Moduli spaces of graphs

Karen Vogtmann

Abstract. We study the combinatorial and topological structure of various moduli spaces of metric graphs, with the aim of highlighting some recent advances in the subject. In order to put the recent advances into perspective we will review basic definitions and some of the historical developments. Many of these historical developments have to do with how spaces of graphs are used to study groups of automorphisms and outer automorphisms of a finitely-generated free group. We also mention connections with Kontsevich's graph complexes and with objects studied in tropical algebraic geometry. We then point out a relation between spaces of graphs and spaces of flat tori, and end by describing moduli spaces of a new class of non-positively curved metric spaces that generalize both metric graphs and flat tori, and that are associated to right-anged Artin groups.

1 Introduction. The study of finite graphs is a well-established branch of mathematics in its own right, dating from the early 18th century when Euler analyzed the seven bridges of Köningsburg. It is also of interest in many other areas of mathematics and science because finite graphs are widely used to represent objects of interest in those areas. There are numerous example, but here we will mention just four of them. In biology, evolutionary relationships between species are represented by finite graphs called phylogenetic trees. In physics, terms in the perturbative expansion of a quantum field theory are indexed by finite graphs called Feynman diagrams. In algebraic geometry, the ways in which a smooth algebraic curve can degenerate to a nodal curve are classified by finite graphs that are dual to systems of simple closed loops in the curve. In algebraic topology the fundamental group connects the topology of finite graphs to algebraic properties of free groups.

In many applications it is natural to assign a positive real number to each edge in a graph. In a phylogenetic tree this number may describe the length of time between speciation events. In a Feynman diagram it may designate the value of a Schwinger parameter for the associated Feynman integral. In an algebraic curve it may encode the rate at which a loop is shrinking as the curve degenerates. We will see presently that it is also useful for the study of free groups and their automorphisms.

A graph with positive real numbers assigned to its edges is can be given the structure of a metric space in which the real number on an internal edge is its length. The Gromov-Hausdorff topology on metric spaces then restricts to a topology on the space of all metric graphs. Much of this article will focus on various subspaces of this space and on closely related spaces. We will be particularly interested in their combinatorial and topological structure.

It would be impossible to survey everything that has been learned about spaces of metric graphs and their applications since I started working on them in the mid-1980's. Instead, in this report I will give a basic introduction to spaces of graphs and then outline a few recent developments and applications that I have personally been involved in. To get there I will also need to explain some of the earlier developments. Many of these developments (both recent and historical) have to do with the group $\mathrm{Out}(F_n)$ of outer automorphisms of a free group, as I will explain. The relevant space of graphs was epinominously named Outer space by Peter Shalen soon after it was introduced, and has been the focus of a great deal of research since then. I would like to apologize here for the fact that fundamental contributions by many, many people are not covered here.

2 Spaces of graphs. In this section we define the most basic objects we would like to understand, namely moduli spaces of finite connected graphs with a fixed Euler characteristic, and work out explicitly what these are in the simplest cases. First we need to say exactly what we mean by a finite graph!

2.1 What is a graph? Informally, a graph consists of a set of vertices connected by edges, where vertices may be connected by several edges (or no edges) and vertices may have single edges (loops) attached. Another way to say this is that a graph is a 1-dimensional CW-complex. The *rank* of a graph is the rank of its (free abelian) first homology group; if the graph is connected this is also the rank of its (free) fundamental group. We will primarily be interested in finite connected graphs with $s \geq 0$ univalent vertices. The pair consisting of a univalent vertex and its adjacent edge is called a *leaf*, and we usually label the leaves with the numbers $\{1, \ldots, s\}$. Graphs G and G' are isomorphic if and only if there is a cellular homeomorphism $G \to G'$ that preserves the leaf-labels.

To make the set of finite connected graphs into a topological space, we assign each edge e a positive real length ℓ_e. The graph can then be given a metric where each open edge e is isometric to an open interval of $\mathbb{R}$ of length $\ell_e > 0$, and the distance between two points is given by the length of the shortest path joining them. The topology on this set of metric graphs is then given by the Gromov-Hausdorff topology on metric spaces. This means that metric graphs are ε-close if, given any finite set of points in one graph, one can find corresponding points in the other such that all pairwise distances are within ε of each other. Notice that nearby graphs are not necessarily homeomorphic (see Figure **??**).

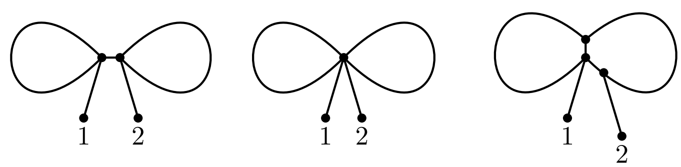

Figure 2.1: Nearby graphs

Notice also that the metric cannot detect bivalent vertices, so we will not allow them in our graphs, i.e. we assume that each vertex which is not at the end of a leaf has valence at least 3. We will use the following term throughout this article.

Definition 2.1. *A graph G is* admissible *if it is finite, connected, has no bivalent vertices, and its s univalent vertices are labeled by the numbers $\{1, \ldots, s\}$.*

Combinatorial graphs. The following combinatorial description of a graph is often very useful. It comes from the observation that if you remove all the leaf vertices and snip all the internal edges of a graph you are left with a collection of *corollas*, where a corolla is a vertex with half-edges attached. Thus a graph can be specified by giving a set H (the half-edges), a partition V of H (with one block for each vertex, consisting of the half-edges attached to the vertex), and an involution ι of H (whose order-two orbits are the internal edges of the graph and fixed points are the leaves). This has the convenient feature of clearly distinguishing between internal edges and leaves. It does not count the number of isolated vertices, but for connected graphs this does not concern us. Our conditions that every univalent vertex is a leaf and there are no bivalent vertices translates to the condition that each block of the partition V of half-edges has at least 3 elements.

Graphs $G = (H, V, \iota)$ and $G' = (H', V', \iota')$ are isomorphic if and only if there is a bijection $H \to H'$ sending V to V' and ι to ι'. Metric graphs are isometric if and only if this bijection also preserves edge-length.

As already mentioned, we will assume our graphs are connected. We will further restrict attention to the subspace consisting of graphs with a fixed rank n and a fixed number s of labeled leaves. The next subsections examine some small examples in detail.

2.2 The space of trees with s leaves. If the rank of a connected graph is 0 the graph is a tree. Since each block of the partition V defining the graph must have at least 3 half-edges, the tree has at least 3 leaves.

Note that the space of all metric graphs (of any rank) with s leaves has a factor $(\mathbb{R}_{>0})^s$ corresponding to leaf lengths. Since we understand this factor very well both geometrically and topologically, from now on we will ignore the lengths of leaves, i.e. we will only assign lengths to the (internal) edges. (Alternatively, we could arbitrarily fix the i-th leaf-lengths to be constants $c_i \in (0, \infty]$.) We denote the resulting parameter space $\mathcal{T}_s$.

For $s = 3$ there is only one combinatorial type of tree, namely a 3-corolla, so $\mathcal{T}_3$ is a point. For $s = 4$ there are four combinatorial types, one corolla and three one-edge trees distinguished by the distribution of their leaf-labels. The space $\mathcal{T}_4$ is thus a union of three rays, paramatrizing the length of the internal edge, meeting at the point corresponding to the corolla (see Figure **??**).

The geometry of $\mathcal{T}_s$ for any s is well-understood. Namely, using Gromov's link condition one can see that $\mathcal{T}_s$ is non-positively curved (i.e. CAT(0)) [**?**]. This implies, in particular, that there are unique geodesics between

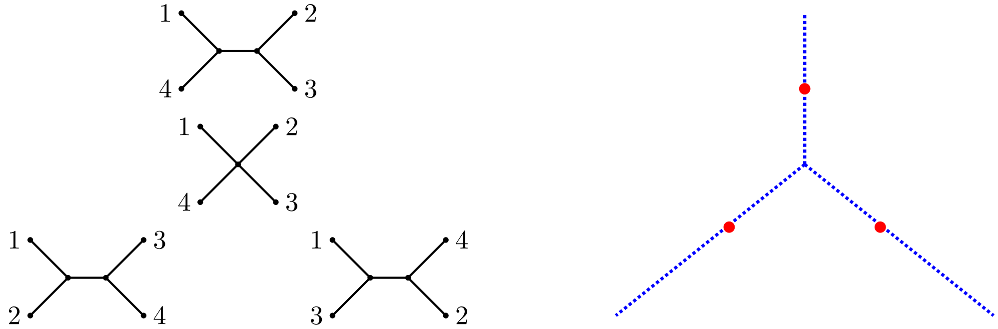

Figure 2.2: Trees with 4 leaves, and the space $\mathcal{T}_4$ with its link $Lk(\mathcal{T}_4)$

points, which can be calculated algorithmically in polynomial time [**?**].

The topology of $\mathcal{T}_s$ is also well understood. For any s, $\mathcal{T}_s$ will have the s-corolla as a cone point, so all of the interesting topological structure is contained in the *link* of the cone point, denoted $Lk(\mathcal{T}_s)$. This link can be realized as the space of trees whose internal edge-lengths sum to 1. Since a tree with four leaves has at most one internal edge, $Lk(\mathcal{T}_4)$ consists of three points, which we think of as a wedge of two 0-spheres. In fact, for any s we have the following theorem.

THEOREM 2.2 ([**?**, **?**]). *For $s \geq 4$, $Lk(\mathcal{T}_s)$ is homotopy equivalent to a wedge of $(s-2)!$ spheres of dimension $(s-4)$.*

The symmetric group Σ_s acts on $Lk(\mathcal{T}_s)$ by permuting the leaf-labels. This induces an action on the non-vanishing homology group $H^{s-2}(Lk(\mathcal{T}_s); \mathbb{R})$, giving a finite-dimensional representation of Σ_s. This representation was determined explicitly by Robinson and Whitehouse in [**?**].

2.3 The moduli space $\mathcal{MG}_{n,s}$ of metric graphs of rank n with s leaves. Motivated by the example of the space of trees, we now alter our definition of metric graphs to ignore leaf lengths.

DEFINITION 2.3. *A* metric graph *is a combinatorial graph $G = (H, V, \iota)$ together with an assignment ℓ of positive real lengths to all (non-leaf) edges.*

DEFINITION 2.4. *The* moduli space of metric graphs $\mathcal{MG}_{n,s}$ *is the space of (connected) metric graphs (H, V, ι, ℓ) with*
- *rank g*
- *s leaves*
- *all vertices of valence at least 3*
- *the sum of the edge-lengths equal to 1*

Two such graphs $G = (H, V, \iota, \ell)$ and $G' = (H', V', \iota', \ell')$ represent the same point in $\mathcal{MG}_{n,s}$ if there is a bijection $H \to H'$ that sends V to V', ι to ι' and ℓ to ℓ'.

For many applications one would like to restrict attention to graphs with no separating edges (called *one-particle-irreducible* (1-PI) graphs in physics). For $n > 0$ the space $\mathcal{MG}_{n,s}$ deformation retracts to the subspace $\mathcal{MG}^{red}_{n,s}$ of 1-PI graphs by shrinking all separating edges (not leaves) linearly to zero and expanding all other edges linearly to main the condition that the sum of the edge-lengths is equal to 1.

2.3.1 Explicit description of $\mathcal{MG}_{2,0}$. This example is illustrated in Figure **??**. There are three combinatorial types of graphs of rank 2 with no leaves, namely the *rose*, the *theta graph* and the *barbell.* The rose has one vertex and two edges, so normalizing the sum of the edge lengths to 1 seems to give an open unit interval, which we think of as the interior of a 1-simplex with a vertex for each edge. However, switching the edge-lengths gives an isometric graph, so the moduli space consists of only half of this interval. The theta graph and the barbell each seem to give the interior of a 2-simplex, but again different points in the simplex give the same metric graph. To get a unique point for each metric graph one must take a fundamental domain for the action of the automorphism group of the combinatorial graph, since this group permutes the edges of the graph, which correspond to the vertices of the simplex.

The reduced moduli space $\mathcal{MG}_{2,0}^r$ consists only of the smaller triangle, as it does not include the barbell graphs.

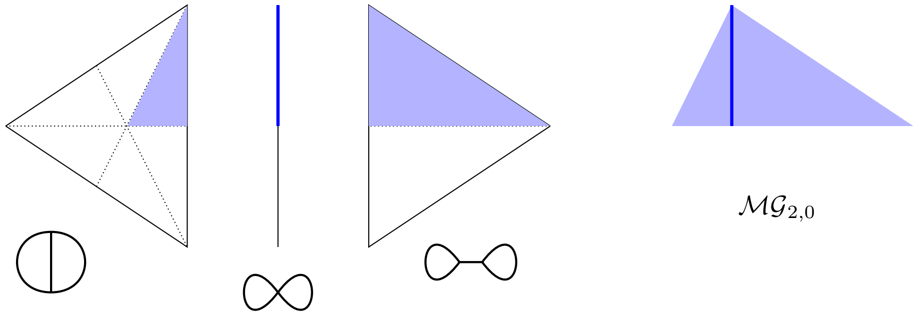

Figure 2.3: The moduli space $\mathcal{MG}_{2,0}$.

2.3.2 Explicit description of $\mathcal{MG}_{1,s}^{red}$. The only combinatorial graphs with rank 1, vertices at least trivalent and no separating edges are circles with at least one leaf attached. If we choose the vertex where the leaf labeled 1 is attached for a basepoint, then each other leaf seems to contribute a factor of S^1 to $\mathcal{MG}_{1,s}^{red}$, since it can be attached at any point on the circle. However, flipping the circle while fixing the basepoint gives an isometric graph, so $\mathcal{MG}_{1,s}$ is actually the quotient of the torus $(S^1)^{s-1}$ by an involution. If s is even, this involution reverses the orientation of the torus, otherwise it preserves orientation. For $s = 2$ we get a half-circle, and for $s = 3$ we get the quotient of the torus $S^1 \times S^1$ by the elliptic involution, resulting in a 2-sphere. In general for s odd the space $\mathcal{MG}_{1,s}^{red}$ can be identified with a *Kummer variety* (of complex dimension $(s-1)/2$).

If we allow separating edges, then there are additional combinatorial types of graphs, namely circles with rooted trees attached instead of single leaves. For $s = 2$ there is only one additional graph, a rooted tree with two leaves. This adds an open half-interval to $\mathcal{MG}_{1,2}^{red}$, glued by its endpoint to one endpoint of the half-circle. This is the first example we have seen of a moduli space of graphs which is not compact. For higher numbers s of leaves the picture becomes more complicated, since it must incorporate the complexity of spaces of rooted trees.

3 Outer space. For $n \geq 2$ the topology of the moduli space $\mathcal{MG}_{n,s}$ becomes more difficult to understand, for example we can compute its homology completely only for very small values of n. Miraculously, equipping graphs with *markings* that identify their fundamental groups with free groups results is a much better-behaved space, as we explain in this section. We first illustrate this in the case of graphs with no leaves, then proceed in the next section to the general case.

Definition 3.1 (Marked graph). *Suppose $n \geq 2$ and let R_n be the rose graph, consisting of a single vertex with n loops attached which are oriented and labeled with the generators of F_n. A* marking *of a graph G is a homotopy equivalence g from R_n to G. Pairs $\{G, g\}$ and $\{G', g'\}$ are* equivalent *if there is a graph isomorphism $f\colon G \to G'$ such that $f \circ g$ is homotopic to g'. A* marked graph *(G, g) is an equivalence class of pairs.*

Remark 3.2. Suppose $f\colon G \to G$ is a graph automorphism. Then $f^{-1} \circ (f \circ g) = g$, so by the above definition $(G, f \circ g)$ and (G, g) are the same marked graph.

To each admissible marked graph (G, g) we can associate an open simplex $\sigma(G, g)$ of marked *metric* graphs by assigning positive real lengths to the edges, normalized so that the sum is equal to 1. Since each vertex of an admissible graph G is at least trivalent and the Euler characteristic of G is $1 - n$, it is easy to compute that $\sigma(G, g)$ has dimension at most $3n - 4$, and $3n - 4$ is achieved if and only if G is trivalent. Marked metric graphs (G, g) and (G', g') are the same if and only if there is an *isometry* $f\colon G \to G'$ with $f \circ g \simeq g'$ In particular, distinct points of $\sigma(G, g)$ are distinct marked metric graphs even if they are isometric as unmarked graphs.

Example 3.3. Let G be a rose with two petals. A homotopy equivalence $g\colon R_2 \to G$ sends each petal of R_2 to a different loop in G. The graph automorphism that switches the loops fixes only the center point of $\sigma(G, g)$, since unless the two loops have the same length this graph automorphism is not an isometry. Similarly, all points in the simplices associated to the theta graph and the barbell graph are distinct marked metric graphs, even though they may be isometric (unmarked) metric graphs.

Since we are assuming $n \geq 2$ there are an infinite number of homotopy classes of homotopy equivalences $R_n \to G$ for any admissible G, so there are also an infinite number of open simplices $\sigma(G, g)$. A face of $\sigma(G, g)$ is obtained by setting some edge lengths equal to zero. If these edges form a *forest*, i.e. a subgraph with no loops, then collapsing them produces a new graph G' with the same fundamental group. Thus if $c\colon G \to G'$ is the map that collapses each edge of the forest to a point, the open simplex $\sigma(c \circ g, G')$ corresponds to a face of $\sigma(G, g)$. We say such a face is a *finite face*. Each finite face of $\sigma(G, g)$ is also a finite face of at least two other simplices $\sigma(G', g')$.

Definition 3.4 (Outer space). *Let $\widehat{\sigma}(G, g)$ denote the union of $\sigma(G, g)$ and all of its finite faces.* Outer space *$\mathcal{CV}_n$ is the union of the $\widehat{\sigma}(G, g)$ for all admissible marked graphs (G, g) of rank n, glued together along common finite faces:*

$$\mathcal{CV}_n := \bigcup_{(G,g)} \widehat{\sigma}(G, g)/\sim \textit{ of finite faces}$$

We also define *reduced* Outer space $\mathcal{CV}_n^{red}$ to be the subspace spanned by graphs with no separating edges. As in the case of the moduli space $\mathcal{MG}_{n,s}$, shrinking all separating edges linearly to zero gives a deformation retraction of $\mathcal{CV}_n$ to $\mathcal{CV}_n^{red}$.

Outer space $\mathcal{CV}_2$ is illustrated in Figure **??**. In this figure the triangles labeled by marked theta-graphs are drawn to fill out a copy of the hyperbolic plane, which represents $\mathcal{CV}_2^{red}$, and it is clear that all of $\mathcal{CV}_2$ can be retracted into $\mathcal{CV}_2^{red}$ by simply "pushing down" the yellow triangles.

Although Figure **??** makes it look as though $\mathcal{CV}_2$ is a simplicial complex, it is not, since all of the vertices are missing as well as the edges not contained in the hyperbolic plane. Adding the missing faces gives the *simplicial closure* CV_2^*, which is defined as the smallest simplicial complex that contains CV_2 as a union of open simplices.

The markings on a simplex $\sigma(G, g)$ in the figure are indicated by arrows and labels on the complement of a maximal tree. Specifically, collapsing the tree and sending each remaining edge to the directed edge-path in R_n indicated by its label gives a map $h\colon G \to R_n$ which is a homotopy inverse to the marking $g\colon R_n \to G$. In fact, we could equally well have defined a marking to be a homotopy equivalence from $G \to R_n$.

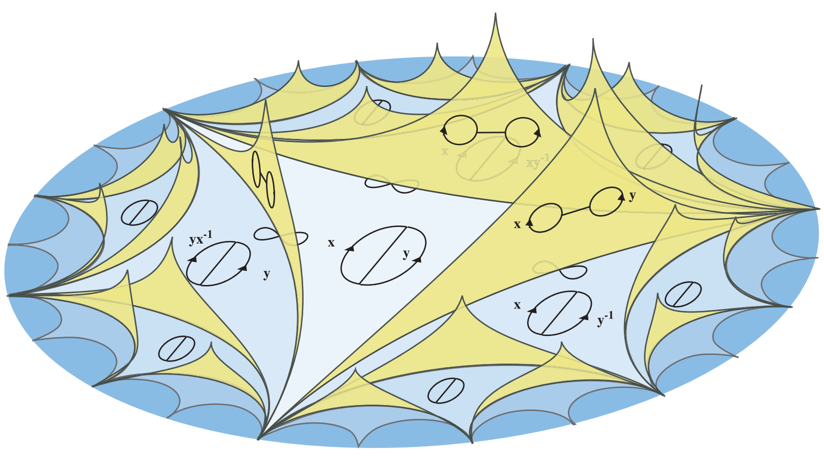

Figure 3.1: Outer space in rank 2

3.1 Action of the group of outer automorphisms of the free group. The original motivation for defining $\mathcal{CV}_n$ was to study the natural action of the group $\mathrm{Out}(F_n)$ of outer automorphisms of the free group F_n on $\mathcal{CV}_n$. To define this action, recall that the petals of R_n are oriented and labeled with the generators of F_n, so $\pi_1(R_n)$ is canonically identified with F_n and any homotopy equivalence $h\colon R_n \to R_n$ induces an outer automorphism $h_*\colon F_n \to F_n$ (it is only well-defined up to inner automorphisms because we do not assume that h

fixes the basepoint). The set of homotopy classes of homotopy equivalences of R_n forms a group $\pi_0(HE(R_n))$, and the map sending the homotopy class of h to the outer automorphism class of h_* is an isomorphism. In symbols the map $[h] \mapsto [h_*]$ induces an isomorphism $\pi_0(HE(R_n)) \cong \mathrm{Out}(F_n)$.

The action is now defined as follows. Given $\alpha \in \mathrm{Out}(F_n)$, let $h_\alpha \colon R_n \to R_n$ be a homotopy equivalence representing α. Then $(G, g) \cdot \alpha = (G, g \circ h_\alpha)$. Note that the action does not change the isometry type of the graph, it only changes the marking. The following fundamental theorem about Outer space and the action of $\mathrm{Out}(F_n)$ was established in 1986.

THEOREM 3.5 ([?]). *For all $n \geq 2$, Outer space $\mathcal{CV}_n$ is contractible of dimension $3n-4$, and the action of $\mathrm{Out}(F_n)$ on $\mathcal{CV}_n$ has finite stabilizers.*

Finiteness of the stabilizers follows since the stabilizer of a point (G, g) is isomorphic to the group of isometries of the metric graph G. This in turn shows that the quotient $\mathcal{CV}_n/\mathrm{Out}(F_n)$ is equal to the moduli space $\mathcal{MG}_{n,0}$ of metric graphs defined in Section **??**.

3.2 Hurewicz's theorem and the cohomology of $\mathrm{Out}(F_n)$**.** In [?] Baumslag and Taylor proved that the kernel IA_n of the natural map from $\mathrm{Out}(F_n)$ to $\mathrm{GL}(n, \mathbb{Z})$ is torsion-free. Torsion-free finite-index subgroups of $\mathrm{GL}(n, \mathbb{Z})$ are plentiful, so pulling back any one of them gives a torsion-free finite-index subgroup of $\mathrm{Out}(F_n)$. Since stabilizers of points in $\mathcal{CV}_n$ are finite, any torsion-free subgroup H therefore freely on $\mathcal{CV}_n$, so the quotient by the action has fundamental group H and the map $\mathcal{CV}_n \to \mathcal{CV}_n/H$ is a covering map. By a theorem of Hurewicz this implies in particular that the cohomology of the quotient is an invariant of H. Since the dimension of $\mathcal{CV}_n$ is $3n-4$, an immediate consequence of Theorem **??** is that the cohomology of any torsion-free subgroup of finite index $\mathrm{Out}(F_n)$ vanishes above dimension $3n-4$. We say $\mathrm{Out}(F_n)$ has *virtual cohomological dimension* at most $3n-4$, and write $\mathrm{VCD}(\mathrm{Out}(F_n)) \leq 3n-4$.

A more sophisticated argument shows that the cohomology of $\mathcal{MG}_n = \mathcal{CV}_n/\mathrm{Out}(F_n)$ with trivial rational coefficients coincides with the group cohomology of $\mathrm{Out}(F_n)$ with trivial rational coefficients, where group cohomology is an invariant that can be defined purely algebraically.

3.3 The sphere complex. In [?] Hatcher gave a different way to describe Outer space in terms of spheres in a doubled handlebody, which we now describe.

Let W_n be the connected sum of n copies of $S^1 \times S^2$, and fix a homotopy equivalence $c : W_n \to R_n$ where R_n is the standard rose. The 3-manifold W_n can be formed by taking two copies of the genus n handlebody $V_n = \#_n(S^1 \times D^2)$ and gluing their boundaries together by the identity, so we call W_n a *doubled handlebody.* A 2-sphere embedded in W_n is *trivial* if it bounds a 3-ball, otherwise it is *essential.* The *sphere complex* $\mathcal{S}(W_n)$ is the simplicial complex whose vertices are isotopy classes of essential 2-spheres, and whose k-simplices are sets of $k+1$ distinct isotopy classes which can be represented by disjoint spheres. Such a set of (isotopy classes of) spheres is called a *sphere system.* A point of $\mathcal{S}(W_n)$ can be specified by giving the vertices $s_0, \ldots, s_k$ of the simplex that contains it together with barycentric coordinates for the point, i.e. positive real *weights* on the individual spheres s_i that sum to 1.

Given any sphere system $S = \{s_0, \ldots, s_k\}$ one can form the *dual graph* G_S. This graph has one vertex for each component of the complement of S in W_n, and one edge for each sphere. The endpoint(s) of this edge are the (one or two) components on either side of the sphere.

A sphere system S is said to be *complete* if the components of $W_n \setminus S$ are all simply-connected. In this case the dual graph G_S can be realized uniquely up to homotopy as a subspace of W_n with one vertex in each component of $W_n \setminus S$ and one edge which intersects its associated sphere in exactly one point. The inclusion $i_S \colon G_S \hookrightarrow W_n$ is a homotopy equivalence, and the weights on the individual spheres can be interpreted as lengths on the edges of G_S, making G_S into a metric graph. Since we have fixed a homotopy equivalence $c \colon W_n \to R_n$, the composition $c \circ i_S$ gives a *marking* on G_S, i.e. $(h_n \circ i_S, G_S)$ is a point of $\mathcal{CV}_n$. Notice that here the markings are going from G to R_n, unlike in the previous section, and $\mathrm{Out}(F_n)$ acts on the left instead of the right.

Note that the subspace $\mathcal{CS}(W_n)$ of $\mathcal{S}(W_n)$ consisting of points represented by complete sphere systems is not a subcomplex, since the simplex containing such a point has missing faces. For instance all of its vertices are missing, since a single sphere cannot cut W_n into simply-connected pieces. On the other hand, the set of points represented by incomplete sphere systems *is* a subcomplex, denoted $\mathcal{S}_\infty(W_n)$.

THEOREM 3.6 (Hatcher [?]). *The map $\mathcal{CS}(W_n) \to \mathcal{CV}_n$ sending a complete sphere system S to its (marked) dual graph $(h_n \circ i_S, G_S)$ is a homeomorphism.*

Hatcher proves that the entire sphere complex $\mathcal{S}(W_n)$ is contractible using a surgery argument to construct a deformation retraction to a point. This deformation retraction restricts to a deformation retraction of $\mathcal{CS}(W_n)$, giving another proof that $\mathcal{CV}_n$ is contractible.

The mapping class group of W_n is the group $\pi_0(\mathrm{Homeo}(W_n))$ of homotopy classes of homeomorphisms of W_n, and it clearly acts on $\mathcal{S}(W_n)$. The natural map $\pi_0(\mathrm{Homeo}(W_n)) \to \mathrm{Out}(F_n)$ is surjective, and a theorem of Laudenbach [?] identifies the kernel as a finite 2-group generated by Dehn twists in 2-spheres. Hatcher uses the fact that such Dehn twists act trivially on essential 2-spheres to identify the action of $\pi_0(\mathrm{Homeo}(W_n))$ on $\mathcal{CS}(W_n)$ with the action of $\mathrm{Out}(F_n)$ on $\mathcal{CV}_n$.

Thus the sphere complex picture gives a complete alternative description of Outer space with its action of $\mathrm{Out}(F_n)$. It also gives a natural way to describe the simplicial closure $\mathcal{CV}_n^*$ of $\mathcal{CV}_n$ (which is defined as the smallest simplicial complex containing $\mathcal{CV}_n$ as a subspace), namely we can identify the simplicial closure with the full sphere complex $\mathcal{S}(W_n)$.

Since each vertex of $\mathcal{CV}_n^*$ is represented by a single sphere in W_n (up to isotopy), one advantage of the sphere complex description is that an edge of a marked graph (G, g) is a well-defined object, independent of which marked graph contains it. A further advantage of is that it easily generalizes to account for graphs with leaves, as we explain in the next section.

4 Spaces of marked graphs with leaves.

We now allow our graphs to have s leaves, labeled with the numbers $1, \ldots, s$. Define the *standard thorned rose* $R_{n,s}$ to be the graph with one vertex, n loops oriented and labeled as before with the generators of F_n and s leaves labeled $1, \ldots, s$. A *marking* of a graph G with s labeled leaves is a homotopy equivalence $g\colon R_{n,s} \to G$ that sends the leaf of R_n labeled i to the corresponding leaf of G. Pairs $\{G, g\}$ and $\{g', G'\}$ are *equivalent* if there is a graph isomorphism $h\colon G \to G'$ preserving the leaf labels such that g' is homotopic to $h \circ g$ by a homotopy fixing the leaves, and a *marked graph* (G, g) is an equivalence class of pairs. We now define $\mathcal{CV}_{n,s}$ to be the space of marked *metric* graphs, where we recall that the metric assigns lengths summing to 1 to the internal edges (not to the leaves).

For $n = 0$ the graph $R_{0,s}$ has no loops, i.e. it is a *corolla* with s labeled leaves. Any two markings of a tree with s labeled leaves are homotopic by a label-preserving homotopy, so in fact there is no difference between a metric tree and a marked metric tree.

The space $\mathcal{CV}_{0,s} = \mathcal{MG}_{0,s}$ of metric trees can be identified with a sphere complex in the following way. Let $W_{0,s}$ be the 3-sphere punctured by removing s disjoint balls. A 2-sphere embedded in $W_{0,s}$ is *trivial* if it either bounds a ball or is isotopic to a boundary sphere, otherwise it is *essential*. The dual graph to a sphere system S is now defined to have one vertex for each component of $W_{0,s} \setminus S$, one edge for each sphere in S, and one leaf for each boundary component B of $W_{0,s}$, where the leaf is attached to the vertex representing the component of $W_{0,s} \setminus S$ containing B. Every sphere system is complete (all complementary components are punctured balls, so are simply connected), and the map taking a sphere system to its dual graph is an isomorphism from the sphere complex $\mathcal{S}(W_{0,s})$ to $\mathcal{MG}_{0,s}$.

For $n \geq 1$, we let $W_{n,s}$ be a doubled handlebody of genus n with s disjoint 3-balls removed. Let $\mathcal{S}(W_{n,s})$ be the sphere complex of $W_{n,s}$ where, as in the case of trees, a sphere is essential if it does not bound a ball and is not isotopic to a boundary component. As in the case $s = 0$, we define a sphere system to be *complete* if all components of its complement are simply-connected; otherwise it is *incomplete*, or *at infinity*. Incomplete sphere systems form a subcomplex $\mathcal{S}_\infty(W_n)$. Hatcher's proof that $\mathcal{S}(W_n)$ is contractible generalizes to give the following theorem.

THEOREM 4.1 (Hatcher [?]). *For any $n \geq 1$ and $s \geq 0$ the sphere complex $\mathcal{S}(W_n)$ is contractible, as is the subspace $\mathcal{S}(W_n) \setminus \mathcal{S}_\infty(W_n)$.*

The dual graph G_S to a sphere system S is defined as in the case of trees, with one vertex for each component of $W_{n,s} \setminus S$, one edge for each sphere in S and one leaf for each boundary component of $W_{n,s}$. To get a marking on the dual graph of a complete system, first fix a homotopy equivalence c from $W_{n,s}$ to $R_{n,s}$ that sends each boundary sphere to a different leaf of $R_{n,s}$; in particular this has the effect of labeling the boundary components of $W_{n,s}$. If S is complete, there is a natural embedding $G_S \hookrightarrow W_{n,s}$ that is unique up to homotopy. The composition of this embedding with the homotopy equivalence $c\colon W_{n,s} \to R_{n,s}$ gives the marking of G_S.

Define $A_{n,s}$ to to be the group of homotopy classes of homeomorphisms of $W_{n,s}$ that fix the boundary, modulo Dehn twists in 2-spheres or, equivalently, the group of homotopy classes of homotopy equivalences of $R_{n,s}$ that fix the leaves. Then $A_{n,0}$ is isomorphic to $\mathrm{Out}(F_n)$ and $A_{n,1}$ is isomorphic to $\mathrm{Aut}(F_n)$. The action of $A_{n,s}$ on

$\mathcal{S}(W_{n,s})$ restricts to an action on complete systems, and the map sending a sphere system to its dual graph gives an isomorphism of the quotient with the moduli space of graphs $\mathcal{MG}_{n,s}$ defined in Section **??**.

We can therefore identify $\mathcal{CV}_{n,s}$ with the subspace of $\mathcal{S}(W_{n,s})$ consisting of complete sphere systems, and the *simplicial completion* $\mathcal{CV}^*_{n,s}$ with the entire sphere complex $\mathcal{S}(W_{n,s})$. In following sections we will pass freely back and forth between the graph and sphere models of $\mathcal{CV}_{n,s}$, using whichever is more convenient at the time.

5 Invariant spines. In Section **??** we pointed out that the dimension of $\mathcal{CV}_n$ is an upper bound for the virtual cohomological dimension VCD(Out(F_n)); similarly the dimension of $\mathcal{CV}_{n,s}$ is an upper bound for the VCD of $A_{n,s}$. However, $\mathcal{CV}_{n,s}$ has a lower-dimensional invariant deformation retract called the *spine* $K_{n,s}$, whose dimension is in fact is equal to the exact VCD. The spine $K_{n,s}$ has the added advantage that its quotient by the action of $A_{n,s}$ is compact. By the classical Švarc-Milnor lemma this implies that $K_{n,s}$ is quasi-isometric to the group $A_{n,s}$, so can be used to compute quasi-isometry invariants such as the Dehn function.

The spine is defined as a subcomplex of the barycentric subdivision of the simplicial complex $\mathcal{CV}^*_{n,s}$.

DEFINITION 5.1 (Spine). *The* spine K_n *of* $\mathcal{CV}_n$ *is the union of those simplices in the barycentric subdivision of* $\mathcal{CV}^*_{n,s}$ *all of whose vertices lie in* $\mathcal{CV}_n$.

Using the identification of $\mathcal{CV}^*_{n,s}$ with the sphere complex $\mathcal{S}(W_{n,s})$, the vertices of the barycentric subdividion $\mathcal{S}'(W_{n,s})$ are sphere systems and the simplices are chains $S_k \supset \ldots \supset S_0$ of sphere systems. The vertices of $K_{n,s}$ are complete sphere systems, and simplices in $K_{n,s}$ are chains such that the smallest system S_0 in the chain is complete. In terms of graphs, $K_{n.s}$ is the geometric realization of the partially ordered set of open simplices $\sigma(G,g)$ of $\mathcal{CV}_{n,s}$, ordered by face relations (i.e. forest collapse). The spine of $\mathcal{CV}^{red}_2$ is illustrated in Figure **??**. The deformation retraction $\mathcal{CV}_{n,s} \to K_{n,s}$, is performed as follows, where we will use the sphere model for $\mathcal{CV}_{n,s}$.

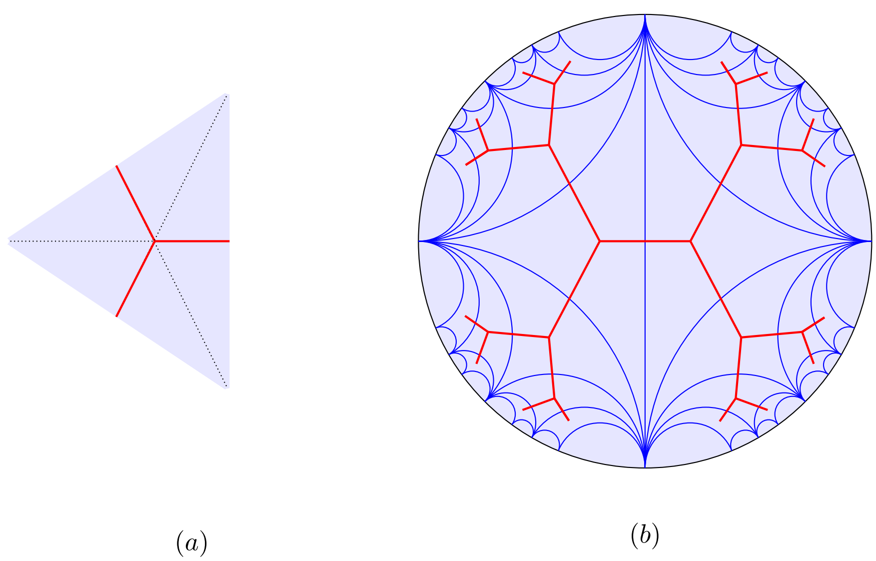

Figure 5.1: (a) Spine of a theta-graph simplex in $\mathcal{CV}_2$. (b) Spine of $\mathcal{CV}^{red}_2$

Let $\tau = S_k \supset \ldots \supset S_0$ be a simplex of $\mathcal{S}'(W_{n,s})$, and suppose τ is not already in K_n. Since τ is not in K_n, some S_i must be incomplete. Let m be the largest index with S_m incomplete. The intersection of τ with $\mathcal{CV}_n$ is obtained by removing the face $S_m \supset \ldots \supset S_0$, and the intersection with K_n is the opposite face $S_k \supset \ldots \supset S_{m+1}$. The deformation retraction retracts $\mathcal{CV}_n \cap \tau = S_k \supset \ldots \supset S_0$ linearly onto its opposite face $S_k \supset \ldots \supset S_{m+1}$.

The fact that K_n is a deformation retract of $\mathcal{CV}_n$ shows that K_n is contractible and the fact that a homeomorphism sends complete systems to complete systems shows that it is invariant under the action of $A_{n,s}$.

5.1 Virtual cohomological dimension. In the graph model, the group $A_{n,s}$ is the group of homotopy equivalences of $R_{n,s}$ that fix the labeled leaves. If $n = 0$ then $A_{0,s}$ is trivial, and if $n = s = 1$ then $A_{1,1} = \mathbb{Z}/2\mathbb{Z}$. For all other n, s the natural map $A_{n,s} \to A_{n,s-1}$ which allows a homotopy equivalence to move the last labeled leaf has kernel F_n. Iterating, we have a short exact sequence

$$1 \to (F_n)^s \to A_{n,s} \to A_{n,0} = \mathrm{Out}(F_n) \to 1. \tag{5.1}$$

Since both F_n and $\mathrm{Out}(F_n)$ are virtually torsion-free, this implies

$$\textsc{vcd}(A_{n,s}) \leq \textsc{vcd}((F_n)^s) + \textsc{vcd}(\mathrm{Out}(F_n)) = s + \textsc{vcd}(\mathrm{Out}(F_n)),$$

Thus to compute $\textsc{vcd}(A_{n,s})$ it suffices to compute $\textsc{vcd}(\mathrm{Out}(F_n))$ and then show that adding s also gives a lower bound for the $\textsc{vcd}$.

It is easy to compute that the dimension of K_n is equal to $2n-3$ so we can conclude that $\textsc{vcd}(\mathrm{Out}(F_n)) \leq 2n-3$. The fact that the $\textsc{vcd}$ of $K_{n,s}$ is at least $2n-3+s$ is shown by exhibiting a subgroup with $\textsc{vcd}$ equal to $2n-3+s$, since the $\textsc{vcd}$ of a group is at least as big as the $\textsc{vcd}$ of any subgroup. If $s=1$, $A_{n,1} = \mathrm{Aut}(F_n)$ has a free abelian subgroup of rank $2n-2$. Specifically, for $i \neq 1$ let $\rho_{i,1}$ denote the automorphism that multiplies the i-th generator of F_n on the right by x_1 and fixes all other generators. Similarly, let $\lambda_{i,1}$ be the automorphism that multiplies the i-th generator by x_1^{-1} on the left. The $\lambda_{i,1}$ and $\rho_{i,1}$ all commute, so generate a free abelian subgroup of $Aut(F_n)$ of rank $2n-2$. For $s=0$ note that the product of all $\lambda_{i,j}$ and $\rho_{i,j}$ is inner, so the image in $\mathrm{Out}(F_n)$ has rank only $2n-3$. For $s \geq 1$ the short exact sequence $1 \to (F_n)^{s-1} \to A_{n,s} \to A_{n,1} = \mathrm{Aut}(F_n) \to 1$ .splits, so $A_{n,s}$ contains a subgroup $(F_n)^{s-1} \rtimes \mathbb{Z}^{2n-2}$, which has the required $\textsc{vcd}$.

5.2 Cube complex structure of the spine. The simplices of the spine $K_{n,s}$ naturally group themselves into cubes, giving $K_{n,s}$ the structure of a *cube complex.* The point is that given two complete systems $S_k \supset S_0$ we can form a chain $S_k \supset S_{k-1} \supset S_1 \supset S_0$ by removing one sphere at a time to obtain a k-simplex of $K_{n,s}$. But in fact we can remove these spheres in any order, giving a set of 2^{k-1} simplices which form a cube. In terms of marked graphs, given a forest $\Phi \subset G$ with k edges, we can collapse these edges in any order to obtain a chain of faces of $\widehat{\sigma}(G,g)$ and hence a simplex of $K_{n,s}$. Thus there is a cube $c(G,g,\Phi)$, for every marked graph (G,g) and forest $\Phi \subset G$ (see Figure **??**). The orientation of the cube is determined by an ordering of the edges in Φ.

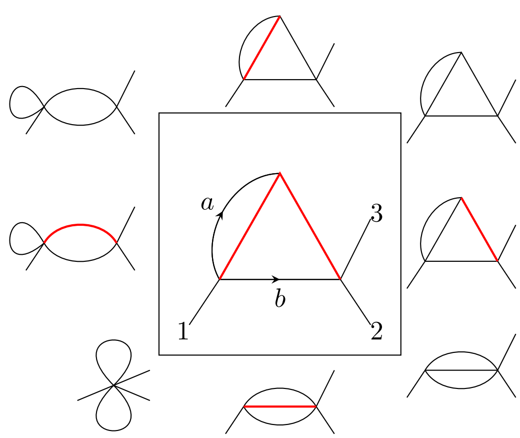

Figure 5.2: A 2-dimensional cube (G,Φ) in $K_{2,3}$. Here Φ is the red subgraph, $\pi_1(G) = F\langle a,b\rangle$ and the leaves are labeled $1,2,3$. The marking is indicated by arrows and labels on the edges in the complement of a maximal tree. The faces of the cube are obtained by either deleting or collapsing edges of Φ.

The group $A_{n,s}$ acts transitively on isomorphism classes of graphs of rank n with s labeled leaves, by simply changing the marking. The stabilizer of a cube $c(G,g,\Phi)$ can be identified with the group $\mathrm{Aut}(G,\Phi)$ of automorphisms of G that preserve Φ. In particular, the quotient of $K_{n,s}$ by the action of $A_{n,s}$ is compact, with one "cell" for each isomorphism class of pairs (G,Φ). This cell is the quotient of a marked cube by its stabilizer, which acts linearly on the boundary of the cube. Thus the quotient of the cube is a cone on a contractible space if some element of $\mathrm{Aut}(G,\Phi)$ induces an orientation-reversing symmetry (i.e. if it gives an odd permutation of the edges of Φ), and a cone on a rational homology sphere otherwise (see [**?**] for more details).

6 Kontsevich's graph complexes and Outer space. In this section we recall a theorem of Kontsevich which shows that the Chevalley-Eilenberg homology of three infinite-dimensional Lie algebras, whose definitions are inspired by considerations from physics, can be computed by various *graph complexes* [?, ?]. We explain how these chain complexes are related to the moduli spaces of graphs we have been studying, providing even more motivation for understanding the rational cohomology of these moduli spaces.

In classical mechanics the phase space of a mechanical system is a symplectic manifold (M, ω), and observables are smooth functions $M \to \mathbb{R}$, forming a commutative algebra $C^\infty(M)$. In quantum mechanics this algebra is replaced by an algebra of operators on a Hilbert space, which is no longer commutative. In [?] Konsevich proposed definitions of formal commutative, associative and Lie manifolds, replacing the geometric object M by free commutative, associative and Lie algebras analogous to $C^\infty(M)$. The Poisson bracket $\{f, g\}$ makes $C^\infty(M)$ into a Lie algebra, and makes sense on the free commutative algebra on $2n$ generators, thought of as polynomial functions on $\mathbb{R}^{2n}$ with symplectic basis $\{p_1, \ldots, p_n, q_1, \ldots, q_n\}$, making this into a Poisson algebra $\mathfrak{c}_n$.

Kontsevich defines analogs $\mathfrak{l}_n$ (respectively $\mathfrak{a}_n$) of this Poisson algebra for free non-unital Lie (respectively associative) algebras $\mathcal{L}_n$ (resp $\mathcal{A}_n$) on generators $\{p_1, \ldots, p_n, q_1, \ldots, q_n\}$, in terms of derivations. To understand the connection, recall that a smooth function $H\colon M \to \mathbb{R}$ gives a 1-form $dH\colon TM \to \mathbb{R}$. The symplectic form ω gives an isomorphism of TM with the cotangent bundle T^*M, so the 1-form dH corresponds to a unique vector field X_H on M. A vector field X, in turn, determines a derivation D_X of $C^\infty(M)$, where $D_X(f)$ is the directional derivative of f in the direction X. Finally, if $X = X_H$ for some H (if X is a *Hamiltonian vector field*) then the flow on M determined by X preserves the symplectic form ω.

Using this correspondence between functions and derivations, Kontsevich proposes that the appropriate analog of $\mathfrak{c}_n$ for $\mathfrak{l}_n$ should be the Lie algebra of derivations D of L_n such that $D(\sum_i [p_i, q_i]) = 0$, and for $\mathfrak{a}_n$ the derivations of A_n that kill $\sum_i p_i q_i - q_i p_i$. For each of $\mathfrak{h}_n = \mathfrak{c}_n, \mathfrak{l}_n$ or $\mathfrak{a}_n$, there are natural inclusions $\mathfrak{h}_n \hookrightarrow \mathfrak{h}_{n+1}$, and $\mathfrak{h}_\infty$ is defined to be the direct limit.

In [?, ?] Kontsevich showed that the Lie algebra homologies (with trivial rational coefficients) of these algebras have close connections with low-dimensional topology. The connection is established by showing that their homology can be computed combinatorially in terms of finite graphs. This uses the observation that each of these Lie algebras $\mathfrak{h}_\infty$ contains a copy of the Lie algebra $\mathfrak{sp}_\infty$ of the infinite symplectic group. The fact that $\mathfrak{sp}_\infty$ is simple and acts on $\mathfrak{h}_\infty$ by the adjoint action then implies that the Chevalley-Eilenberg chain complex $(\Lambda^*\mathfrak{h}_\infty, d)$ is quasi-isomorphic to (i.e. has the same homology as) the subcomplex of $\mathfrak{sp}_\infty$ invariants. Using classical invariant theory, these invariants can be identified with linear combinations of finite graphs whose vertices are decorated by generators of the (commutative, Lie or associative) cyclic operad. The differential can be interpreted as the operation of collapsing a non-loop edge and combining the generators on either end using the symplectic form. For details, see the exposition of Kontsevich's theorem in [?].

The Chevalley-Eilenberg homology of $\mathfrak{h}_\infty$ is a Hopf algebra $H_*(\mathfrak{h}_\infty)$, whose primitive part $PH_*(\mathfrak{h}_\infty)$ is computed by the subcomplex indexed by connected graphs. Using the connection with graph complexes Kontsevich proves

THEOREM 6.1 (Kontsevich [?]). *$PH_k(\mathfrak{h}_\infty)$ is the direct sum of $PH_k(\mathfrak{sp}_\infty)$ and*

- *(if $\mathfrak{h}_\infty = \mathfrak{l}_\infty$)* $\bigoplus_{n\geq 2} H^{2n-2-k}(\mathrm{Out}(F_n); \mathbb{Q})$
- *(if $\mathfrak{h}_\infty = \mathfrak{a}_\infty$)* $\bigoplus_{n>0, 2-2g-n<0} H^{4g-4-2n-k}(Mod(S_{g,n}) : \mathbb{Q})$
- *(if $\mathfrak{h}_\infty = \mathfrak{c}_\infty$)* $\bigoplus_{n\geq 2} (graph\, homology)_k^{(n)}$

In all versions of the graph complex the differential collapses a non-loop edge, so preserves the dimension of the first homology group of the graph, and the entire graph complex breaks up into a direct sum of subcomplexes consisting of graphs with the same first homology. If one further restricts to connected graphs, this is the rank of the fundamental group of the graph, giving rise to the direct sum decompositions in the statement.

In the commutative case $\mathfrak{c}_\infty$ there is no additional structure at the vertices of a graph, so the rank n piece of the primitive subcomplex simply has one generator for each graph with fundamental group F_n, i.e. there is one generator for each cell in the moduli space $\mathcal{MG}_n$. It is useful to think "upstairs", where there is one generator for each orbit of open simplices in $\mathcal{CV}_n$, i.e. one generator for each orbit in the relative chain complex $C_*(\mathcal{CV}_n^*, \mathcal{CV}_n^\infty)$. A simplex in $\mathcal{CV}_n^*$ is oriented by ordering its vertices, corresponding to ordering the edges of the graph. This coincides with Kontsevich's notion of "even" orientation on a graph, so the differential on $C_*(\mathcal{CV}_n^*, \mathcal{CV}_n^\infty)$ coincides with Kontsevich's differential on the even commutative graph complex, and Kontsevich's chain complex becomes

$$C_*(\mathcal{CV}_n^*, \mathcal{CV}_n^\infty) \otimes_{\mathbb{Z}[\mathrm{Out}(F_n)]} \mathbb{Z}.$$

Here $\mathrm{Out}(F_n)$ acts trivially on $\mathbb{Z}$. Kontsevich also defines an odd version of each graph complex, whose differential in this case coincides with the non-trivial action of $\mathrm{Out}(F_n)$ on $\mathbb{Z}$ given by the determinant map, i.e. $\alpha \in \mathrm{Out}(F_n)$ acts by the sign of the determinant of its image in $\mathrm{GL}(n, \mathbb{Z})$.

Of particular interest to us is the Lie graph complex. Here the vertices of a generator of the graph complex are decorated by (cyclic) bracket expressions of the p_i and q_i. These are subject to anti-symmetry and Jacobi relations, and it becomes less clear how to relate them to $\mathcal{CV}_n$. However, in [**?**] we showed that they can be understood in terms of cubes in the spine K_n of $\mathcal{CV}_n$. Recall that K_n is a cube complex, with one cube for each pair (G, Φ) consisting of an admissible graph of rank n and a subforest $\Phi \subset G$. Define the *forested graph complex* to be the chain complex with k-chains generated by pairs (G, Φ) where G is trivalent of rank n and $\Phi \subset G$ is a forest with k ordered edges. The differential removes edges of Φ one at a time, multiplying the i-th edge by $(-1)^i$ and adding.

THEOREM 6.2 ([**?**]). *Kontsevich's Lie graph complex is isomorphic to the forested graph complex.*

One can then reprove the Lie version of Kontsevich's theorem by showing that the forested graph complex is quasi-isomorphic to the cochain complex generated by the cellular chain complex of $K_n/\mathrm{Out}(F_n)$...which computes the rational cohomology of $\mathrm{Out}(F_n)$.

The associative graph complex is also related to $\mathcal{CV}_n$, but I will be even briefer here. The decoration at a vertex of an associative graph consists of a cyclic ordering of the half-edges adjacent to the vertex. A marked hyperbolic surface $S = S_{g,s}$ for $s > 0$ contains an embedded graph G_S consisting of points "equidistant from at least two punctures." The inclusion $G_S \hookrightarrow S$ induces an isomorphism on π_1, and the embedding gives a cyclic ordering of the edges of the graph at each vertex. Forgetting the cyclic ordering gives a marked graph, i.e. a point of $\mathcal{CV}_n$. In fact $\mathcal{CV}_n$ is a union of (overlapping) Teichmüller spaces of surfaces with fundamental group F_n, i.e. Teichmüller spaces $\mathcal{T}_{g,s}$ with $s > 0$ and $2g + s - 1 = n$. The mapping class group $Mod(S_{g,s})$ embeds in $\mathrm{Out}(F_n)$ as the subgroup that preserves the conjugacy classes corresponding to loops around the punctures. There are usually many copies of a given $\mathcal{T}_{g,s}$ in $\mathcal{CV}_n$. They are permuted by the action of $\mathrm{Out}(F_n)$, and the stabilizer of a single copy is isomorphic to $Mod(S_{g,s})$.

7 Finding cohomology classes

7.1 Cohomology of $\mathcal{MG}_{n,s}$ for n=0 and n=1. Recall from Section **??** that for $s \geq 4$, $\mathcal{MG}_{0,s}$ is homotopy equivalent to a wedge of $(s-2)!$ spheres of dimension $(s-4)$. Therefore the reduced homology of $\mathcal{MG}_{0,s}$ vanishes except in dimension $(s-4)$, where it has dimension $(s-2)!$.

In Section **??** we also learned that for $s \geq 1$, $\mathcal{MG}_{1,s}^{red}$ (which is a deformation retract of $\mathcal{MG}_{1,s}$ so has the same cohomology) is the quotient of an $(s-1)$-torus by an order 2 involution. The cohomology of the torus is an exterior algebra on $(s-1)$ one-dimensional generators x_i, and the order 2 involution sends each generator to its negative, so kills all of the odd-dimensional cohomology. In particular, if s is odd, there is a single class in $H^{s-1}(\mathcal{MG}_{1,s}) = \Lambda^{s-1}(x_1, \ldots, x_{s-1})$, which we will denote α_s.

7.2 Cohomology computations using the spine. For $n \geq 2$ things get more complicated. Since $K_{n,s}$ is an equivariant deformation retract of $\mathcal{CV}_{n,s}$, the quotient $K_{n,s}/A_{n,s}$ is a deformation retract of $\mathcal{MG}_{n,s}$; in particular it has the same cohomology. The cubical structure of $K_{n,s}$ gives rise to a relatively small chain complex for the quotient, with one generator for each orbit of cubes with no orientation-reversing symmetry, i.e. each pair (G, Φ) which has no automorphism that induces an odd permutation of the edges of Φ (see [**?**]).

For small values of n the computation can then be done by hand, and for slightly larger n it can be done by computer. This was done for $n \leq 6$ and $s = 1$ in [**?**]. The result is that the reduced homology is trivial in all cases except for $n = 4$, where there is one non-trivial class in $H_4(\mathcal{MG}_{4,1})$. In particular this gives a non-trivial rational cohomology class for $\mathrm{Aut}(F_4)$, since $H_4(\mathcal{MG}_{4,1}; \mathbb{Q})) = H_4(\mathrm{Aut}(F_4); \mathbb{Q})$.

Generators of the cellular chain complex for $K_4/\mathrm{Out}(F_4)$ are the images $\overline{c}(G, \Phi)$ of marked cubes $c(G, g, \Phi)$ in K_4. The non-trivial class in $H_4(\mathcal{MG}_{4,1})$ can be represented by a cycle formed by the sum of all generators $\overline{c}(G, \Phi)$ such that G is the union of two disjoint circles connected by three edges, and Φ is a 4-edge forest with two edges in each circle (see Figure **??**). The image of this cycle in $H_4(\mathcal{MG}_{4,0})$ formed by forgetting the basepoint is also non-trivial, giving a non-trivial class in $H_4(\mathrm{Out}(F_4); \mathbb{Q})$.

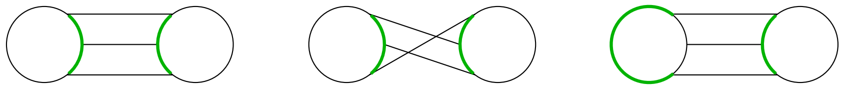

Figure 7.1: Pairs (G, Φ) determining three of the cube-quotients that contribute to the non-trivial homology class in $H_4(\mathrm{Out}(F_4); \mathbb{Q})$.

7.3 Morita cycles. There is a natural generalization of the cycle in $H_4(\mathcal{MG}_{4,0})$ constructed as follows. Consider all graphs G consisting of two circles with $2k+1$ edges joining them, and all forests Φ with $4k$ edges contained in the circles. Each such pair (G, Φ) determines a generator of the cellular chain complex for $K_{2k+2}/\mathrm{Out}(F_{2k+2})$, and the sum of all these generators is a cycle. We remark that the case $k = 1$ illustrated in Figure **??** is deceptive because the graphs have too much symmetry; in general different ways of joining the circles by edges will result in non-isomorphic graphs, which all contribute to the cycle. It is known that these cycles are non-trivial in homology for $k = 1, 2$ and 3, and it is conjectured that they are all non-trivial. These are the graphical representations of the *Morita cycles*, which were discovered by S. Morita using the connection discovered by Kontsevich between the cohomology of the infinite-dimensional Lie algebra $\mathfrak{l}_\infty$ and the cohomology of $\mathrm{Out}(F_n)$ [**?**] described in Section **??**.

7.4 Assembly maps. The graphical picture of Morita cycles suggests a procedure for producing new potential homology classes from known classes. The idea is that a graph can be assembled from a collection of smaller graphs with leaves by gluing the leaves together in pairs. For instance the graphs involved in the Morita cycles are assembled from two genus one graphs, each with $2k+1$ leaves, by gluing their leaves together in pairs. If $G \in \mathcal{MG}_{N,S}$ is assembled from graphs $G_i \in \mathcal{MG}_{n_i,s_i}$, there is an obvious map of groups

$$A_{n_1,s_1} \times \ldots \times A_{n_k,s_k} \to A_{N,S},$$

since elements of $A_{n,s}$ are given by homotopy equivalences that fix the leaves. If we take homology with trivial rational coefficients, the Künneth formula gives a map

$$H_*(A_{n_1,s_1}; \mathbb{Q}) \otimes \ldots \otimes H_*(A_{n_k,s_k}; \mathbb{Q}) \to H_*(A_{N,S}, \mathbb{Q}).$$

Therefore if we have classes in $H_*(A_{n_1,s_1}; \mathbb{Q})$ we can use them to find cycles in $H_*(A_{N,S}, \mathbb{Q})$. For example the Morita class $\mu_{2k+1} \in H_{4k}(\mathrm{Out}(F_{2k+2}))$ is the image of $\alpha_{2k+1} \otimes \alpha_{2k+1}$, where α_{2k+1} is the generator of $H_{2k}(\mathcal{MG}_{1,2k+1})$ described in Section **??**. A class in $H_i(\mathrm{Out}(F_N); \mathbb{Q})$ that is assembled from the α_i's is called a *generalized Morita cycle*. The first of these which is not an original Morita cycle is in $H_8(\mathrm{Out}(F_7)$; it is assembled from 4 copies of α_3 arranged in a tetrahedral pattern.

Note that the symmetric group Σ_s acts on $\mathcal{MG}_{n,s}$ by permuting the labels on the leaves of the graphs. This gives an action on $H_*(\mathcal{MG}_{n,s})$, and in [**?**] elementary representation theory of symmetric groups was used to analyze assembly maps and find consequences for the cohomology of $\mathcal{MG}_{n,s}$. In particular, the cohomology of $\mathcal{MG}_{2,s}$ with trivial rational coefficients was computed for all s, including its decomposition into irreducible representations of Σ_s. This used the fact that the map $\mathrm{Out}(F_2) \to GL(2, \mathbb{Z})$ induced by abelianization is an isomorphism, and the cohomology of $\mathrm{GL}(2, \mathbb{Z})$ can be expressed in terms of modular forms, by the classical Eichler-Shimura theorem.

8 Homology stability. The finite number of low-rank computations that we did in the last section can be bootstrapped to give an infinite number of computations using the phenomenon of homological stabiity. A sequence of groups or spaces X_n and maps $f_n \colon X_n \to X_{n+1}$ is said to satisfy *homological stability* if for each $i \geq 0$ there is some $N = N(i)$ such that f_n induces an isomorphism $H_i(X_n) \to H_i(X_{n+1})$ for all $n \geq N$. If $\{X_n, f_n\}$ satisfies homological stability, then the i-th *stable homology group* is the group $H_i(X_N)$, which is sometimes denoted $H_i(X_\infty)$. If N depends linearly on i, i.e. $N(i) = mi + b$, then the sequence is said to be homologically stable with *slope m*. A spectral sequence argument originally due to Quillen has been used to prove homological stability for many natural sequences of groups, including general linear groups, symplectic groups and surface mapping class groups, using as input the natural actions of these groups on various simplicial complexes. Quillen's argument has been vastly generalized in recent years, notably by Randall-Williams and Wahl [**?**].

There are natural inclusions $f_n \colon \mathrm{Aut}(F_n) \to \mathrm{Aut}(F_{n+1})$ given by sending an automorphism of F_n to the automorphism of F_{n+1} that has the same effect on the first n generators but fixes the last generator. These maps satisfy homology stability. The first complete proof of this was given in [?]. The stabilization range for $\mathrm{Aut}(F_n)$ was shown to be at most linear with slope 2 for integral coefficients, and slope 3/2 for rational coefficients. In [?] the rational slope for $\mathrm{Aut}(F_n)$ was improved to 5/4, and it is reasonable to conjecture that the slope should actually be 1.

There are also natural inclusions $\mathcal{MG}_{n,1} \to \mathcal{MG}_{n+1,1}$. Here we think of a graph in $\mathcal{MG}_{n,1}$ as having a basepoint at the vertex where the unique leaf is attached; the map is then defined by adding a loop at the basepoint. There is a nice proof of rational homology stability for $\mathrm{Aut}(F_n)$ that relies on the fact that $H_*(\mathrm{Aut}(F_n);\mathbb{Q}) \cong H_*(\mathcal{MG}_{n,1};\mathbb{Q})$, so in order to prove stability for $\mathrm{Aut}(F_n)$ it suffices to show that the inclusion $\mathcal{MG}_{n,1} \to \mathcal{MG}_{n+1,1}$ induces an isomorphism on i-th homology for n sufficiently large. To do this, we first look at upstairs at the map $\mathcal{CV}_n \to \mathcal{CV}_{n+1}$ that sends a marked graph (G, g') to (G', g'), where as before G' is obtained from G by adding a loop at the basepoint, and g' extends g to R_{n+1} by sending the last petal of R_{n+1} homeomorphically to the last loop.

We now define the *degree* of a basepointed graph G as $\sum_{v\neq b} |v| - 2$, where b is the basepoint and $|v|$ denotes the valence of v. For example, the degree of a single-stemmed rose is 0, and the degree of single-stemmed rose wedged with a theta graph is 1. Define $\mathcal{CV}_{n,1}(d)$ to be the subspace of $\mathcal{CV}_{n,1}$ consisting of marked graphs (G, g) such that G has degree at most d. The main theorem in [?] is the following:

THEOREM 8.1. *The degree d subspace $\mathcal{CV}_{n,1}(d) \subset \mathcal{CV}_{n,1}$ is $(d-1)$-connected.*

The proof in [?] uses a graphical version of Cerf theory. A more direct combinatorial proof was given later in [?].

Since the action of $\mathrm{Aut}(F_n)$ changes only the marking, the subspace $\mathcal{CV}_{n,1}(d)$ is invariant under the action. Since the action is proper and $\mathcal{CV}_{n,1}(d)$ is $(d-1)$-connected, the i-th homology of the quotient $\mathcal{MG}_{n,1}(d)$ is equal to the i-th homology of $\mathrm{Aut}(F_n)$ with rational coefficients for $i < d$.

An easy Euler characteristic argument now shows

LEMMA 8.2. *If G has rank $n \geq 2d$ then any graph of degree d must have a loop at the basepoint.*

Since the only difference between $\mathcal{MG}_{n,1}(d)$ and $\mathcal{MG}_{n+1}(d)$ for $n \geq 2d$ is the number of loops at the basepoint, these spaces are homeomorphic, so have the same homology, i.e. for $i < d \leq n/2$

$$H_i(\mathrm{Aut}(F_n);\mathbb{Q}) = H_i(\mathcal{MG}_{n,1}(d);\mathbb{Q}) = H_i(\mathcal{MG}_{n+1,1}(d);\mathbb{Q}) = H_i(\mathrm{Aut}(F_{n+1};\mathbb{Q}).$$

Note we also have $H_i(\mathcal{MG}_{n,1}) = H_i(\mathcal{MG}_{n+1,1})$ for $i < n/2$. A more subtle argument shows that for $n \geq 3d/2$ the inclusion $\mathcal{MG}_{n,1}(d) \to \mathcal{MG}_{n+1,1}(d)$ is at least a homotopy equivalence, so we still have an isomorphism on homology. To get slope $5d/4$ we show that some generators of the chain complex for $\mathcal{MG}_{n,1}$ cancel in pairs.

To get homology stability with integer coefficients in [?], a form of Quillen's method was applied, using the action of $\mathrm{Aut}(F_n)$ on the simplicial closure $\mathcal{CV}^*_{n,1}$. Later Bestvina gave a variation of the geometric argument above to prove integral homology stability for $\mathrm{Aut}(F_n)$, using the added ingredient of homology stability for sequences of signed symmetric groups [?].

In [?] it was shown that the groups $A_{n,s}$ for $s \geq 1$ also satisfy homology stability as $n \to \infty$. This was shown using Quillen's method, using natural actions of $A_{n,s}$ on various highly-connected simplicial complexes. That paper also purported to show that the map $\mathrm{Aut}(F_n) \to \mathrm{Out}(F_n)$ and the maps $A_{n,s} \to A_{n,s+1}$ induce an isomorphisms on i-th homology for n sufficiently large, but an error was pointed out by N. Wahl. The error was corrected in [?], showing that $H_i(A_{n,s};\mathbb{Q})$ is independent of both n and s for all n sufficiently large. In particular, some of the calculations we have done for small values of n determine low-dimensional homology for all n. The few nontrivial classes we found are not in the stable range, however, and we could only conclude that $H_i(\mathrm{Aut}(F_\infty)) = 0$ for $i = 1, 2$ and 3. In particular the Morita classes $\mu_k \in H_{4k}(\mathrm{Aut}_{2k+2};\mathbb{Q})$ are not in the known stable range. In [?] we showed they cannot be pushed into the stable range, either:

THEOREM 8.3. *The Morita classes $\mu_k \in H_{4k}(\mathrm{Aut}_{2k+2};\mathbb{Q})$ vanish after one stabilization, ie the image of μ_k in $H_{4k}(\mathrm{Aut}_{2k+3};\mathbb{Q})$ is trivial.*

In [?] we conjectured that in fact the groups $\mathrm{Aut}(F_n)$ have no stable rational homology. This was confirmed in a major breakthrough by Søren Galatius, who proved the following theorem in 2006.

THEOREM 8.4 (Galatius [?]). *The inclusions $\Sigma_n \hookrightarrow \mathrm{Aut}(F_n)$ induce an isomorphism on stable homology. In particular, the stable homology with rational coefficients is trivial.*

9 Unstable homology. Although $\mathrm{Out}(F_n)$ has no stable rational homology, we know that unstable homology exists by our low-dimensional calculations. However it becomes extremely challenging to do exact homology calculations when the rank gets large.

The Euler characteristic is easier to compute, since by the magic of Euler characteristic it can be calculated on the chain complex level, i.e. it only involves counting the number of cells of each dimension. Using the chain complex coming from the cubical structure of the reduced spine of $\mathcal{CV}_{n,s}$, this means counting the number of pairs (G, Φ) where

- G is an admissible graph with s labeled leaves
- Φ is a forest in G consisting of internal edges and
- there are no automorphisms of G that induce an odd permutation of the edges of Φ.

For very low values of n this can be done by hand. For $n = 8, 9, 10$ and 11 Morita, Sakasai and Suzuki used a different method to compute the Euler characteristic based on Kontsevich's theorem, which involved counting symplectic invariants of the Chevalley-Eilenberg chain complex for $\mathfrak{l}_n$ [**?**]. Their computations produced the startling result that for $n = 9, 10$ and 11 the Euler characteristic is negative, and seems to be growing quite fast: for $n = 11$ it is equal to -1202. This means that there must be odd-dimensional homology in these ranks, and lots of it. Recall that all of the homology produced by the generalized Morita construction is even-dimensional.

The first (and up to now the only published) non-trivial odd-dimensional homology class that we know of was found by L. Bartholdi, who used extensive computer power to show that $H_{11}(\mathcal{MG}_{7,0}; \mathbb{Q}) = H_{11}(\mathrm{Out}(F_7); \mathbb{Q}) = \mathbb{Q}$ [**?**]. The appearance of this class was quite a surprise because the Euler characteristic calculation did not detect it. It turned out that there is another class, in $H_8(\mathrm{Out}(F_7); \mathbb{Q})$, that cancels with the class in $H_{11}(\mathrm{Out}(F_7); \mathbb{Q})$ when you compute the Euler characteristic. Bartholdi told me in a recent conversation that he can now show that the second class is equal to the generalized Morita class described in Section **??**.

10 The rational Euler characteristic of the groups $A_{n,s}$. If a group Γ has torsion, then any $K(\Gamma, 1)$ space has cells in infinitely many dimensions, so its Euler characteristic cannot be defined on the chain level. If Γ contains a torsion-free subgroup H of finite index, however, there may actually be a finite $K(H, 1)$. This occurs, for instance, if Γ acts properly and cocompactly on a contractible space X, since any torsion-free subgroup H of finite index acts freely on X, so the quotient X/H is a compact $K(H, 1)$ space. The *rational Euler characteristic* $\chi(\Gamma)$ is then defined to be

$$\chi(\Gamma) = \frac{\chi(X/H)}{[\Gamma : H]},$$

where $[\Gamma : H]$ is the index of H in Γ. It is not difficult to see that this number is independent of the choice of H. For this and other basic information about $\chi(\Gamma)$ we refer to Brown's book on cohomology of groups [**?**].

Warning: in spite of its name, the rational Euler characteristic of a group Γ is not the same as the alternating sum of the dimensions of the $H_i(\Gamma, \mathbb{Q})$, even when that number is finite (in fact $\chi(\Gamma)$ is often not an integer). The rational Euler characteristic has computational advantages over the alternating sum of the dimensions. For example it behaves well with respect to group extensions, mimicking the behavior of the Euler characteristic of topological spaces with respect to fibrations. Namely, if $1 \to A \to B \to C \to 1$ is a short exact sequence of groups and the Euler characteristic is defined for each of A, B and C, then $\chi(B) = \chi(A) + \chi(C)$. The rational Euler characteristic is also related to number theory; see [**?**] for more information about this connection. On the other hand, it does not necessarily detect rational cohomology; for instance $\chi(\mathrm{GL}(n, \mathbb{Z})) = 0$ for all $n \geq 2$, even though $\mathrm{GL}(n, \mathbb{Z})$ has very interesting (and mysterious) rational cohomology.

Recall that under the map $\mathrm{Out}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$, any torsion-free finite-index subgroup of $\mathrm{GL}(n, \mathbb{Z})$ pulls back to a torsion-free finite-index subgroup of $\mathrm{Out}(F_n)$ which acts cocompactly on the spine K_n. Therefore the rational Euler characteristic of $\mathrm{Out}(F_n)$ is defined. Similarly, the rational Euler characteristics of the groups $A_{n,s}$ are also defined, so the group extensions

$$1 \to (F_n)^{s-1} \to A_{n,s} \to \mathrm{Out}(F_n) \to 1$$

reduce the problem of computing $\chi(A_{n,s})$ to that of computing $\chi(\mathrm{Out}(F_n))$.

10.1 The Euler characteristic of $\mathrm{Out}(F_n)$. To compute $\chi(\mathrm{Out}(F_n))$ we use the action of $\mathrm{Out}(F_n)$ on the spine K_n of $\mathcal{CV}_n$. If Γ acts cellularly and cocompactly on a contractible CW-complex X with finite stabilizers, then

$$\chi(\Gamma) = \sum_{[\sigma]} \frac{(-1)^{\dim(\sigma)}}{|stab(\sigma)|}$$

where the sum is over the orbits $[\sigma]$ of cells σ, and $|stab(\sigma)|$ is the order of the stabilizer of σ (again, see [?]). For the action of $\mathrm{Out}(F_n)$ on the spine K_n this translates to

$$\chi(\mathrm{Out}(F_n)) = \sum_{(G,\Phi)} \frac{(-1)^{e(\Phi)}}{|\mathrm{Aut}(G,\Phi)|} = \sum_G \frac{1}{|\mathrm{Aut}(G)|} \sum_{\Phi \subset G} (-1)^{e(\Phi)} \tag{10.1}$$

where (G, Φ) runs over all isomorphism classes of pairs consisting of a connected graph G with fundamental group F_n and forest Φ, $e(\Phi)$ is the number of edges in Φ, and $\mathrm{Aut}(G, \Phi)$ is the (finite) group of automorphisms of G that preserve Φ.

In [?] Equation (??) was used to produce a generating function for $\chi(\mathrm{Out}(F_n))$, i.e. a formal power series $T(z)$, where the coefficient of z^n is $\chi_n = \chi(\mathrm{Out}(F_{n+1}))$. (The shift in index is convenient for manipulating the power series.) The terms in the power series were defined recursively, and were computed for $n \le 10$. Since the answers were all negative and seemed to be growing very quickly we conjectured that they were always negative and were growing more than exponentially fast. Further computation by D. Zagier (private communication in 1988) was carried out for $n \le 100$, and supported that conjecture, but it was not actually resolved until 2022, when the following theorem was proved.

THEOREM 10.1 ([?]). *For $n \ge 2$ the rational Euler characteristic $\chi(\mathrm{Out}(F_n))$ is negative. Its asymptotic growth rate is given by*

$$\chi(\mathrm{Out}(F_n)) \sim -\frac{\Gamma(n-\frac{3}{2})}{\sqrt{2\pi}\log^2 n} \sim -\left(\frac{n}{e}\right)^n \frac{1}{(n\log n)^2}.$$

Here $f(n) \sim g(n)$ means $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 1$, and the second asymptotic equivalence can be shown by using Sterling's asymptotic formula for the Γ function.

To prove this theorem, we first show that the function $\sqrt{2\pi}n^n e^{-n}$ has an asymptotic expansion in the scale $\phi_k = (-1)^k\Gamma(n+\frac{1}{2}-k)$ whose coefficients are closely related to $\chi(\mathrm{Out}(F_n))$. Specifically,

$$\sqrt{2\pi}n^n e^{-n} \sim \sum_{k\ge1} \widehat{\chi}_k \phi_k \quad \text{as } n \to \infty \tag{10.2}$$

where $\widehat{\chi}_k$ is the coefficient of z^k in the formal power series $\exp(T(z)) = \exp(\sum_n \chi(\mathrm{Out}(F_{n+1}))z^n)$. Exponentiating the generating function $T(z)$ is a standard trick in generating functions, which eliminates the need to restrict to connected graphs. We can then recover $T(z)$ from $\widehat{T}(z)$ by taking the formal logarithm.

We next produce another asymptotic expansion

$$\sqrt{2\pi}n^n e^{-n} \sim \sum_{k\ge1} \nu_k \phi_k \quad \text{as } n \to \infty \tag{10.3}$$

whose coefficients can be expressed in terms of the coefficients in an asymptotic expansion of the Lambert W-function. The Lambert W-function is a classical function with a long history, implicitly defined by the equation $W(x)e^{W(x)} = x$. Since two asymptotic expansions of the same function in the same scale must be equal, we can use results of Volker about the coefficients of the expansion of $W(x)$ to show the $\widehat{\chi}_k$ are negative, and singularity analysis to determine the asymptotic growth rate of the $\widehat{\chi}_k$. We then observe that this implies the same statements about χ_k.

To produce the asymptotic expansion (??) we again started from Equation (??). We can write $\widehat{T}(z)$ more compactly by setting $\tau(G) = \sum_{\Phi\subseteq G}(-1)^{e(\Phi)}$ and $|G| = e(G) - v(V)$; then

$$\widehat{T}(z) = \sum_n \widehat{\chi}_n z^n = \sum_n \left(\sum_{|G|=n} \frac{\tau(G)}{|\mathrm{Aut}(G)|} \right) z^n = \sum_G \frac{\tau(G)}{|\mathrm{Aut}(G)|} z^{|G|}. \tag{10.4}$$

where we are summing over all admissible graphs G, not necessarily connected.

We then proved the following combinatorial proposition.

PROPOSITION 10.2. *If G contains a cycle, then $\sum_{\gamma\subseteq G}\tau(\gamma)(-1)^{e(G/\!/\gamma)}=0$, where $G/\!/\gamma$ is the graph obtained from G by collapsing each component of γ to a point.*

This gives the following (trivial) equation in the ring of formal power series, which is strongly reminiscent of the formula (**??**) for $\widehat{T}(z)$:

$$1=\sum_{G}\sum_{\gamma\subseteq G}\frac{\tau(\gamma)(-1)^{e(G/\!/\gamma)}}{|\mathrm{Aut}(G)|}z^{|G|},$$

where the only non-trivial term is contributed by the empty graph. We think of this equation as an asymptotic formula for the constant function 1 in the scale n^{-k} as $n\to\infty$, change scales to the scale $\psi_k=\phi_k/\sqrt{2\pi}n^ne^{-n}$ then multiply by $\sqrt{2\pi}n^ne^{-n}$ to get the asymptotic expansion **??**.

The intuition behind the formula in Proposition **??** comes from the Connes-Kreimer Hopf algebra of core graphs, which was invented for understanding renormalization in quantum field theory calculations. Here a core graph is a graph with no separating edges, also called a 1-particle irreducible (1-PI) graph. The function τ is multiplicative on disjoint unions, so is a character on this Hopf algebra. The function $\sigma(G)=(-1)^{e(G)}$ is another such character. The set of characters forms a group, with product $\tau\star\sigma(G)=\sum_{\gamma\subseteq G}\tau(\gamma)\sigma(G/\!/\gamma)$, where γ runs over all subgraphs of G, including the trivial subgraph and the whole subgraph. The characters τ and σ are inverses in the character group of this Hopf algebra, which is a reformulation of Proposition **??** in the case of 1-PI graphs.

The motivation for using the scale ϕ_k in our asymptotic expansion comes from the relation between the Γ function and graph enumeration studied in [**?**]. There the connection is drawn between asymptotic expansions and graphical enumeration by using Laplace's method to find an asymptotic expansion of

$$\psi_k(n)=\frac{\Gamma(n-k+\frac{1}{2})}{\sqrt{2\pi}e^{-n}n^n}$$

in the scale $\{n^{-k}\}$ as $n\to\infty$.

10.2 The Euler characteristic of $\mathcal{MG}_n$. Although the rational Euler characteristic of $A_{n,s}$ is easier to compute than the actual Euler characteristic, it does not necessarily give information about the topology of the moduli space $\mathcal{MG}_{n,s}$; for instance $\chi(\mathrm{GL}(n,\mathbb{Z}))=0$ in spite of the fact that the quotient of the symmetric space by $\mathrm{GL}(n,\mathbb{Z})$ has a large amount of very interesting cohomology.

In [**?**] we proved the following theorem for $\mathcal{MG}_n$.

THEOREM 10.3. *For $n\geq 2$ the asymptotic growth rate of the Euler characteristic $\chi(\mathcal{MG}_n)$ is is given by*

$$\chi(\mathcal{MG}_n)\sim -e^{-\frac{1}{4}}\left(\frac{n}{e}\right)^n\frac{1}{(n\log n)^2}.$$

Since the rational cohomology of $\mathrm{Out}(F_n)$ is the same as that of $\mathcal{MG}_n$, we conclude that there is a vast amount of rational cohomology in $\mathrm{Out}(F_n)$ as n increases, mostly in odd dimensions, in spite of the fact that there is no stable rational homology.

Notice that the asymptotic growth rates of $\chi(\mathcal{MG}_n)$ and $\chi(\mathrm{Out}(F_n))$ differ by $e^{-\frac{1}{4}}$. Thus the difference in the amount of cohomology grows as n gets larger, answering a question of Morita, Sakasai and Suzuki in [**?**] whether the two numbers are asymptotically the same. This contrasts with the behavior of Euler characteristics for surface mapping class groups, where the rational Euler characteristic and alternating sum of the Betti numbers are in fact the same asymptotically.

To prove this theorem we used a different approach approach to finding the generating function for $\chi(\mathcal{MG}_n)$, which is similar to that used by both Penner and Kontsevich to calculate the Euler characteristics of surface mapping class groups (a calculation which was done originally by Harer and Zagier using a more combinatorial method). The idea comes more directly from quantum field theory, as we explain in the next subsection. In the rest of this subsection we give the setup for applying the idea in the case that interests us.

For a compact CW-complex X, the Euler characteristic can be computed either by counting cells or using the dimensions of the homology or chain groups, i.e.

$$\chi(X)=\sum_{\sigma}(-1)^{\dim(\sigma)}=\sum_{i}(-1)^{i}\mathrm{rank}(H^{i}(X;Z))$$
$$=\sum_{i}(-1)^{i}\dim H^{i}(X;\mathbb{Q})=\sum_{i}(-1)^{i}\dim C^{i}(X,\mathbb{Q}),$$

where $C^i(X,\mathbb{Q})$ is any chain complex that computes the rational homology of X. For the spine $K_{n,s}$ we have such a chain complex (see Section**??**), with one generator for each pair (G,Φ) with no odd automorphisms. We call such a pair *even*. Therefore we have the simple-looking formula

$$\chi(\mathcal{MG}_{n,s})=\sum_{(G,\Phi)}(-1)^{e(\Phi)},$$

where the sum is over all even pairs (G,Φ), where G is admissible of rank n with s labeled leaves.

We will focus on the case $s=0$. We want to produce a generating function for the numbers $\chi(\mathcal{MG}_n)$. In order to analyze the asymptotics, we want to relate it to the generating function $T(z)$ for $\chi(\mathrm{Out}(F_n))$ that we produced in [**?**], since we already know the asymptotics of that. To do this we use a counting lemma that makes our formula for $\chi(\mathcal{MG}_n)$ more complicated, but easier to relate to $T(z)$:

$$\chi(\mathcal{MG}_n)=\sum_{(G,\Phi)}\frac{1}{|\mathrm{Aut}(G,\Phi)|}\sum_{\alpha\in\mathrm{Aut}(G,\Phi)}(-1)^{e_\alpha(\Phi)}, \tag{10.5}$$

where now we are summing over all pairs (G,Φ) with G admissible of rank n and $(-1)^{e_\alpha}(\Phi)$ is the number of edge-orbits of α in Φ.

10.3 Counting graphs using ideas from quantum field theory. By Equation (**??**) the problem of computing the Euler characteristic of $\mathcal{MG}_n$ reduces to that of counting isomorphism classes of pairs consisting of a graph and forest weighted by the automorphism group of the pair, with appropriate sign.

The way to relate such counting problems to quantum field theory is well-known to physicists, but perhaps not to most mathematicians. We illustrate it by the much simpler problem of counting (not necessarily connected) graphs with no leaves or bivalent vertices, weighted by their automorphism groups. Define $\mathcal{G}$ to be the set of all such graphs, so we are interested in the numbers $\gamma(n)=\sum_G\frac{1}{|\mathrm{Aut}(G)|}$ where G runs over all $G\in\mathcal{G}$ of rank n. We want to produce a formal power series $F(\hbar)=\sum_G\gamma(n)\hbar^n$. By standard techniques in combinatorics, if you can find a generating function that counts all such graphs, you can take the formal logarithm to get a generating function that only counts connected graphs.

To fix notation, let $v(G)$ be the number of vertices of G, $e(G)$ the number of edges and $|G|=e(G)-v(G)$. Decompose $\mathcal{G}$ into the disjoint union $\sqcup_{e,v}\mathcal{G}_{e,v}$, where $\mathcal{G}_{e,v}$ is the (finite) subset of $\mathcal{G}$ consisting of graphs with e edges and v vertices. Then

$$F(\hbar)=\sum_{G\in\mathcal{G}}\frac{1}{|Aut(G)|}\hbar^{|G|}=\sum_{e,v}\sum_{G\in\mathcal{G}_{v,e}}\frac{1}{|Aut(G)|}\hbar^{e-v}.$$

If we fix an identification of the half-edges of G with $\{1,\dots,2e\}$ we get a *labeled graph*; we denote the set of labeled admissible graphs by $\ell\mathcal{G}$. The symmetric group Σ_{2e} acts on the set of labeled graphs with e edges by permuting the labels. The orbit of a labeled graph is the set of labeled graphs isomorphic to the underlying combinatorial graph G, and the stabilizer is $\mathrm{Aut}(G)$. By the orbit-stabilizer theorem, we get

$$\sum_{G\in\mathcal{G}_{v,e}}\frac{1}{|Aut(G)|}=\frac{|\ell\mathcal{G}_{v,e}|}{(2e)!},$$

so we now want to count labeled graphs.

The problem of counting labeled graphs is simpler than that of counting graphs. To count labeled graphs with $k=2e$ half-edges, we first count corollas with at least 3 leaves, then exponentiate to get collections of corollas, then count ways to match their edges to get labeled graphs.

It is very easy to count corollas with $k \geq 3$ leaves, weighted by their automorphism groups: there is only one isomorphism type for each k, with automorphism group the full symmetric group Σ_k. Putting this information into a generating function gives $\sum_{k\geq 3} \frac{1}{k!}x^k$, which is equal to $e^x - 1 - x - \frac{x^2}{2}$. Exponentiating $y(e^x - 1 - x - \frac{x^2}{2})$ gives a two-variable generating function, where the coefficient of $y^v \frac{x^{2e}}{(2e)!}$ is the number of collections of v corollas with a total of $2e$ leaves. There are $(2e-1)!!$ ways to match $2e$ leaves in pairs so, using the notation $[X]P$ to mean the coefficient of the monomial X in the formal power series P, we have

$$\begin{aligned} |\ell\mathcal{G}_{v,e}| &= (2e-1)!![y^v \frac{x^{2e}}{(2e)!}] \exp(y(e^x - 1 - x - \frac{x^2}{2})) \\ &= (2e-1)!!(2e)![y^v x^{2e}] \exp(y(e^x - 1 - x - \frac{x^2}{2})). \end{aligned}$$

The generating function $F(\hbar)$ now looks like

$$\begin{aligned} F(\hbar) &= \sum_{e,v\geq 0} \big[(2e-1)!![y^v x^{2e}] \exp(y(e^x - 1 - x - \frac{x^2}{2}))\big]\hbar^{e-v} \\ &= \sum_{e\geq 0} \big[(2e-1)!![x^{2e}] \exp(\frac{1}{\hbar}(e^x - 1 - x - \frac{x^2}{2}))\big]\hbar^{e} \end{aligned}$$

The final form of our formal power series $F(\hbar)$ can be interpreted as the path integral of a 0-dimensional scalar quantum field theory with potential $V(x) = -(e^x - 1 - x - \frac{x^2}{2})$ and action $S(x) = \frac{-x^2}{2} - V(x) = e^x - 1 - x$, i.e. with partition function

$$\mathcal{Z}(\hbar) = \frac{1}{\sqrt{2\pi\hbar}} \int_{\mathbb{R}} e^{\frac{-1}{\hbar}(e^x - 1 - x)} dx$$

and free energy $\log \mathcal{Z}(\hbar)$. In this simple case the integral actually converges for $\hbar > 0$ to a function that is closely related to the Gamma function, and the coefficients of $F(\hbar)$ are the Bernoulli numbers $B(n)$ divided by $n(n-1)$, whose asymptotic behavior is well known.

In more general graph counting problems the path integral will usually not be convergent, and it is more challenging to analyze the asymptotic behavior of its coefficients.

Recall that for the Euler characteristic of $\mathcal{MG}_n$ we want to count pairs (G, Φ) weighted by their automorphism groups and with appropriate sign. The process is similar: we first count trees, then forests, then ways to glue the leaves of the forest together in pairs. A complicating factor is the fact the sign in our formula depends on the number of *orbits* of edges under an automorphism α, instead of simply the number of edges in the forest, so we must keep track of automorphisms all the way through. In order to incorporate automorphisms into our count, we use Joyal's theory of species, which allows one to manipulate generating functions that count pairs consisting of an object together with an automorphism.

Once we have produced our generating function, we are able to isolate terms that give the generating function for $\chi(\mathrm{Out}(F_n))$. We then analyze the asymptotic behavior of the remaining terms to prove the theorem.

10.4 Euler characteristics of graph complexes. What is actually computed in Section **??** above is the orbifold Euler characteristic of Kontsevich's even commutative graph complex (see, e.g., [**?**]). Since the (odd) Lie graph complex computes the rational cohomology of $\mathcal{MG}_n$, its Euler characteristic is the Euler characteristic of $\mathcal{MG}_n$. The sign and asymptotics of the Euler characteristic of the even Lie graph complex are also determined in [**?**]. More recently Borinsky and Zagier gave a generating function for the Euler characteristics of the even and odd commutative graph complexes and determined the asymptotic behavior of their coefficients [**?**]. They also showed that their results combined with work of Getzler and Kapranov determine the asymptotics for the associative graph complex. Thus the asymptotics of the Euler characteristics and orbifold Euler characteristics of all six of the graph complexes Kontsevich defined in his original paper have now been determined.

11 Tropical geometry. In this section we discuss the relation between $\mathcal{MG}_{n,s}, \mathcal{CV}_{n,s}$ and moduli spaces of tropical curves, as defined by algebraic geometers. In order to conform to their notation, in this section (only!) we replace n by g and s by n, writing $\mathcal{MG}_{g,n}$ instead of $\mathcal{MG}_{n,s}$ and $\mathcal{CV}_{g,n}$ instead of $\mathcal{CV}_{n,s}$. We begin by recalling the classical moduli spaces $\mathcal{M}_{g,n}$.

11.1 Teichmüller space, moduli space, and the Deligne-Mumford compactification. Let $S = S_{g,n}$ be a closed orientable surface of genus g with n labeled points, $\mathcal{T}(S)$ the Teichmüller space of S and $Mod(S)$ the mapping class group of S, i.e. the group of isotopy classes of homeomorphisms of S that fix the labeled points. Teichmüller space $\mathcal{T}(S)$ is contractible and $Mod(S)$ acts on $\mathcal{T}(S)$ with finite stabilizers. Depending on your field in mathematics, you may think of the quotient $\mathcal{M}_{g,n}$ as the space of hyperbolic metrics on S after S is punctured at the labeled points, or as the space of Riemann surface (i.e. complex) structures on S with singularities at the labeled points, or as the moduli space of algebraic curves of genus g with n "marked" points. Deligne and Mumford defined a compactification $\overline{\mathcal{M}}_{g,n}$ of $\mathcal{M}_{g,n}$ from the last perspective by allowing the algebraic curves to degenerate and develop nodal singularities.

From the perspective of complex structures, a nodal curve is obtained from a Riemann surface X by collapsing each curve in a set $\{\gamma_1, \ldots, \gamma_k\}$ of disjoint simple closed curves in X to a point. This forms a collection of Riemann surfaces X_{g_i,n_i} joined at these points, called a *nodal Riemann surface.* (More correctly, the multicurve $\{\gamma_1, \ldots, \gamma_k\}$ is a set of isotopy classes of simple closed curves, and being disjoint means there is a set of representatives of these isotopy classes that are pairwise disjoint.) The boundary $\partial\overline{\mathcal{M}}_{g,n} = \overline{\mathcal{M}}_{g,n} \setminus \mathcal{M}_{g,n}$ is the union of boundary strata which intersect transversely, with one stratum for each topological type of nodal Riemann surface. This topological type is encoded by the *dual graph* to the multicurve $\mu = \{\gamma_1, \ldots, \gamma_k\}$; this has one vertex for each X_i, weighted by the genus of X_i, one edge for each γ_i, (connecting the components on either side of γ_i), and one leaf for each labeled point x_i (connected to the component containing x_i). Top-dimensional strata correspond to mapping class group orbits of a single curve, i.e. to dual graphs with one edge (which may be a loop). A stratum S is contained in the closure of S' if and only if the dual graph to S' is obtained from that of S by collapsing an edge and, if the edge is a loop, increasing the weight on the resulting vertex by 1. The associated *boundary complex* $B(\overline{\mathcal{M}}_{g,n}, \mathcal{M}_{g,n})$ has a vertex for each stratum, and a k-cell for each set of $k+1$ mutually intersecting strata, so is identified with $\mathcal{MG}^*_{g,n}$.

In [?] Bers showed that this nodal picture lifts to Teichmüller space, giving an enlargement $\overline{\mathcal{T}}(S)$ called *augmented Teichmüller space* whose quotient by $Mod(S)$ is the Deligne-Mumford compactification $\overline{\mathcal{M}}_{g,s}$. The stratification of $\partial\overline{\mathcal{M}}_{g,s}$ lifts to a stratification of $\partial\overline{\mathcal{T}}(S) = \overline{\mathcal{T}}(S) \setminus \mathcal{T}(S)$, where the top-dimensional strata are products of Teichmüller spaces of Riemann surfaces obtained by collapsing a single curve, (as opposed to a $Mod(S)$-orbit of curves). The closures of top-dimensional strata intersect non-trivially if and only if the relevant curves are disjoint (up to isotopy), in which case the intersection is a product of Teichmüller spaces of the components of the associated nodal Riemann surface. Thus in this case the boundary complex $B(\overline{\mathcal{T}}(S), \mathcal{T}(S))$ is the well-known *curve complex* $\mathcal{C}(S)$ of the surface. Furthermore, since all of the strata and all of their non-trivial intersections are contractible, the boundary $\partial\overline{\mathcal{T}}(S)$ is actually homotopy equivalent to $\mathcal{C}(S)$.

11.2 Tropicalization. If $g = 0$, $\mathcal{M}_{0,n}$ is a smooth closed subvariety of $(\mathbb{C}^*)^{\binom{n}{2}-2}$, and there is a well-defined notion of "tropicalization" which produces a new space $\mathcal{M}^{trop}_{0,n}$. The space $\mathcal{M}^{trop}_{0,n}$ turns out to be isomorphic to the space of trees $\mathcal{T}_n$ defined in Section **??**, with link $Lk(\mathcal{T}_n) = \mathcal{MG}_{0,n}$. In genus $g > 1$ (or $g = 1, n > 1$) the space $\mathcal{M}_{g,n}$ is not even compact, much less a projective variety, but algebraic geometers generalized the notion of tropicalization enough to produce a space $\mathcal{M}^{trop}_{g,n}$ for all g and n. This space has a cone point, whose link is the compact space $\mathcal{MG}^*_{g,n} = \mathcal{CV}^*_{g,n}/A_{g,n}$. In other words, the tropicalization of $\mathcal{M}_{g,n}$ is the cone on the boundary complex $B(\overline{\mathcal{M}}_{g,n}, \mathcal{M}_{g,n})$.

In [?] Chan, Galatius and Payne observed that for $n = 0$ the space $\mathcal{MG}^*_g$ is the same as the quotient of the curve complex $\mathcal{C}(S_g)$ of a surface of genus g by the action of the mapping class group $Mod(S_g)$. They then used the fact that the rational cohomology of $\mathcal{MG}^*_{g,n}$ relative to $\mathcal{MG}^\infty_{g,n}$ is isomorphic to Kontsevich's even commutative graph homology to study the cohomology of the mapping class group. In particular, by combining results of Willwacher on graph homology and Brown on the Grothendiek-Teichmüller Lie algebra they found many new cohomology classes in $Mod(S_g)$.

Since the curve complex is the boundary complex of $\overline{\mathcal{T}}(S)$ and the tropicalization of $\mathcal{M}(S)$ is the cone on the boundary complex of $\overline{\mathcal{M}}(S)$, Ramadas proposed that by analogy one should consider the cone on the curve complex $\mathcal{C}(S)$ to be the tropicalization of Teichmüller space $\mathcal{T}(S)$, [?]. Chan, Melo and Viviani had previously suggested the cone on $\mathcal{CV}^*_n$ as a candidate for the tropicalization of Teichmüller space, based on the fact that there is the following nice map from $\mathcal{T}(S)$, thought of as a space of marked hyperbolic surfaces, to this cone [?]. Each hyperbolic surface has a (possibly empty) set of curves whose length is less than the Margulis constant. This guarantees that the curves in the set are disjoint, and the map sends the surface to the dual graph of the resulting curve system. The dual graph becomes a metric graph by assigning edge lengths inversely related to the lengths of the curves.

The proposal of Ramadas leaves open the question of whether the cone on $\mathcal{CV}_n^*$ is the boundary complex of some other natural space, and should be considered to be the tropicalization of that space. The answer lies in considering handlebodies, which interpolate between graphs and surfaces since they are "thickened" graphs and have surfaces as boundary.

11.3 Complex handlebodies and $\mathcal{CV}_{g,n}^*$. Let $V = V_{g,n}$ be a handlebody of genus g with n labeled points on its boundary. A *complex handlebody* is a space X homeomorphic to V together with a complex structure on its boundary (up to isotopy). A *marking* of a complex handlebody X is a homeomorphism to V. The *Teichmüller space* $\mathcal{T}(V)$ is the space of marked complex handlebodies up to isotopy. There is an obvious map $\mathcal{T}(V) \to \mathcal{T}(\partial V)$ given by restricting the marking to the boundary surface. In fact this map is a homeomorphism, which follows since any marking of ∂V isotopic to the identity can be extended to all of V. A *meridian curve* in V is a simple closed curve in ∂V that avoids the labeled points and bounds an embedded disk in V. We enlarge $\mathcal{T}(V)$ to a subset $\overline{\mathcal{T}}(V)$ of $\overline{\mathcal{T}}(\partial V)$ by adding only the strata that correspond to sets of meridian curves. Thus the boundary complex of $\overline{\mathcal{T}}(\partial)$ relative to $\mathcal{T}(\partial V)$ is the disk complex of the handlebody instead of the curve complex of its boundary surface.

If instead of marking complex handlebodies by homeomorphisms we mark them by homotopy equivalences we obtain a space $h\mathcal{T}(V)$ of *homotopy-marked* complex handlebodies. We can enlarge it to a space $\overline{h\mathcal{T}}(V)$ by adding *nodal* homotopy-marked complex handlebodies, suitably defined. Note that a Dehn twist in a meridian curve on the boundary of V extends to a homeomorphism of V (also called a Dehn twist) that has no effect on the fundamental group of V. The set of Dehn twists in meridian curves generates a subgroup of the mapping class group $Mod(V)$ called the *twist subgroup* $Tw(V)$, and we have

THEOREM 11.1 ([?]). *Let $V = V_{g,n}$ be a complex handlebody. Then $\overline{h\mathcal{T}}(V)$ is the quotient of $\overline{\mathcal{T}}(V)$ by the twist subgroup $Tw(V)$.*

Recall that a *Schottky group* of genus g is a free subgroup of $PSL(2,\mathbb{C})$ generated by g loxodromic elements. In the case $n = 0$, the space $h\mathcal{T}(V)$ can be identified with the space of marked Schottky groups and $\overline{h\mathcal{T}}(V)$ with extended (marked) Schottky space as defined by Gerritzen and Herrlich in [?].

Neither $\overline{\mathcal{T}}(\partial V)$ nor $\overline{\mathcal{T}}(V)$ is a complex manifold, as points in the boundary do not have complex neighborhoods. Surprisingly, however, we have

THEOREM 11.2 ([?]). *For all $g > 0$ and $n \geq 0$, $\overline{h\mathcal{T}}(Vg,n)$ is a simply-connected complex manifold. The boundary $\partial\overline{h\mathcal{T}}(Vg,n) = \overline{h\mathcal{T}}(Vg,n) \setminus h\mathcal{T}(Vg,n)$ is a union of strata that are transversely-intersecting complex manifolds, and the associated boundary complex is isomorphic to $\mathcal{CV}_{g,n}^*$.*

Thus our candidate for the tropicalization of $h\mathcal{T}(Vg,n)$ is the cone on $\mathcal{CV}_{g,n}^*$.

12 Spaces of graphs and spaces of flat tori.

12.1 $\mathrm{Out}(F_n)$ and $\mathrm{GL}(n,\mathbb{Z})$. Since commutators subgroups are invariant under all automorphisms, the abelianization map $F_n \to \mathbb{Z}^n$ induces a homomorphism $\mathrm{Out}(F_n) \to \mathrm{Out}(\mathbb{Z}^n) = \mathrm{GL}(n,\mathbb{Z})$. This map is easily seen to be surjective, so the kernel IA_n can be thought of as measuring the difference between the two groups. In 1935 Magnus gave a finite generating set for IA_n and asked whether IA_n is finitely-presentable [?]. In [?] Krstic and McCool proved that IA_3 is *not* finitely presentable, but the question remains open for $n > 3$ as of this writing. If it is finitely presentable, the second homology $H_2(IA_n)$ would necessarily be finitely generated. The fact that the rational Euler characteristic is negative for all n (see Section ??) implies that the rational cohomology of IA_n is not finitely generated in *some* dimension. In fact Bestvina, Bux and Margalit have proved that for all n, $H_{2n-4}(IA_n)$ is not finitely generated [?].

12.2 The symmetric space for the special linear group. The classical symmetric space $Q_n = \mathrm{SL}(n,\mathbb{R})/SO(n)$ is a contractible space on which the group $\mathrm{SL}_n(\mathbb{Z})$ acts with finite stabilizers. It can be described in several different ways.

- Q_n is the space of marked lattices in $\mathbb{R}^n$ up to rotation and scaling, where a marking is a choice of basis;
- Q_n is the space of marked flat n-dimensional tori of volume 1, where a marking is a choice of basis for the fundamental group; or
- Q_n is the space of homothety classes of positive-definite quadratic forms on $\mathbb{R}^n$.

The description of Q_n as a space of marked flat tori resonates with our definition of $\mathcal{CV}_n$ as a space of marked metric graphs, and this analogy was one of the main inspirations for the original definition of $\mathcal{CV}_n$ (the other was the analogy with Teichmúller space with its action of the mapping class group.)

12.3 The Jacobian map. The map $\mathrm{Out}(F_n) \to \mathrm{GL}(n, \mathbb{Z})$ is reflected on the level of spaces by a map $J : \mathcal{CV}_n \to Q_n$. This is called the *Jacobian* map by analogy with the classical Jacobian map from the moduli space of genus g algebraic curves to the space of $g \times g$ complex matrices with positive definite imaginary part (not to be confused with the Abel-Jacobi map that embeds an algebraic curve into its Jacobian). The map J in the case $n = 3$ was studied in detail in [?].

To define J we use the description of Q_n as a space of positive-definite quadratic forms and identify $H_1(R_n)$ with $\mathbb{R}^n$. If (G, g) is a point in $\mathcal{CV}_n$, let $\mathbb{R}^E$ be the Euclidean space with basis the set of edges of G, and (positive definite) quadratic form q given by the diagonal matrix with i-th entry equal to the length of the i-th edge. Then $H_1(G)$ is a subspace of $\mathbb{R}^E$, so the restriction of q is a positive definite form on $H_1(G)$. But $H_1(G)$ is identified with $H_1(R_n)$ by the marking $g : R_n \to G$, and $H_1(R_n)$ is canonically identified with $\mathbb{R}^n$, so pulling back the restriction of q to $H_1(R_n)$ gives positive definite quadratic form on $\mathbb{R}^n$.

The map J extends to the simplicial closure $\mathcal{CV}_n^*$, with image contained in the set Q_n^* of positive *semi*-definite forms on $\mathbb{R}^n$. To see this, note that we can describe a point in $\mathcal{CV}_n^*$ as an equivalence class of marked metric graphs (G, g) where some of the edge-lengths are allowed to be zero. The equivalence relation is generated by setting (G, g) equal to (g', G') if G' can be obtained from G by collapsing a zero-length non-loop edge, and g' is the composition of g with this collapse. One can show that the restriction of the (now semi-definite) diagonal form q to $H_1(G)$ is well-defined on equivalence classes, so the marking gives a positive semi-definite form on $\mathbb{R}^n$.

12.4 The Torelli map. The map J is equivariant with respect to the actions of $\mathrm{Out}(F_n)$ on $\mathcal{CV}_n^*$ and $GL(n, \mathbb{Z})$ on Q_n^*, i.e. for any $\alpha \in \mathrm{Out}(F_n)$ with image $\overline{\alpha} \in GL(n, \mathbb{Z})$ the following diagram commutes:

$$\begin{tikzcd}
\mathcal{CV}_n^* \arrow[r, "J"] \arrow[d, "\alpha"] & Q_n^* \arrow[d, "\overline{\alpha}"] \\
\mathcal{CV}_n^* \arrow[r, "J"] & Q_n^*
\end{tikzcd}$$

The induced map on the quotients $\mathcal{MG}_n^* \to Q_n^*/GL(n, \mathbb{Z})$ is called the *Torelli map* by algebraic geometers.

F. Brown gave a very concrete description of the Torelli map in terms of adjacency matrices for graphs, and used that to pull back differential forms from the locally symmetric space for SL_n to “differential forms" on $\mathcal{MG}_n^*$ [?]. Here the definition of differential form takes into account the fact that $\mathcal{MG}_n^*$ is not a manifold by using its stratification by the interiors of simplices. In addition Brown proved a “Stokes theorem" for the pulled-back forms. His proof involved extending the forms to a certain *compactification* of $\mathcal{MG}_n$ which is less singular than $\mathcal{MG}_n^*$. This compactification is the quotient of an enlargement $\widehat{\mathcal{CV}}_n$ of $\mathcal{CV}_n$ analogous to the Borel-Serre bordification of symmetric space. This space $\widehat{\mathcal{CV}}_n$ was originally defined by Bestvina and Feighn in [?] by a type of of blowup procedure, and in [?] it is shown to be homeomorphic to the space obtained by removing a uniform neighborhood of the boundary $\mathcal{CV}_n^\infty$ from $\mathcal{CV}_n^*$.

13 Beyond graphs and tori: RAAGs and Gamma-complexes.

13.1 Right-angled Artin groups. Spaces of graphs are used to study free groups and their automorphisms, and spaces of flat tori are used to study free abelian groups and their automorphisms. Free groups and free abelian groups are examples of a more general class of groups called *right-angled Artin groups* (RAAGs), which occupy a prominent place in both group theory and low-dimensional topology. In geometric group theory they are a rich source both of interesting examples and counterexamples. In topology they were a key ingredient in the solution by Agol and Wise of some of Thurston's conjectures about the topology of 3-manifolds.

In general a RAAG is a finitely-generated group with a presentation specifying that some of the generators commute. If none commute, we have a free group; if they all commute we have a free abelian group. The presentation is conveniently expressed by giving a graph Γ whose vertices are the generators, with an edge for each pair of commuting generators. The resulting RAAG is denoted A_Γ. If Γ is the disjoint union of Γ_1 and Γ_2, then A_Γ is the free product of A_{Γ_1} and A_{Γ_2}. If Γ is the *join* of Γ_1 and Γ_2 (i.e. every vertex in Γ_1 is connected to every vertex in Γ_2, then A_Γ is the Cartesian product $A_{\Gamma_1} \times A_{\Gamma_2}$. But in general A_Γ does not decompose neatly into products and free products.

13.2 Salvetti complexes. There is a classical model space with fundamental group A_Γ called the *Salvetti complex* $\mathbb{S}_\Gamma$, which generalizes both the n-peraled rose and the n-torus. It is a cell complex with a single vertex and one edge (which is necessarily a loop) for each generator. We then glue in a square for each commuting pair of generators, forming a torus. We now have a complex with fundamental group A_Γ but we keep going, gluing in a k-cube for each k-clique (complete subgraph with k vertices) contained in Γ, forming a k-torus in $\mathbb{S}_\Gamma$. The point of adding all of the higher-dimensional cubes is that the universal cover $\widetilde{\mathbb{S}}_\Gamma$ then has the structure of a cube complex that satisfies Gromov's *link condition*, so is non-positively curved (CAT(0)) and therefore contractible. In particular, since $A_\Gamma = \pi_1(\mathbb{S}_\Gamma)$ acts freely on its universal cover, the quotient $\mathbb{S}_\Gamma$ is a $K(A_\Gamma, 1)$ space and can be used, for example, to compute the group cohomology of A_Γ.

The Salvetti $\mathbb{S}_\Gamma$ can also be described as a subcomplex of the n-torus T^n, where n is the number of vertices of Γ. Namely, let C_n be the unit cube spanned by the standard basis vectors $e_1, \dots, e_n$ of $\mathbb{R}^n$, which we think of as corresponding to the vertices of Γ. Let C_Γ be the union of the faces of C that are spanned by basis vectors corresponding to cliques in Γ. The n-torus T^n is the quotient of C_n obtained by identifying opposite faces., and $\mathbb{S}_\Gamma$ is the image of C_Γ in T^n.

13.3 Outer space for RAAGs. Subgroups of free groups and free abelian groups are easy to classify, but RAAGs in general have much more subtle subgroup structure; for example all hyperbolic 3-manifold groups occur (virtually) as subgroups of RAAGS, by the famous work of Agol and Wise mentioned above. Automorphisms of RAAGs must preserve the isomorphism type of subgroups, so the structure of these automorphism groups is also very subtle. In [**?**] we built a space $\mathcal{O}_\Gamma$ generalizing both $\mathcal{CV}_n$ and the symmetric space Q_n on which $\mathrm{Out}(A_\Gamma)$ acts with finite stabilizers, and proved that it is contractible.

Points in $\mathcal{O}_\Gamma$ are *marked* Γ*-spaces* (X, g). Here X is a non-positively curved metric space with fundamental group A_Γ that generalizes both metric graphs and flat tori, and g is a homotopy equivalence from the Salvetti $\mathbb{S}_\Gamma$ to X. More specifically, a Γ-space is a metric space that is isometric to a *twisted* Γ*-complex*, where a Γ*-complex* is a cube complex adapted to Γ that satisfies Gromov's link condition, and the adjective "twisted" refers to the fact that some of "cubes" may be parallelotopes. In the next section we give more details about Γ-complexes, twisted Γ-complexes and Γ-spaces.

13.3.1 Gamma-complexes. The basic example of a Γ-complex is the cube complex underlying the Salvetti complex $\mathbb{S}_\Gamma$. (We will abuse terminology slightly to call such a cube complex a Salvetti, even though its edges are not labeled and oriented.) Every Γ-complex X has (at least one) contractible subcomplex T_X which can be collapsed along hyperplanes to give a Salvetti. If Γ is discrete a Γ-complex X is a connected graph with no separating edges, and T_X is a maximal tree, which can be collapsed to give the rose. If Γ is complete, the only Γ-complex is the torus, which is already a Salvetti. If Γ is the join of a discrete graph and a complete graph, then A_Γ is the product of a free group with a free abelian group, a Γ-complex X is the product of a graph with a torus with its natural cube complex structure and T_X is the product of a maximal tree in the graph with the torus. For general Γ the structure of a Γ-complex is more subtle, reflecting the structure of the graph Γ.

The set of marked Salvetti complexes is a discrete set, and the group $\mathrm{Out}(A_\Gamma)$ acts transitively on marked Salvettis by changing the marking. As remarked above, Salvettis have the structure of non-positively curved cube complexes. To make a contractible space with an $\mathrm{Out}(A_\Gamma)$ action we need to interpolate between marked Salvettis by both changing the metric (as in the case of flat tori) and changing the homeomorphism type (as in the case of metric graphs).

To explain what a Γ-complex is in general, it is helpful to have a set of generators for $\mathrm{Aut}(A_\Gamma)$ which can get us from one marked Salvetti to any other. Such a set was given in work of Laurence [**?**] and Servatius [**?**].

13.3.2 The Laurence-Servatius generators for the automorphism group of a RAAG. We will describe the Laurence-Servatius automorphisms for $\mathrm{Aut}(A_\Gamma)$ by saying what they do on generators of A_Γ, i.e. vertices v of Γ. We note that a map on generators extends to a homomorphism of A_Γ if and only if commuting generators are sent to commuting elements of A_Γ. Commuting generators are adjacent in Γ, and we call the subgraph of Γ spanned by all w adjacent to v the *link* of v and denoted it $\mathrm{lk}(v)$. The subgraph spanned by $\mathrm{lk}(v)$ and v is the *star* of v, denoted $\mathrm{st}(v)$.

The first three types of Laurence-Servatius generators are as follows

1. The map sending $v \to v^{-1}$ and fixing all other generators induces an automorphism called an *inversion.*
2. An automorphism of the graph Γ sends edges to edges, so preserves commuting relations, inducing an automorphism of A_Γ called (unsurprisingly) a *graph automorphism.*

3. If v and m commute and $\mathrm{st}(v) \subset \mathrm{st}(m)$, then the map τ sending $v \to vm$ and fixing the other generators induces a homomorphism of A_Γ called a *twist*, with multiplier m. The map $v \to vm^{-1}$ is an inverse to τ so a twist is an automorphism.

The other Laurence-Servatius generators are called *folds* and *partial conjugations*, but we will give a larger set of generators that includes them both; these are called Γ-*Whitehead automorphisms*, by analogy with J.H.C. Whitehead's generating set for the automorphism group of a free group. The advantage is that a single Γ-Whitehead automorphism can do several twists and folds at the same time, giving a more efficient way to decompose an automorphism. In what follows we will abbreviate Γ-Whitehead by ΓW.

A ΓW automorphism is determined by a vertex m of Γ (called the multiplier) and a partition $\mathcal{P}$ of the set $\Gamma^\pm$ of vertices of Γ and their inverses called a Γ-*Whitehead partition based at* m, which is defined as follows. The graph spanned by vertices in $\Gamma \setminus \mathrm{st}(m)$ is a union of connected components, some of which may be singletons. Write $\Gamma \setminus \mathrm{st}(m) = C_1 \sqcup \ldots \sqcup C_k \sqcup y_1 \sqcup \ldots, y_\ell$ where each C_i has at least two vertices and each y_i is a singletons, and decompose $\Gamma^\pm \setminus \{m, m^{-1}\}$ into *m-indecomposable pieces*

$$C_1^\pm \sqcup C_k^\pm \sqcup \{y_1\} \sqcup \{y_1^{-1}\}, \ldots \sqcup \{y_\ell\} \sqcup \{y_\ell^{-1}\}.$$

A ΓW-partition $\mathcal{P}$ based at m is a partition of $\Gamma^\pm$ into three parts

$$\mathcal{P} = \mathrm{lk}(m)^\pm \sqcup P_0 \sqcup P_1,$$

where P_0 contains m and a proper subset of the m-indecomposable pieces and P_1 contains m^{-1} and the rest of the m-indecomposable pieces (see Figure **??** for an example). The associated ΓW-automorphism $\alpha(\mathcal{P}, m)$ is defined by

$$\alpha(\mathcal{P}, m) : \begin{cases} v \mapsto v & \text{if } v \in \mathrm{st}(m) \\ v \mapsto m^{-1}v & \text{if } v^{-1} \in P_0, v \in P_1 \\ v \mapsto vm & \text{if } v \in P_0, v^{-1} \in P_1, v \neq m \\ v \mapsto m^{-1}vm & \text{if } v \in C_i \subset P_0 \\ v \mapsto v & \text{if } v \in C_i \subset P_1 \end{cases}$$

For example, if $P_0 = \{y_1\}$ then $\alpha(\mathcal{P}, m)$ sends $y_1 \mapsto y_1 m$ and fixes all other generators; this is called a *fold*. If $P_0 = C_1^\pm$ then $\alpha(\mathcal{P}, m)$ conjugates each vertex in C_1 by m and fixes all other generators; this is called a *partial conjugation.*

The subgroup of $\mathrm{Out}(A_\Gamma)$ generated by twists and inversions is called the *twist subgroup*, denoted $\mathrm{Tw(A_\Gamma)}$. The subgroup generated by ΓW automorphisms, inversions and graph automorphisms is called the *untwisted subgroup*, denoted $U(A_\Gamma)$. If Γ is complete then $\mathrm{Out}(A_\Gamma) = \mathrm{Tw(A_\Gamma)} = \mathrm{GL(n, \mathbb{Z})}$. If Γ is discrete, then $\mathrm{Out}(A_\Gamma) = U(A_\Gamma) = \mathrm{Out}(F_n)$, Techniques used to study $\mathrm{Out}(F_n)$ can often be adapted to work for $U(A_\Gamma)$. In particular, a simplicial complex K_Γ modeled on the spine of $\mathcal{CV}_n$ was constructed for $U(A_\Gamma)$ in [**?**]. The proof that full outer space $\mathcal{O}_\Gamma$ is contractible depends on the fact proved there that K_Γ is contractible. Vertices of K_Γ are Γ-complexes, so we begin by explaining what they are.

13.3.3 Two-vertex Gamma-complexes. Given a ΓW partition $\mathcal{P}$ based at m we can construct a Γ-complex $\mathbb{S}^\mathcal{P}$ with two vertices as follows. Let $\mathbb{S}_{\mathrm{lk}(m)}$ be the Salvetti for $\mathrm{lk}(m)$, and $T_\mathcal{P}$ the cylinder $\mathbb{S}_{\mathrm{lk}(m)} \times [0, 1]$, with two vertices $b_1 \in \mathbb{S}_{\mathrm{lk}(m)} \times \{0\}$ and $b_2 \in \mathbb{S}_{\mathrm{lk}(m)} \times \{1\}$. This cylinder will be the subcomplex of $\mathbb{S}^\mathcal{P}$ that is collapsed to a single copy of $\mathbb{S}_{\mathrm{lk}(m)}$ to recover $\mathbb{S}_\Gamma$, so all other edges of $\mathbb{S}^\mathcal{P}$ are directed and labeled with vertices of Γ. If $v \in P_i$ (resp $v^{-1} \in P_i$) attach the terminal vertex (resp. initial vertex) of the corresponding edge to b_i. This completes the 1-skeleton of $\mathbb{S}^\mathcal{P}$. Note that if $v \in \mathrm{lk}(m)$ there are two loops labeled v, one in the copy of $\mathbb{S}_{\mathrm{lk}(m)}$ at the top of the cylinder and one in the copy at the bottom. Now attach a k-dimensional cube whenever you see the 1-skeleton of a k-cube cube that is labeled by commuting generators; this guarantees that $\mathbb{S}^\mathcal{P}$ satisfies Gromov's link condition, so is non-positively curved. Finally, label the edge in $T_\mathcal{P}$ with the label $(\mathcal{P}, m)$ so that all edges in $\mathbb{S}^\mathcal{P}$ are labeled. See Figure **??**

Let $c_\mathcal{P} \colon \mathbb{S}^\mathcal{P} \to \mathbb{S}_\Gamma$ be the map that collapses $T_\mathcal{P}$ to one copy of $\mathbb{S}_{\mathrm{lk}(m)}$; we call this a *hyperplane collapse*, and we call the inverse operation of inserting $T_\mathcal{P}$ into $\mathbb{S}_\Gamma$ a *blowup*. Since m commutes with everything in $\mathrm{lk}(m)$, we observe that $\mathbb{S}^\mathcal{P}$ also contains a cylinder $T_m = \mathbb{S}_{\mathrm{lk}(m)} \times [0, 1]$ isomorphic to $T_\mathcal{P}$, with vertical edge labeled m instead of $\mathcal{P}$. Collapsing T_m instead of $T_\mathcal{P}$ gives a new Salvetti, and the composition $c_\mathcal{P} \circ c_m^{-1} \colon \mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ induces the ΓW-automorphism $\alpha(P, m)$ on A_Γ.

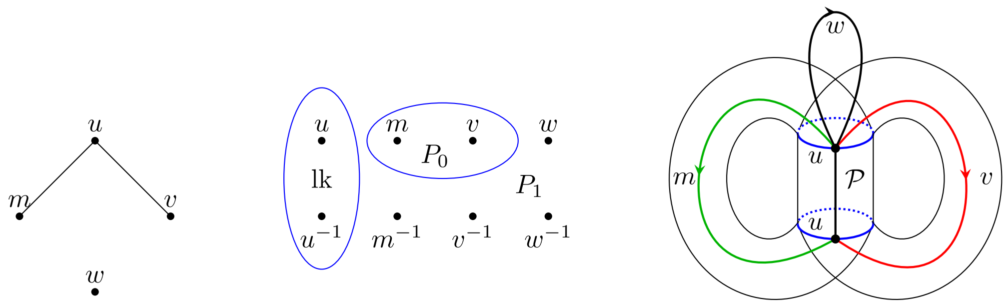

Figure 13.1: A graph Γ, a ΓW-partition $\mathcal{P} = (\mathrm{lk}(m)^{\pm}|\{m, v\}|\{m^{-1}, v^{-1}, w, w^{-1}\})$ and the blowup $\mathbb{S}^{\mathcal{P}}$.

Thus by including two-vertex blowups in our spine K_Γ we can pass from any marked Salvetti $(\mathbb{S}_\Gamma, h)$ to any other $(\mathbb{S}_\Gamma, \alpha \circ h)$ by blowups and collapses as long as α is in the subgroup generated by ΓW-automorphisms. In fact, since inversions and graph automorphisms fix the identity-marked Salvetti, we can allow α to be any automorphism in the *untwisted subgroup* $U(A_\Gamma)$.

13.3.4 Gamma complexes with more than two vertices. To eventually get a complex K_Γ which is not only connected but actually contractible we will have to include Γ-complexes with more than 2 vertices. In general a Γ-complex X with k vertices is constructed from a set of $k-1$ ΓW-partitions. These must be pairwise compatible, where $\mathcal{P} = (\mathrm{lk}(m)^{\pm}|P_0|P_1)$ and $\mathcal{Q} = (\mathrm{lk}(n)^{\pm}|Q_0|Q_1)$ are *compatible* if either m and n commute or $P_i \cap Q_j = \emptyset$ for some i, j. Given a pairwise compatible set $\Pi = \{\mathcal{P}_1, \ldots, \mathcal{P}_k\}$ we can construct a complex $\mathbb{S}^{\Pi}$ with k vertices, that can be collapsed to $\mathbb{S}_\Gamma$ by a series of hyperplane collapses. The construction is natural but a bit too technical to include here, but it may be useful to keep in mind the case when Γ is discrete. In this case $X = \mathbb{S}^{\Pi}$ is a graph and $T = T_\Pi$ is a maximal tree. Removing an edge of T separates T into two subtrees, T_0 and T_1. The edges of X that are not in T are oriented and labeled with vertices of Γ, so that $\Gamma^{\pm}$ is partitioned into two sets, depending on where the associated edge starts or terminates. Distinct edges e and e' cut up T in two different ways $T_0 \sqcup T_1$ and $T_0' \sqcup T_1'$, with some T_i disjoint from some T_j'.

13.3.5 Recap. The space K_Γ is a simplicial complex whose vertices are Γ-complexes with untwisted markings. This means that composing the marking with the collapse map $\mathbb{S}^{\Pi} \to \mathbb{S}_\Gamma$ gives a map $\mathbb{S}_\Gamma \to \mathbb{S}_\Gamma$ that induces an element of $U(A_\Gamma)$. Two vertices of K_Γ are connected by an edge if one can be obtained from the other by collapsing a hyperplane, and a k-simplex is given by a chain of k hyperplane collapses. The group $U(A_\Gamma)$ acts on K_Γ by changing the marking. The usefulness of K_Γ for the study of $\mathrm{Out}(A_\Gamma)$ is due to the following theorem.

THEOREM 13.1 ([**?**]). *The simplicial complex K_Γ is contractible. The group $U(A_\Gamma)$ acts on K_Γ cocompactly. The stabilizer of a vertex in K_Γ is isomorphic to the group of cube-complex isomorphisms of the associated Γ-complex, so in particular is a finite group.*

13.4 Realization. In [**?**] we proved a *realization theorem* analogous to the classical ?? theorem for the action of $\mathrm{GL}(n, \mathbb{Z})$ on Q_n or the Culler-Khramtsov-Zimmermann realization theorem for the action of $\mathrm{Out}(F_n)$ on $\mathcal{CV}_n$.

THEOREM 13.2. *Let $U^0(A_\Gamma)$ be the subgroup of $U(A_\Gamma)$ generated by Γ-Whitehead automorphisms and inversions. Every finite subgroup of $U^0(A_\Gamma)$ fixes a vertex of K_Γ.*

As a corollary we get the following information about the subgroup structure of $U^0(A_\Gamma)$.

COROLLARY 13.3. *There are only finitely many conjugacy classes of finite subgroups of $U^0(A_\Gamma)$*

Using Theorem **??** it should possible to learn more about finite subgroups of $U^0(A_\Gamma)$ by studying automorphism groups of Γ-complexes, as has been done for finite subgroups of $\mathrm{Out}(F_n)$ by realizing them as automorphisms of graphs (see, e.g. [**?**]).

The restriction to $U^0(A_G)$ is unfortunate and is an artifact of the proof, which relies heavily on the structure of Γ-complexes. There is a unique vertex of K_Γ whose marking is not just a homotopy equivalence, but a homeomorphism, and this vertex is fixed by any graph automorphism. But it is still not clear how to find a fixed point for a general element of $U(A_\Gamma)$ which may be a product of graph automorphisms and elements of $U^0(A_\Gamma)$.

13.4.1 Twisted Gamma-complexes and Gamma-spaces As remarked above, a Γ-complex always has a non-positively curved (NPC) metric where all cubes are unit cubes. However, if Γ is not discrete (for instance if Γ is complete) we need to allow more general metrics, though we still want to maintain non-positive curvature.

We impose the following rules for a metric on a Γ-complex to be allowable. All edges with the same label must have the same length. If the edges of a square S in a Γ-complex are labeled by generators v, w with $\mathrm{st}(v) \subset \mathrm{st}(w)$ (or by $(\mathcal{P}, v)$ and w or by $(\mathcal{P}, v)$ and $(\mathcal{Q}, w)$), then S is allowed to be isometric to a parallelogram, otherwise S must be isometric to a rectangle. Moreover, if $\mathrm{st}(v) \subset \mathrm{st}(w)$ all parallelograms with edge labels v and w or v and $(\mathcal{Q}, w)$ must be isometric, and all parallelograms with edge labels $(\mathcal{P}, v)$ and w or $(\mathcal{P}, v)$ and $(\mathcal{Q}, w)$ must be isometric. The metrics on the two-dimensional faces of a cube of any dimension then determine a unique parallelotope metric on the cube. A Γ-complex equipped with such a metric is a non-positively curved space called a *twisted* Γ*-complex.*

Finally, a Γ*-space* is a metric space isometric to a twisted Γ-complex, but without a specifed combinatorial structure...it is not *a priori* divided into parallelotopes. Outer space $\mathcal{O}_\Gamma$ is the space of all marked Γ-spaces, where the marking is any homotopy equivalence to the standard Salvetti. The proof that $\mathcal{O}_\Gamma$ is contractible depends on contractibility of K_Γ and goes in three stages.

1. Form a new space Σ_Γ by allowing edge-lengths (but not the angles) of Γ-complexes to vary, and taking homothety classes of the resulting metric spaces. Then show that Σ_Γ deformation retracts to K_Γ. This step is straightforward.

2. Enlarge Σ_g to form a new space $\mathcal{T}_\Gamma$ by allowing the angles between some edges to skew, but still insist on untwisted markings. Show that Σ_Γ is a deformation retract of $\mathcal{T}_\Gamma$. This step is reasonably straightforward. The idea is to equivariantly rotate skewed edges until they become orthogonal. The difficulty lies in ensuring that the intermediate Γ-complexes remain non-positively curved.

3. Map $\mathcal{T}_\Gamma$ to $\mathcal{O}_\Gamma$ by forgetting the combinatorial structure. Then
   - show this map is surjective, and
   - show this map is a fibration and the fibers are contractible.

   Both of these require some work.

In the end we have the following theorem.

THEOREM 13.4 ([?]). *The space $\mathcal{O}_\Gamma$ is finite-dimensional and contractible, and* $\mathrm{Out}(A_\Gamma)$ *acts on $\mathcal{O}_\Gamma$ with finite stabilizers.*

A natural question to ask, now that we have a contractible space with a nice $\mathrm{Out}(A_\Gamma)$-action, is whether there is a realization theorem for any finite subgroup of $\mathrm{Out}(A_\Gamma)$ acting on $\mathcal{O}_\Gamma$. At least for finite subgroups of $\mathrm{Out}^0(A_\Gamma)$, the group generated by all generators except graph automorphisms, we hope to have a (positive) answer soon.

14 Some previous survey articles. As mentioned in the introduction, there is a great deal more known about spaces of graphs and automorphisms of free groups than is included in this report. Over the years I have written a number of survey articles which cover some of this, and include references to much more. Here is a brief guide to what is contained in these articles, in chronological order.

2002 [?] The first survey article I wrote was for a 2000 conference in Haifa, Israel on geometric group theory. This includes the definition of Outer space $\mathcal{CV}_n$ and summarizes the state of knowledge about $\mathrm{Out}(F_n)$ at the turn of the century. It includes a long list of questions (many of which have since been answered), mostly concentrating on analogies with arithmetic groups acting on symmetric spaces and mapping class groups acting on Teichmüller space.

2006 [?] This is an article written for the proceedings of the 2006 Internation Congress in Madrid. It discusses different ways to define Outer space, including their relative advantages and how they generalize, and gives examples of how they have been used to study the group $\mathrm{Out}(F_n)$.

2015 [?] These are lecture notes from a 3-lecture minicourse on right-angled Artin groups and their automorphisms given at the 2013 edition of the conference Groups St. Andrews.

2015 [?] This is an article written for the AMS Current Events Bulletin that gives an exposition of Bestvina and Algom-Kfir's work on the Lipschitz metric for Outer space. This is a non-symmetric metric which is analogous to Thurston's Lipschitz metric on Teichmüller space.

2017 [?] These are course notes for an 8-week course given at the University of Warwick in 2017. It includes basic definitions, then discusses assembly maps, the action of the symmetric group on cohomology, and Kontsevich's graph complexes.

2018 [?] This article was written for the proceedings of the 2016 European Congress in Berlin. It includes descriptions of various connections between spaces of graphs and other areas of mathematics, biology and physics.

2025 [?] These are notes from a minicourse given at CIRM in Luminy in 2023. It includes basic definitions and an exposition of some of early results on $\mathrm{Out}(F_n)$, starting from Culler's fixed point theorem for the action on $\mathcal{CV}_n$.

Finally, here are a couple of resources written by other people for general information about Outer space and automorphisms of free groups.

- [?] Park City Math Institute lectures from 2012 by Mladen Bestvina.
- [?] Class notes from Alex Wright's 2024 graduate student seminar at the University of Michigan
- [?] The first part of a textbook by Lee Mosher, posted on the arXiv in 2020.

Acknowledgments. I owe deep thanks to every one of my coauthors over the years. I am also very grateful to the Isaac Newton Institute in Cambridge for support during the writing of this report.

References

[1] O. Baker, *The Jacobian map on outer space*, phd thesis, Cornell University, July 2011, https://hdl.handle.net/1813/30769.

[2] L. Bartholdi, *The rational homology of the outer automorphism group of* F_7, New York J. Math., 22 (2016), pp. 191–197, http://nyjm.albany.edu:8000/j/2016/22_191.html.

[3] G. Baumslag and T. Taylor, *The centre of groups with one defining relator*, Math. Ann., 175 (1968), pp. 315–319, https://doi.org/10.1007/BF02063216.

[4] L. Bers, *Spaces of degenerating Riemann surfaces*, in Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973), Ann. of Math. Stud., No. 79, Princeton Univ. Press, Princeton, NJ, 1974, pp. 43–55.

[5] M. Bestvina, *Geometry of outer space*, in Geometric group theory, vol. 21 of IAS/Park City Math. Ser., Amer. Math. Soc., Providence, RI, 2014, pp. 173–206, https://doi.org/10.1090/pcms/021/06.

[6] M. Bestvina, *Homological stability of* $Aut(F_n)$ *revisited*, in Hyperbolic geometry and geometric group theory, vol. 73 of Adv. Stud. Pure Math., Math. Soc. Japan, Tokyo, 2017, pp. 1–11, https://doi.org/10.2969/aspm/07310001.

[7] M. Bestvina, K.-U. Bux, and D. Margalit, *Dimension of the Torelli group for* $\mathrm{Out}(F_n)$, Invent. Math., 170 (2007), pp. 1–32, https://doi.org/10.1007/s00222-007-0055-0.

[8] M. Bestvina and M. Feighn, *The topology at infinity of* $\mathrm{Out}(F_n)$, Invent. Math., 140 (2000), pp. 651–692, https://doi.org/10.1007/s002220000068.

[9] L. J. Billera, S. P. Holmes, and K. Vogtmann, *Geometry of the Space of Phylogenetic Trees*, Advances in Applied Mathematics, 27 (2001), pp. 733 – 767, https://doi.org/10.1006/aama.2001.0759.

[10] M. Borinsky, *Renormalized asymptotic enumeration of Feynman diagrams*, Annals of Physics, 385 (2017), pp. 95 – 135.

[11] M. Borinsky and K. Vogtmann, *The Euler characteristic of* $\mathrm{Out}(F_n)$, Comment. Math. Helv., 95 (2020), pp. 703–748, https://doi.org/10.4171/cmh/501.

[12] M. Borinsky and K. Vogtmann, *The Euler characteristic of the moduli space of graphs*, Adv. Math., 432 (2023), pp. Paper No. 109290, 58, https://doi.org/10.1016/j.aim.2023.109290.

[13] M. Borinsky and K. Vogtmann, *Computing Euler characteristics using quantum field theory*, in Geometric methods in group theory—papers dedicated to Ruth Charney, vol. 34 of Sémin. Congr., Soc. Math. France, Paris, 2025, pp. 1–15.

[14] M. Borinsky and D. Zagier, *On the euler characteristic of the commutative graph complex and the top weight cohomology of* $\Updownarrow_g$, 2024, https://arxiv.org/abs/2405.04190.

[15] C. Bregman, R. Charney, and K. Vogtmann, *Outer space for RAAGs*, Duke Math. J., 172 (2023), pp. 1033–1108, https://doi.org/10.1215/00127094-2023-0007.

[16] C. Bregman, R. Charney, and K. Vogtmann, *Finite groups of untwisted outer automorphisms of RAAGs*, Algebr. Geom. Topol., 25 (2025), pp. 2413–2440, https://doi.org/10.2140/agt.2025.25.2413.

[17] F. Brown, *Invariant differential forms on complexes of graphs and Feynman integrals*, SIGMA Symmetry Integrability Geom. Methods Appl., 17 (2021), pp. Paper No. 103, 54, https://doi.org/10.3842/SIGMA.2021.103.

[18] K. S. Brown, *Cohomology of groups*, vol. 87 of Graduate Texts in Mathematics, Springer-Verlag, New York, 1994. Corrected reprint of the 1982 original.

[19] K.-U. Bux, P. Smillie, and K. Vogtmann, *On the bordification of outer space*, J. Lond. Math. Soc. (2), 98 (2018), pp. 12–34, https://doi.org/10.1112/jlms.12124.

[20] M. Chan, S. r. Galatius, and S. Payne, *Tropical curves, graph complexes, and top weight cohomology of* M_g, J. Amer. Math. Soc., 34 (2021), pp. 565–594, https://doi.org/10.1090/jams/965.

[21] M. Chan, M. Melo, and F. Viviani, *Tropical Teichmüller and Siegel spaces*, in Algebraic and combinatorial aspects of tropical geometry, vol. 589 of Contemp. Math., Amer. Math. Soc., Providence, RI, 2013, pp. 45–85, https://doi.org/10.1090/conm/589/11742.

[22] R. Charney, N. Stambaugh, and K. Vogtmann, *Outer space for untwisted automorphisms of right-angled Artin groups*, Geometry and Topology, 21 (2017), pp. 1131 – 1178, https://doi.org/10.2140/gt.2017.21.1131.

[23] J. Conant, A. Hatcher, M. Kassabov, and K. Vogtmann, *Assembling homology classes in automorphism groups of free groups*, Comment. Math. Helv., 91 (2016), pp. 751–806, https://doi.org/10.4171/CMH/402.

[24] J. Conant and K. Vogtmann, *Morita classes in the homology of* $\mathrm{Aut}(F_n)$ *vanish after one stabilization*, Groups Geom. Dyn., 2 (2008), pp. 121–138, https://doi.org/10.4171/GGD/33.

[25] G. Crowe, P. Helms, and K. Vogtmann, *The geometry and topology of automorphism groups of free groups*, in Groups acting on fractals, hyperbolicity and self-similarity, vol. 63 of Panor. Synthèses, Soc. Math. France, Paris, 2025, pp. 69–91.

[26] M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Inventiones Mathematicae, 84 (1986), pp. 91 – 119, https://doi.org/10.1007/bf01388734.

[27] S. r. Galatius, *Stable homology of automorphism groups of free groups*, Ann. of Math. (2), 173 (2011), pp. 705–768, https://doi.org/10.4007/annals.2011.173.2.3.

[28] L. Gerritzen and F. Herrlich, *The extended Schottky space*, J. Reine Angew. Math., 389 (1988), pp. 190–208, https://doi.org/10.1515/crll.1988.389.190.

[29] A. Hatcher, *Homological stability for automorphism groups of free groups*, Comment. Math. Helv., 70 (1995), pp. 39–62, https://doi.org/10.1007/BF02565999.

[30] A. Hatcher and K. Vogtmann, *Cerf theory for graphs*, J. London Math. Soc. (2), 58 (1998), pp. 633–655, https://doi.org/10.1112/S0024610798006644.

[31] A. Hatcher and K. Vogtmann, *Rational homology of Aut(F_n)*, Mathematical Research Letters, (1998).

[32] A. Hatcher and K. Vogtmann, *Homology stability for outer automorphism groups of free groups*, Algebr. Geom. Topol, 4 (2004), pp. 1253 – 1272.

[33] A. Hatcher, K. Vogtmann, and N. Wahl, *Erratum to: "Homology stability for outer automorphism groups of free groups [Algebr. Geom. Topol.* **4** *(2004), 1253–1272 (electronic); mr2113904] by Hatcher and Vogtmann*, Algebr. Geom. Topol., 6 (2006), pp. 573–579, https://doi.org/10.2140/agt.2006.6.573, https://doi-org.proxy.library.cornell.edu/10.2140/agt.2006.6.573.

[34] M. Kontsevich, *Feynman diagrams and low-dimensional topology*, First European Congress of Mathematics, 2 (1992), pp. 97 – 121.

[35] M. Kontsevich, *Formal (non)commutative symplectic geometry*, in The Gel'fand Mathematical Seminars, 1990–1992, Birkhäuser Boston, Boston, MA, 1993, pp. 173–187.

[36] S. Krstic and J. McCool, *The non-finite presentability of IA(F3) and GL2(Z[t, tâĹŠ1])*, Inventiones Mathematicae, 129 (1997), pp. 595 – 606, https://doi.org/10.1007/s002220050174.

[37] M. R. Laurence, *A generating set for the automorphism group of a graph group*, J. London Math. Soc. (2), 52 (1995), pp. 318–334, https://doi.org/10.1112/jlms/52.2.318.

[38] W. Magnus, *Über n-dimensionale Gittertransformationen*, Acta Math., 64 (1935), pp. 353–367, https://doi.org/10.1007/BF02545673, https://doi-org.proxy.library.cornell.edu/10.1007/BF02545673.

[39] R. McEwen and M. C. B. Zaremsky, *A combinatorial proof of the degree theorem in Auter space*, New York J. Math., 20 (2014), pp. 217–228.

[40] E. Miller, M. Owen, and J. S. Provan, *Polyhedral computational geometry for averaging metric phylogenetic trees*, Adv. in Appl. Math., 68 (2015), pp. 51–91, https://doi.org/10.1016/j.aam.2015.04.002.

[41] S. Morita, *Structure of the mapping class groups of surfaces: a survey and a prospect*, in Proceedings of the Kirbyfest (Berkeley, CA, 1998), vol. 2 of Geom. Topol. Monogr., Geom. Topol. Publ., Coventry, 1999, pp. 349–406, https://doi.org/10.2140/gtm.1999.2.349.

[42] S. Morita, T. Sakasai, and M. Suzuki, *Integral Euler characteristic of* $\operatorname{Out} F_{11}$, Exp. Math., 24 (2015), pp. 93–97, https://doi.org/10.1080/10586458.2014.956373.

[43] L. Mosher, *The topology, geometry and dynamics of free groups. part i: Outer space, fold paths, and the nielsen/whitehead problems*, 2020, https://arxiv.org/abs/2001.02776.

[44] R. Ramadas, *Thurston obstructions and tropical geometry*, Bull. Lond. Math. Soc., 57 (2025), pp. 2404–2428, https://doi.org/10.1112/blms.70102.

[45] R. Ramadas, R. Silversmith, K. Vogtmann, and R. R. Winarski, *Handlebodies, outer space, and tropical geometry*, 2025, https://arxiv.org/abs/2507.02440.

[46] O. Randal-Williams and N. Wahl, *Homological stability for automorphism groups*, Adv. Math., 318 (2017), pp. 534–626, https://doi.org/10.1016/j.aim.2017.07.022.

[47] A. Robinson and S. Whitehouse, *The tree representation of* Σ_{n+1}, Journal of Pure and Applied Algebra, 111 (1996), pp. 245–253, https://doi.org/10.1016/0022-4049(95)00116-6.

[48] H. Servatius, *Automorphisms of graph groups*, J. Algebra, 126 (1989), pp. 34–60, https://doi.org/10.1016/0021-8693(89)90319-0.

[49] J. Smillie and K. Vogtmann, *A generating function for the Euler characteristic of (Fn)*, Journal of Pure and Applied Algebra, 44 (1987), pp. 329 – 348.

[50] J. Smillie and K. Vogtmann, *Automorphisms of graphs, p-subgroups of* $\mathrm{Out}(F_n)$ *and the Euler characteristic of* $\mathrm{Out}(F_n)$, J. Pure Appl. Algebra, 49 (1987), pp. 187–200, https://doi.org/10.1016/0022-4049(87)90128-9.

[51] K. Vogtmann, *Local structure of some* $\mathrm{Out}(F_n)$*-complexes*, Proc. Edinburgh Math. Soc. (2), 33 (1990), pp. 367–379, https://doi.org/10.1017/S0013091500004818.

[52] K. Vogtmann, *Automorphisms of free groups and outer space*, Geometriae Dedicata, 94 (2002), pp. 1 – 31.

[53] K. Vogtmann, *The cohomology of automorphism groups of free groups*, in International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1101–1117.

[54] K. Vogtmann, *$GL(n,\mathbb{Z})$, $Out(F_n)$ and everything in between: automorphism groups of RAAGs*, in Groups St Andrews 2013, vol. 422 of London Math. Soc. Lecture Note Ser., Cambridge Univ. Press, Cambridge, 2015, pp. 105–127.

[55] K. Vogtmann, *On the geometry of outer space*, Bull. Amer. Math. Soc. (N.S.), 52 (2015), pp. 27–46, https://doi.org/10.1090/S0273-0979-2014-01466-1.

[56] K. Vogtmann, *Outer space and automorphisms of free groups*, AMS Open Math Notes, 110672 (2017), https://www.ams.org/open-math-notes.

[57] K. Vogtmann, *The topology and geometry of automorphism groups of free groups*, in European Congress of Mathematics, Eur. Math. Soc., Zürich, 2018, pp. 181–202.

[58] A. Wright, *Outer automorphism groups of free groups*, AMS Open Math Notes, 111470 (2025), https://www.ams.org/open-math-notes.