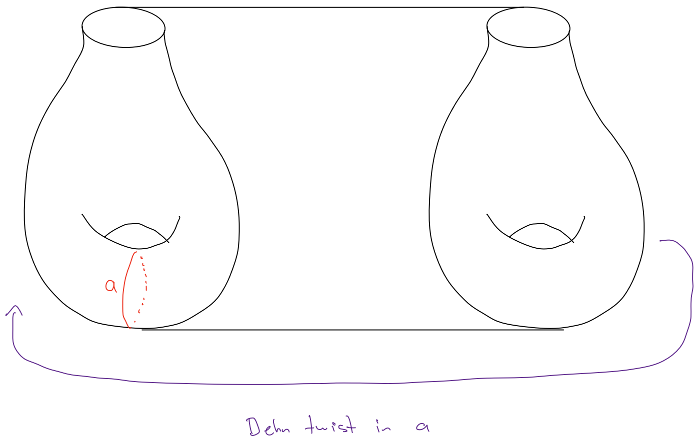

FIGURE 132

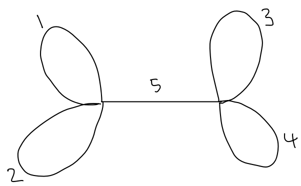

FIGURE 133

40.2. **The geometry of AF_n.** The simplicial complex AF_n is similar to FF_n but has an $\mathrm{Aut}(F_n)$ action rather than an $\mathrm{Out}(F_n)$ action.

- Vertices are proper free factors of F_n. (Unlike in FF_n, these are actual free factors, not conjugacy classes of free factors.)
- Simplices come from nested collections.

We often use the 1-skeleton, which is a graph. In the $n = 3$ case, it contains many hexagons as in Figure 134, which are reminiscent of apartments in a building.

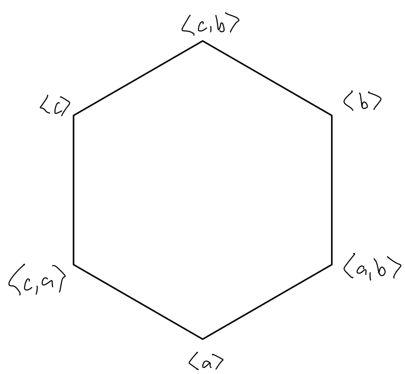

FIGURE 134

For a while it was open if AF_n is hyperbolic. But it is not hyperbolic [BBW23].

There is a short exact sequence

$$1 \to F_n \to \mathrm{Aut}(F_n) \to \mathrm{Out}(F_n) \to 1$$

which is analogous to the Birman exact sequence

$$1 \to \pi_1(\Sigma_{g,0}) \to MCG(\Sigma_{g,1}) \to MCG(\Sigma_{g,0}) \to 1.$$

So $\mathrm{Aut}(F_n)$ is analogous to the mapping class group of a punctured surface. Since curve graphs are hyperbolic even when there are punctures, one would have expected AF_n to be hyperbolic via this analogy.

So, why isn't AF_n hyperbolic? Consider a surface Σ with one boundary component and a basepoint on the boundary component. As in Figure 135, let g denote a Dehn twist about a peripheral curve. Let f be a pseudo-Anosov that fixes the basepoint.

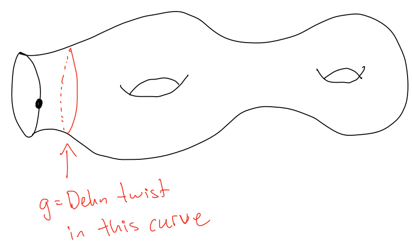

FIGURE 135

Note that $fg = gf$. One can see that f acts loxodromically on AF_n because it acts loxodromically on FF_n and there is a map $AF_n \to FF_n$. There isn't any obvious free factor fixed by g, and although it isn't obvious, it can be seen that g also acts loxodromically.

It turns out that $\langle f, g\rangle$ gives rise to a 2 dimensional quasi-flat in AF_n. (A 2 dimensional quasi-flat is a quasi-isometric embedding of $\mathbb{R}^2$ or equivalently of $\mathbb{Z}^2$.) For every element of F_n which arises as a boundary curve of a surface as above with fundamental group isomorphic to F_n one actually gets a "book" of quasi-flats, as in Figure 136. Different books are basically disjoint, and each book can be thought of as a product region.

Problem 40.4. Does AF_n become hyperbolic if all these books are electrified? Is it relatively hyperbolic with respect to books?

(Recall that the electrification of a space along a collection of subspaces is obtained by adding a cone point corresponding to each subspace, and a adding a segment of length 1 from each point in the subspace to the corresponding cone point.)

Right now we don't even know the answer to:

Problem 40.5. Does AF_n have 3-dimensional quasi-flats?

The guess is no; there is no known obstruction to hyperbolicity in AF_n other than books.

One can modify AF_n by adding an edge between each pair A, B of vertices if there exists a surface Σ as above with a marked point on its boundary and with $\pi_1(\Sigma) = F_n$

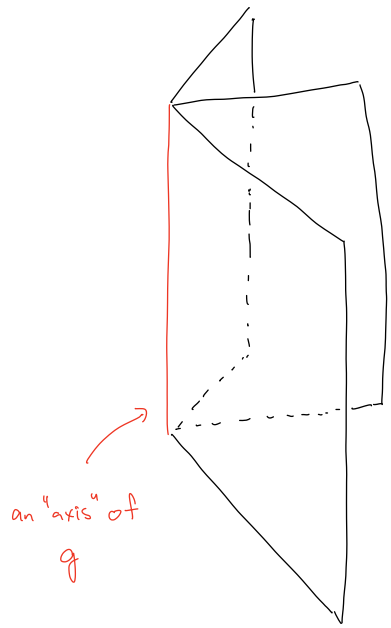

FIGURE 136

such that both A and B can be represented as subsurfaces with an embedded arc to the basepoint, as in Figure 137. (This is related to a construction of Dowdall and Taylor.)

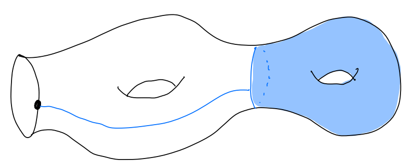

FIGURE 137

Problem 40.6. Is this modification of AF_n hyperbolic?

The modification kills all books since the pseudo-Anosov f is now elliptic, since it permutes the free factors arising as subsurfaces on the surface f is defined on. The Dehn twist g similarly becomes elliptic in the modification. It is possible the modification kills more than just the books, but the guess is that it doesn't.

40.3. **Asymptotic dimension.** Mapping class groups have finite asymptotic dimension.

Problem 40.7. Does $\mathrm{Out}(F_n)$ have finite asymptotic dimension?

The conjecture is "yes". The following definition is due to Gromov.

Definition 40.8. Let X be a metric space. We say $\operatorname{asdim} X \leqslant n$ if and only if for all $R > 0$ (think of this as large) there is a covering $\mathcal{U}$ of X (by sets that don't necessarily have to be open) such that

- $\sup_{U \in \mathcal{U}} \operatorname{diam}(U) < \infty$, and
- every R-ball in X intersects at most $n+1$ elements of $\mathcal{U}$.

A covering of $\mathbb{R}^2$ by large bricks, as in Figure 138, shows $\operatorname{asdim} \mathbb{R}^2 \leqslant 2$, and in fact $\operatorname{asdim} \mathbb{R}^n = n$. Quasi-isometric spaces have the same asymptotic dimension.

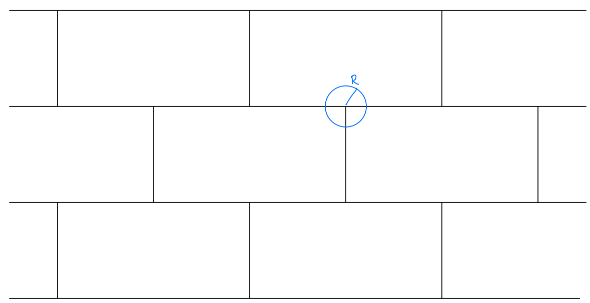

FIGURE 138

Theorem 40.9 (Bell-Fujiwara)**.** $\operatorname{asdim} \mathcal{C}(\Sigma) < \infty$.

This result on curve graphs, with work, implies $\operatorname{asdim} MCG(\Sigma) < \infty$.

Problem 40.10. Do any of FF_n, FZ_n, FS_n have asdim $< \infty$?

Here FZ_n is the $\mathbb{Z}$-splitting complex, which is similar to FS_n except the F_n actions on trees can have edge stabilizers which are $\mathbb{Z}$.

The theorem of Bell-Fujiwara is modelled on:

Theorem 40.11 (Gromov)**.** *Every hyperbolic group has* asdim $< \infty$*.*

Gromov shows that asdim is bounded in terms of the number of vertices in a ball of radius approximately 2δ in the Cayley graph. In curve graphs, balls have infinitely many vertices. The solution to this problem is to use Masur-Minsky's results on tight geodesics.

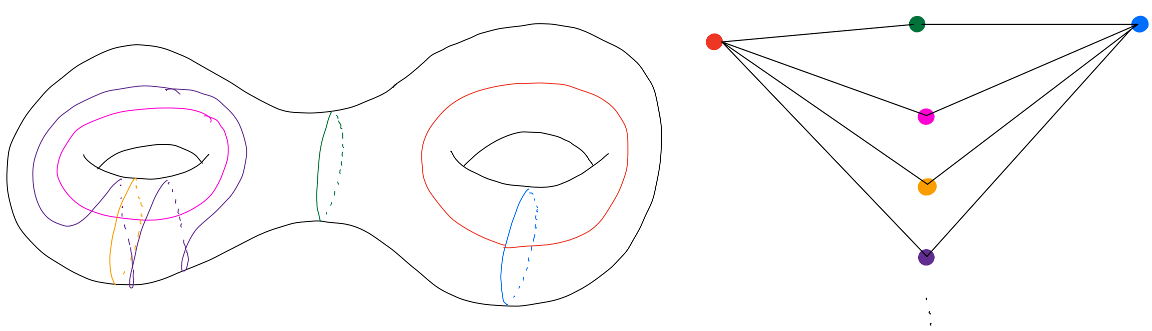

FIGURE 139. On the left are curves on a surface, and on the right is the corresponding piece of the curve graph

In $\mathcal{C}\Sigma$ there can be infinitely many geodesics joining two vertices. For example, the red and blue curves in Figure 139 are at distance 2, and there are infinitely many curves adjacent to both in the curve graph. Thus there are infinitely many geodesics between the red curve and the blue curve. There is however one preferred one, which passes through the green curve, and we call that one preferred geodesic tight. The green curve is preferred because it is homotopic to the boundary of a regular neighborhood of the union of the red and blue curves. In other examples, such a neighbourhood might have finitely many boundary curves, and one uses curves with multiple components.

A tight geodesic is defined as a geodesic in which, for any three consecutive curves, the middle is the boundary of a regular neighbourhood of the union of the other two. Masur-Minsky show that any two points in the curve graph are joined by at least one and at most finitely many tight geodesics.

Problem 40.12. Is there an analogous theory of tight geodesics in (a space quasi-isometric to) FF_n, FZ_n or FS_n?

The need to slightly modify the spaces and instead consider a quasi-isometric space can be clearly seen already in FF_3. There $[\langle a \rangle]$ and $[\langle b \rangle]$ are at distance 2. The geodesics between them pass through a point $[\langle a, gbg^{-1} \rangle]$. This gives infinitely many geodesics, with no reason to prefer one over another. The solution to this particular problem should be to add an edge from $[\langle a \rangle]$ to $[\langle b \rangle]$.

Appendix A. Hints for some exercises

Hint A.1 (Exercise 5). *If w has that form, it is easy to see $a_1, \dots, a_{n-1}, w$ is a basis. So assume w is not of that form. Consider the map from the subdivided rose labelled by $a_1, \dots, a_{n-1}, w$ to the standard rose. We can fold this "from the ends of w" until we reach a rose labelled by $a_1, \dots, a_{n-1}, v$, where v is a reduced word which begins and ends with a power of a_n. If the powers on a_n and the beginning and ending are the same, show that no more folding is possible. If the powers are opposite, show that after some initial folding then no more folding is possible. (The initial folding turns the petal labelled v into a "petal on a stem".)*

Hint A.2 (Exercise 6). *Consider a maximal tree in the cover. The edges not in this tree form a basis, so the fact that H is finitely generated implies there are only finitely many edges not in this tree.*

Hint A.3 (Exercise 7). *Say $F_n = \langle a_1, \dots, a_n \rangle$. Consider the map ϕ from a subdivided n-rose to the m-rose R_m, where the loops of the domain rose read the words $h(a_i)$. We will partially fold ϕ, by doing as many folds as possible that don't change the fundamental group of the graph. That is, we only let ourselves do type 1 and type 2 folds, and we don't stop until no more type 1 or type 2 folds are possible. This produces a map $\psi : \Gamma \to R_m$, which may not be an immersion since we haven't allowed ourselves to do type 3 and type 4 folds. Show that as we fold ψ, we only make type 3 and type 4 folds, and note that these types of folds do not change the vertex set. Keep in mind that by construction $\pi_1(\Gamma) = F_n$.*

Hint A.4 (Exercise 8). *The definition gives that $f(m)$ is m or the identity or m^2. The Abelianization of the automorphism doesn't have determinant ± 1 if $f(m) = e$ or $f(m) = m^2$. Converse, if $f(m) = m$, one can explicitly write down an inverse.*

Hint A.5 (Exercise 9). *The key is to prove what is called the "basis exchange property" in the theory of matroids: If A and B are distinct maximal trees, then there is an edge a of $A - B$ and an edge b of $B - A$ such that $(A - a) \cup b$ is a maximal tree.*

Hint A.6 (Exercise 11). *Using the identification between models of the universal cover, $f_*(\alpha)$ and $g_*(\alpha)$ have actions that are same up to conjugation by this identification. Then do the same proof as above.*

Hint A.7 (Exercise 13). *This is false for graphs like*

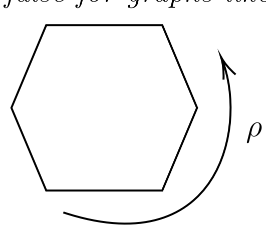

due to possible rotation, but this is not a counterexample since the valence at each vertex is at most 2. But this is true for graphs such as

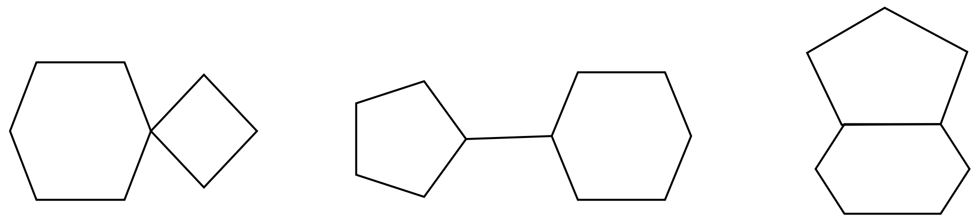

since ρ needs to preserve each immersed loop. Every relevant Γ has one of these three graphs as a subgraph.

Hint A.8 (Exercise 14)**.** *Pick an edge e that is in exactly one of F_1, F_2, and say the vertices of e are v and w. Let β be an immersed loop in $\Gamma - e$ passing through v, and let γ be an immersed loop in $\Gamma - e$ passing through w. Let α be the "dumbbell" shaped loop that traverses β starting and ending at v, crosses e, traverses γ, and crosses e to return to its starting point. Show that one of α, β, γ has the desired property.*

Hint A.9 (Exercise 16)**.** *First note that $g = a_1a_2a_2a_1a_2$ is primitive. One way to see this is via the automorphism $f : F_2 \to F_2$ defined by $f(a_1) = a_1a_2, f(a_2) = a_2a_1a_2$, computing that $f^2(a_1) = g$.*

Next note that g is not conjugate to $a_1^2a_2^3$, because both are cyclically reducded words and they are not cyclic permutations of each other.

Finally, recall that Lemma 14.12 says there is only one conjugacy class of primitive element that maps to $(2, 3)$.

Hint A.10 (Exercise 22)**.** *It is obvious though how to start defining the path, since the definition indicates how to get Γ_ϵ from Γ_0. One approach is to show the start of this path is linear in the simplex, and to continue that path until the first time t_1 that it leaves the simplex.*

Hint A.11 (Exercise 26)**.** *For this exercise, especially part (2) which might seem to be more general, it's very important that the spaces considered are trees. Trees have the property that between every pair of points there is a unique embedded (injective) path, and that every path between that pair of points must contain that unique embedded path as a subset. (You can think of that as a version of the intermediate value theorem.) Spaces with these properties are called "real trees" or "$\mathbb{R}$-trees", and everything works for those spaces too.*

Hint A.12 (Exercise 27)**.** *It should be enough to divide up $[0, 1]^k$ according to the relative sizes of the t_i and the $1 - t_i$.*

Hint A.13 (Exercise 34)**.** *Show lengths of all loops vary linearly along the line.*

Hint A.14 (Exercise 37)**.** *Since there are finitely many orbits of simplices, it suffices to show that thick part intersect the closure of any simplex is compact.*

Hint A.15 (Exercise 38)**.** *Arzela-Ascoli implies that the set of $\phi : \Gamma \to \Gamma$ that are e^B Lipschitz is compact. Nearby $\Gamma \to \Gamma$ are homotopic.*

Hint A.16 (Exercise 39)**.** *Use the discussion immediately before Lemma 25.4.*

Hint A.17 (Exercise 42)**.** *First show that every matrix in $GL(2,\mathbb{Z})$ is either finite order, or conjugate to a matrix of the form*

$$\begin{pmatrix} \pm 1 & n \\ 0 & \pm 1 \end{pmatrix},$$

where the diagonal entries have the same sign and $n \neq 0$, or has $|\lambda_1| > 1$. Then show that if Φ is reducible the homology of the reducing graph is a loop which is multiplied by ± 1 in homology, and hence $\mathrm{Ab}(\Phi)$ has the form

$$\begin{pmatrix} \pm 1 & n \\ 0 & \pm 1 \end{pmatrix}.$$

Finally, show that the automorphisms corresponding to

$$\begin{pmatrix} \pm 1 & n \\ 0 & \pm 1 \end{pmatrix}$$

are reducible, as in Exercise 23.2. Having done all this you will have "if and only if" criteria for being elliptic and for being parabolic, and so an "if and only if" criteria for being hyperbolic follows immediately since every automorphism is one of the three types.

Hint A.18 (Exercise 43)**.** *Use that τ_Φ is the largest growth rate. (That is true in general; here it can be verified by using the previous exercise to reduce to the case where Φ is irreducible and hence represented by a train track map.)*

To show $\tau_\Phi \geqslant |\lambda_1|$ consider the action on the Abelianization $\mathbb{Z}^2$.

To show $\tau_\Phi \leqslant |\lambda_1|$ consider the linear action f of $\mathrm{Ab}(\Phi)$ on $\mathbb{R}^2/\mathbb{Z}^2$ minus the origin. Note that $f_* = \Phi$. By deforming this map, show that there is a map g from $\mathbb{R}^2/\mathbb{Z}^2$ minus an ε ball around the origin to itself, which is $(|\lambda_1| + \delta)$-Lipschitz (with $\delta \to 0$ as $\varepsilon \to 0$), and with $g_* = \Phi$.

Use that the torus minus the ϵ ball around the origin has a $C(\epsilon)$-Lipschitz map to the standard rose.

Hint A.19 (Exercise 46)**.** *Actually much more is true: The assumption that A can't be conjugated into the subgroup generated by any $n-1$ of the standard generators implies there is at least one edge of each label in the image.*

Hint A.20 (Exercise 49)**.** *Let X be a proper free factor with $\alpha \in X$. Start with the core of the X cover of Γ and add loops corresponding to a basis of a complement to X.*

Hint A.21 (Exercise 50)**.** *It is the barycentric subdivision of the boundary of an n-simplex, and hence is homeomorphic to a sphere. See Figures 140, 141.*

Hint A.22 (Exercise 51)**.** *There is a bijection between the vertices of $\mathcal{P}_n$ and the vertices of FF_n that come from rank 1 free factors, given by $[\{x, x^{-1}\}] \mapsto [\langle x \rangle]$.*

Hint A.23 (Exercise 53)**.** *If $w \in \langle a_1, \ldots, a_{n-1} \rangle$, then $w = (wa_n)(a_n^{-1})$.*

Hint A.24 (Exercise 54)**.** *It is helpful to first do Exercises 48 and 49. The Whitehead graph of $\langle w \rangle$ will be the complete graph, so cannot have a cut vertex.*

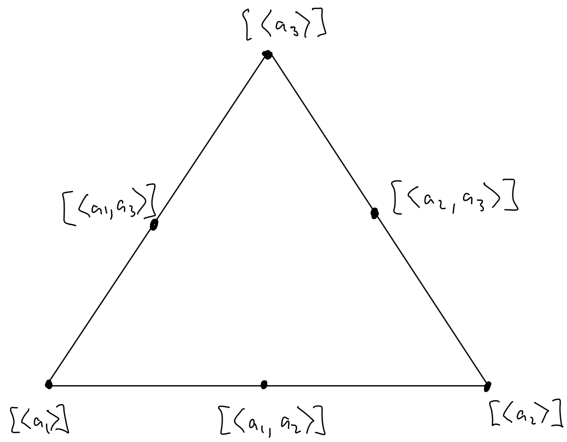

Figure 140

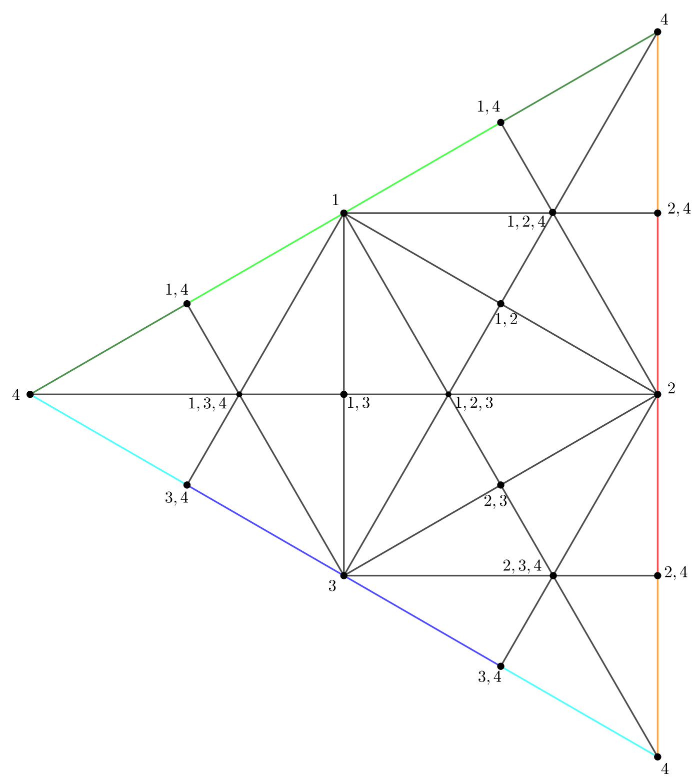

Figure 141. Edges with the same color get identified, and vertices with the same label get identified. A vertex labelled “i,j” indicates $[\langle a_i, a_j\rangle]$, etc.

Hint A.25 (Exercise 55)**.** *Let X be a proper free factor with $\alpha \in X$. Start with the core of the X cover of Γ and add loops corresponding to a basis of a complement to X.*

Hint A.26 (Exercise 58). *Define a coarse inverse map by* $v \mapsto c_v$. *Make some non-canonical definition on edges; the map does not have to be continuous. Apply Exercise 57, keeping in mind that to show that a map* f *is coarsely Lipschitz it suffices to give a uniform bound on* $d(f(x), f(y))$ *when* $d(x, y) \leqslant 1$. *For the application to outer space, Lemma 26.6 may be helpful.*

Hint A.27 (Exercise 59). *Show that the fixed point set in* T *of* g *is a subtree.*

Hint A.28 (Exercise 60). *Compare to Exercise 14 and keep in mind Exercise 35.7.*

Hint A.29 (Exercise 61). *Consider a minimal action of* F_n *on a tree* T. *If the action is free (i.e., has trivial point stabilizers), then* T/F_n *defines a point in outer space. If the action is not free, we still have a natural map* $F_n \to \pi_1(T/F_n)$. *By Bass-Serre Theory (Example 35.7), the kernel of this action is a free factor* A. *Modify* T/F_n *by wedging on a tiny rose with fundamental group* A.

Hint A.30 (Exercise 62). *The picture of* CV_2 *is a version of the Farey complex minus vertices with fins added on. So* FS_2 *has a similar picture with missing faces and vertices added in.*

Hint A.31 (Exercise 64). *You will need Example 35.7.*

Hint A.32 (Exercise 65). *First show* FS_n *is quasi-isometric to the graph whose vertices correspond to 1-edge free splitting, and where you join two vertices by an edge if there exists a 2-edge free splitting which collapses to both of them. Then apply the criterion in Exercise 58.*

Hint A.33 (Exercise 69). *To see that distinct points* $\Gamma_1, \Gamma_2 \in CV_n$ *have distinct images, use that* $d(\Gamma_1, \Gamma_2)$ *and* $d(\Gamma_2, \Gamma_1)$ *are both positive.*

References

[Aim] *Aimpl: Rigidity properties of free-by-cyclic groups*, http://aimpl.org/freebycyclic/.

[AS11] Javier Aramayona and Juan Souto, *Automorphisms of the graph of free splittings*, Michigan Math. J. **60** (2011), no. 3, 483–493.

[Asc] Dario Ascari, *A fine property of Whitehead's algorithm*, preprint, arXiv:2110.11936 (2021).

[BB] Mladen Bestvina and Martin Bridson, *Rigidity of the free factor complex*, arXiv preprint arXiv:2306.05941.

[BBP23] Tara Brendle, Nathan Broaddus, and Andrew Putman, *The mapping class group of connect sums of $S^2 \times S^1$*, Trans. Amer. Math. Soc. **376** (2023), no. 4, 2557–2572.

[BBW23] Mladen Bestvina, Martin Bridson, and Richard Wade, *On the geometry of the free factor graph for aut(f_n)*, arXiv preprint arXiv:2312.03535 (2023).

[BCV23] Corey Bregman, Ruth Charney, and Karen Vogtmann, *Outer space for RAAGs*, Duke Math. J. **172** (2023), no. 6, 1033–1108.

[Besa] Mladen Bestvina, *Folding graphs and applications, d'après Stallings*, https://www.math.utah.edu/~bestvina/5520/folding.pdf.

[Besb] ———, *Introduction to* $\mathrm{Out}(F_n)$, https://www.youtube.com/playlist?list=PLuleUFkasbo1ILZUt64kNK988_vqKTg_N.

[Besc] ———, *Notes on the geometry of outer space: Preliminary version*, https://www.math.utah.edu/~bestvina/notes.pdf.

[Besd] ———, *Outer limits*, https://www.math.utah.edu/~bestvina/eprints/bestvina.feighn..outer_limits.pdf.

[Bes11] ———, *A Bers-like proof of the existence of train tracks for free group automorphisms*, Fund. Math. **214** (2011), no. 1, 1–12.

[Bes14] ———, *Geometry of outer space*, Geometric group theory, IAS/Park City Math. Ser., vol. 21, Amer. Math. Soc., Providence, RI, 2014, pp. 173–206.

[BF00] Mladen Bestvina and Mark Feighn, *The topology at infinity of* $\mathrm{Out}(F_n)$, Invent. Math. **140** (2000), no. 3, 651–692.

[BF14a] ———, *Hyperbolicity of the complex of free factors*, Adv. Math. **256** (2014), 104–155.

[BF14b] ———, *Subfactor projections*, J. Topol. **7** (2014), no. 3, 771–804.

[BFH00] Mladen Bestvina, Mark Feighn, and Michael Handel, *The Tits alternative for* $\mathrm{Out}(F_n)$*. I. Dynamics of exponentially-growing automorphisms*, Ann. of Math. (2) **151** (2000), no. 2, 517–623.

[BFH05] ———, *The Tits alternative for* $\mathrm{Out}(F_n)$*. II. A Kolchin type theorem*, Ann. of Math. (2) **161** (2005), no. 1, 1–59.

[BH92] Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51.

[BR15] Mladen Bestvina and Patrick Reynolds, *The boundary of the complex of free factors*, Duke Math. J. **164** (2015), no. 11, 2213–2251.

[Bri00] P. Brinkmann, *Hyperbolic automorphisms of free groups*, Geom. Funct. Anal. **10** (2000), no. 5, 1071–1089.

[BSV18] Kai-Uwe Bux, Peter Smillie, and Karen Vogtmann, *On the bordification of outer space*, J. Lond. Math. Soc. (2) **98** (2018), no. 1, 12–34.

[BV00] Martin R. Bridson and Karen Vogtmann, *Automorphisms of automorphism groups of free groups*, J. Algebra **229** (2000), no. 2, 785–792.

[BV01] ———, *The symmetries of outer space*, Duke Math. J. **106** (2001), no. 2, 391–409.

[BV23] Michael Borinsky and Karen Vogtmann, *The Euler characteristic of the moduli space of graphs*, Adv. Math. **432** (2023), Paper No. 109290, 58.

[Cha07] Ruth Charney, *An introduction to right-angled Artin groups*, Geom. Dedicata **125** (2007), 141–158.

[Cla05] Matt Clay, *Contractibility of deformation spaces of G-trees*, Algebr. Geom. Topol. **5** (2005), 1481–1503.

[CV86] Marc Culler and Karen Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math. **84** (1986), no. 1, 91–119. MR 830040

[DD89] Warren Dicks and M. J. Dunwoody, *Groups acting on graphs*, Cambridge Studies in Advanced Mathematics, vol. 17, Cambridge University Press, Cambridge, 1989.

[FM12a] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.

[FM12b] Stefano Francaviglia and Armando Martino, *The isometry group of outer space*, Adv. Math. **231** (2012), no. 3-4, 1940–1973.

[FP92] Edward Formanek and Claudio Procesi, *The automorphism group of a free group is not linear*, J. Algebra **149** (1992), no. 2, 494–499.

[Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.

[Hat22] ———, *Topology of numbers*, American Mathematical Society, Providence, RI, [2022] ©2022.

[Hen20] Sebastian Hensel, *A primer on handlebody groups*, Handbook of group actions. V, Adv. Lect. Math. (ALM), vol. 48, Int. Press, Somerville, MA, [2020] ©2020, pp. 143–177.

[HH17] Arnaud Hilion and Camille Horbez, *The hyperbolicity of the sphere complex via surgery paths*, J. Reine Angew. Math. **730** (2017), 135–161.

[HM13] Michael Handel and Lee Mosher, *The free splitting complex of a free group, I: hyperbolicity*, Geom. Topol. **17** (2013), no. 3, 1581–1672.

[HM19] ———, *The free splitting complex of a free group, II: Loxodromic outer automorphisms*, Trans. Amer. Math. Soc. **372** (2019), no. 6, 4053–4105.

[Hor16] Camille Horbez, *A short proof of Handel and Mosher's alternative for subgroups of* $Out(F_N)$, Groups Geom. Dyn. **10** (2016), no. 2, 709–721.

[HV] Allen Hatcher and Karen Vogtmann, *The complex of free factors of a free group*, `https://pi.math.cornell.edu/~hatcher/Papers/FCn.pdf`.

[HW19] Michael Heusener and Richard Weidmann, *A remark on Whitehead's cut-vertex lemma*, J. Group Theory **22** (2019), no. 1, 15–21.

[KKN21] Marek Kaluba, Dawid Kielak, and Piotr W. Nowak, *On property (T) for* $Aut(F_n)$ *and* $SL_n(\mathbb{Z})$, Ann. of Math. (2) **193** (2021), no. 2, 539–562.

[KR14] Ilya Kapovich and Kasra Rafi, *On hyperbolicity of free splitting and free factor complexes*, Groups Geom. Dyn. **8** (2014), no. 2, 391–414.

[Lev09] Gilbert Levitt, *Counting growth types of automorphisms of free groups*, Geom. Funct. Anal. **19** (2009), no. 4, 1119–1146.

[Mar95] Reiner Martin, *Non-uniquely ergodic foliations of thin-type, measured currents and automorphisms of free groups*, ProQuest LLC, Ann Arbor, MI, 1995, Thesis (Ph.D.)–University of California, Los Angeles.

[MT18] Joseph Maher and Giulio Tiozzo, *Random walks on weakly hyperbolic groups*, J. Reine Angew. Math. **742** (2018), 187–239.

[Nit] Martin Nitsche, *Computer proofs for Property (T), and SDP duality*, preprint, arXiv:2009.05134 (2022).

[Ser03] Jean-Pierre Serre, *Trees*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003, Translated from the French original by John Stillwell, Corrected 2nd printing of the 1980 English translation.

[Sko] Richard Skora, *Deformations of length functions in groups*, `https://levitt.users.lmno.cnrs.fr/SkoraDeformation.pdf`.

[Sta] *How equivalent are the homeomorphism and diffeomorphism groups of 3-manifolds?*, `https://math.stackexchange.com/q/1194926`.

[Sta83] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565.

[SW79] Peter Scott and Terry Wall, *Topological methods in group theory*, Homological group theory (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser., vol. 36, Cambridge Univ. Press, Cambridge-New York, 1979, pp. 137–203. MR 564422

[Vog] Karen Vogtmann, *Outer space and automorphisms of free groups*, `https://www.ams.org/open-math-notes/omn-view-listing?listingId=110672`.

[Vog02] ———, *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.

[Vog08] ———, *What is... outer space?*, Notices Amer. Math. Soc. **55** (2008), no. 7, 784–786.

[Vog15] ———, *On the geometry of outer space*, Bull. Amer. Math. Soc. (N.S.) **52** (2015), no. 1, 27–46. MR 3286480

[Wil] Henry Wilton, *Group actions on trees*, `https://www.dpmms.cam.ac.uk/~hjrw2/Talks/trees.pdf`.