Pacific Journal of Mathematics

A SPECTRAL SEQUENCE DETERMINING THE HOMOLOGY OF $\mathrm{Out}(F_n)$ IN TERMS OF ITS MAPPING CLASS SUBGROUPS

MATTHEW HORAK

Volume 227 No. 1 September 2006

PACIFIC JOURNAL OF MATHEMATICS
Vol. 227, No. 1, 2006

A SPECTRAL SEQUENCE DETERMINING THE HOMOLOGY OF Out(F_n) IN TERMS OF ITS MAPPING CLASS SUBGROUPS

MATTHEW HORAK

We construct a covering of the spine of the Culler–Vogtmann outer space Out(F_n) by complexes of ribbon graphs. By considering the equivariant homology for the action of Out(F_n) on this covering, we construct a spectral sequence converging to the homology of Out(F_n) that has its E^1 terms given by the homology of mapping class groups and their subgroups. This spectral sequence can be seen as encoding all of the information of how the homology of Out(F_n) is related to the homology of mapping class groups and their subgroups

1. Introduction

Much is known about the cohomology of mapping class groups of surfaces. (*All surfaces considered in this work are assumed orientable.*) Let Σ be a surface with boundary, and let $P\Gamma(\Sigma)$ be the group of isotopy classes, relative to the boundary, of homeomorphisms of Σ that fix the boundary pointwise. We call $P\Gamma(\Sigma)$ the pure mapping class group of Σ. Harer [1985] proved that the k-th integral homology group of $P\Gamma(\Sigma)$ is independent of the genus and number of boundary components of Σ if the genus of Σ is at least $3k$. Later, Ivanov [1989] and Harer [1993] improved these bounds, and Harer was able to find the exact location at which the rational homology stabilizes. He also computed in [Harer 1986] the virtual cohomological dimension (VCD) of $P\Gamma(\Sigma)$ and showed that this group has no rational homology at its VCD. Madsen and Weiss [2002] have determined the entire stable integral cohomology algebra of pure mapping class groups. In particular, their result verifies the conjecture of Mumford that the stable rational cohomology algebra is a polynomial algebra with a single generator in each even dimension.

For outer automorphism groups of free groups, much less is known. Culler and Vogtmann [1986] have compute the VCD of Out(F_n) by considering the action of this group on a contractible simplicial complex known as the spine of outer space. Recently, Hatcher and Vogtmann [2004] have shown that the k-th integral homology of Out(F_n) is independent of n if $n \geq 2k + 5$, but the exact stability

MSC2000: primary 20F65; secondary 57M07.
Keywords: free group automorphisms, mapping class groups, ribbon graphs.

range remains unknown. Indeed, there are no nontrivial stable rational homology or cohomology classes known for $\mathrm{Out}(F_n)$. For a good survey of current knowledge about $\mathrm{Out}(F_n)$ and $\mathrm{Aut}(F_n)$, see [Vogtmann 2002].

Since the mapping class groups of surfaces appear as subgroups of $\mathrm{Out}(F_n)$, it is natural to try to understand the homology of $\mathrm{Out}(F_n)$ in terms of the homology of mapping class groups. This paper represents an attempt to clarify this relationship. For a punctured surface, the mapping class group is simply the group of isotopy classes of orientation preserving homeomorphisms of the surface. The group of all isotopy classes of homeomorphisms of a punctured surface will be called the extended mapping class group of that surface, so that extended mapping class groups contain orientation reversing homeomorphisms.

We construct a first quadrant spectral sequence that converges to $H_*(\mathrm{Out}(F_n))$, many of whose terms consist of the homology of mapping class groups. The spectral sequence arises from a covering of the spine of outer space by a collection of subcomplexes called ribbon graph subcomplexes. We prove that the nerve of this covering is contractible. The spectral sequence mentioned above is the equivariant homology spectral sequence of the action of $\mathrm{Out}(F_n)$ on this nerve.

All of the terms on the E^1 page of this spectral sequence are given by the homology simplex stabilizers. For a 0-simplex, the stabilizer is simply the extended mapping class group of a punctured surface Σ, or equivalently the stabilizer of the set conjugacy classes in $F_n \cong \pi_1 \Sigma$ that correspond to positively and negatively oriented curves about the punctures of Σ. For higher-dimensional simplices, stabilizers are given by the generalized stabilizers $\mathcal{A}_{U,G}$ of m-tuples of conjugacy classes, which are studied in [McCool 1975]. (These groups are finite-index subgroups of the ordinary stabilizers of certain sets of conjugacy classes in F_n.) We prove:

Theorem. *For any* $\mathrm{Out}(F_n)$*-module* M *there is a spectral sequence of the form*

$$E^1_{pq} = \bigoplus_{\sigma \in \Delta_p} H_q(G_\sigma; M_\sigma) \Rightarrow H_{p+q}(\mathrm{Out}(F_n); M),$$

where Δ_0 *is the set of homeomorphism classes of punctured orientable surfaces with fundamental group* F_n *and where for a vertex* $v \in \Delta_0$ *corresponding to surface* Σ, *the stabilizer* G_v *is the extended mapping class group* $\mathrm{MCG}^{\pm}(\Sigma)$. *Moreover, for* $p > 0$, *each* G_σ *is a generalized stabilizer of the form* $\mathcal{A}_{U_\sigma, H_\sigma}$.

The rest of this paper is organized as follows. In Sections 2 and 3 we review the definitions of outer space, the spine of outer space, ribbon graphs and some related objects. In Section 4, we construct a covering of the spine of outer space by subcomplexes of ribbon graphs. Section 5 is devoted to the proof of the fact that the nerve of this covering is contractible. In Section 6 we determine simplex stabilizers for the action of $\mathrm{Out}(F_n)$ on the nerve. The analysis of the equivariant homology spectral sequence for this action appears in final two sections where we prove the above theorem and use Harer's stability theorems to find rough upper bounds on the dimensions of some portions of the E^∞ page of the spectral sequence. These bounds limit the possible contribution that the mapping class subgroups of $\mathrm{Out}(F_n)$ can make to the homology of $\mathrm{Out}(F_n)$.

2. Outer space

For convenience and to set notation, we briefly review the construction in [Culler and Vogtmann 1986] of outer space and its spine. A *graph* is a connected, one-dimensional CW-complex. We will consider only finite graphs with all vertices having valence at least 3. A *subforest* of a graph Γ is a subgraph of Γ that contains no circuits; a forest is a disjoint union of trees.

Fix an integer $n \geq 2$. Denote by R_0 the standard n-petal rose; R_0 has one vertex and n edges. Fix an identification $\pi_1(R_0) = F_n$. A *marking* on a graph is a homotopy equivalence, $g : R_0 \to \Gamma$. We define an equivalence relation on the set of markings by setting $(\Gamma_1, g_1) \sim (\Gamma_2, g_2)$ if

there is a graph isomorphism $h : \Gamma_1 \to \Gamma_2$ such that $g_2 \simeq h \circ g_1$, that is, such that the diagram

(1)
$$\begin{tikzcd} & \Gamma_1 \arrow[dd, "h"] \\ R_0 \arrow[ur, "g_1"] \arrow[dr, "g_2"'] & \\ & \Gamma_2 \end{tikzcd}$$

commutes up to free homotopy. An equivalence class of markings is called a *marked graph* and can be denoted by (Γ, g). The marking g identifies $\pi_1(\Gamma)$ with F_n up to composition with an inner automorphism.

The marked graph (Γ, g) is usually represented by a labeled graph as follows. Fix an identification of $\pi_1(R_0)$ with F_n. Choose a spanning tree T in Γ and a homotopy inverse to g that collapses T to the vertex of R_0 and maps each edge of $\Gamma - T$ to a reduced edge path in R_0. A directed edge $\vec{e}$ in the complement of T corresponds, via this homotopy equivalence, to an of element in F_n. Label e with a direction and the corresponding element of F_n. Note that the same marked graph can be represented by many different labeled graphs, depending on the choice of T and the particular representative of (Γ, g). For a marked rose, the spanning tree must consist of the single vertex, so we get a label for each directed edge. The set of labels on a marked rose is a basis of F_n, which is determined up to conjugacy. Two labeled roses correspond to equivalent marked roses if and only if their edges are labeled by conjugate bases of F_n.

If Φ is a forest in the marked graph (Γ, g), then collapsing each component of Φ to a point produces another marked graph, denoted by $(\Gamma/\Phi, q \circ g)$, where q is the quotient map collapsing each component of Φ to a point. Passing from (Γ, g) to $(\Gamma/\Phi, q \circ g)$ is called a *forest collapse*. There is a partial order on the set of marked graphs with fundamental group F_n defined by $(\Gamma_1, g_1) \leq (\Gamma_2, g_2)$ if there is a forest collapse taking (Γ_2, g_2) to (Γ_1, g_1). The geometric realization of the poset of marked graphs is the spine of outer space and is denoted by K_n.

The group $\mathrm{Out}(F_n)$ acts on K_n by changing the markings of the underlying graphs. Explicitly, for $\psi \in \mathrm{Out}(F_n)$,

$$(\Gamma, g) \cdot \psi := (\Gamma, g \circ |\psi|), \tag{2}$$

where $|\psi| : R_0 \to R_0$ is a homotopy equivalence inducing an automorphism of $F_n = \pi_1(R_0)$ that represents the outer automorphism class ψ. Culler and Vogtmann observe that this action is cocompact and that vertex stabilizers are finite.

Culler and Vogtmann also define a larger space, called *outer space*, consisting of metric marked graphs. This space has the disadvantages of not being a simplicial complex and the $\mathrm{Out}(F_n)$ action not being cocompact. The complex K_n can be constructed as a simplicial spine onto which of outer space deformation retracts.

3. Ribbon graphs

There are similar constructions for mapping class groups that use marked ribbon graphs rather than ordinary marked graphs. A *ribbon graph* is a graph Γ together with, at each vertex v, a cyclic ordering of the set $h(v)$ of half-edges incident to v. The collection of cyclic orderings at the vertices is called a *ribbon structure* for Γ, and is denoted by $\mathbb{O}$. The term "ribbon graph" is used because one can construct a bounded surface from a ribbon graph $(\Gamma, \mathbb{O})$ by fattening its edges to ribbons. We give a formal construction of this surface after Definition 3.1, but informally, the surface is constructed from $(\Gamma, \mathbb{O})$ by replacing each edge by a ribbon and gluing the ribbons together at their ends according to the cyclic order of the corresponding half-edges. The gluing is done in such a way as to produce an oriented surface. Figure 1 shows this process for two different ribbon structures on a rose with 2 edges. In these figures, ribbon structures are specified by the given embeddings of a neighborhood of the vertices into the plane. The ribbon graph $(\Gamma, \mathbb{O}_1)$ produces a pair of pants while $(\Gamma, \mathbb{O}_2)$ produces a torus with one boundary component.

The boundary curves of the surface produced from $(\Gamma, \mathbb{O})$ correspond to reduced edge paths in Γ that follow the cyclic ordering at the vertices in the sense of the following definition. Following [Mulase and Penkava 1998], we view a directed edge $\vec{e}$ as an ordering (e^+, e^-) of the half-edges e^+ and e^- comprising e.

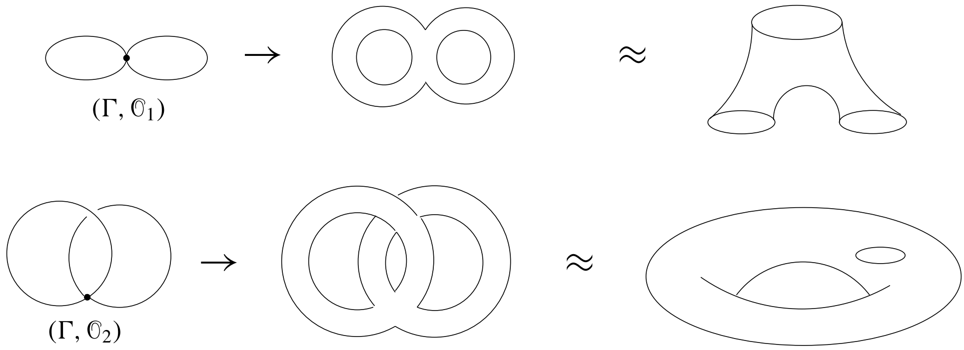

Figure 1. Fattenings of ribbon graphs.

Definition 3.1. A *boundary cycle* in the ribbon graph $(\Gamma, \mathbb{O})$ is a directed reduced edge cycle,

$$(\vec{e}_1, \vec{e}_2, \dots, \vec{e}_{l-1}, \vec{e}_l = \vec{e}_1)$$

such that for each i the half-edges e_i^+ and e_{i+1}^- are incident to the same vertex, and in the cyclic ordering at that vertex, e_{i+1}^- directly follows e_i^+.

For our purposes, it will be more convenient to work with punctured surfaces, so we now give a precise construction of a punctured surface $|\Gamma, \mathbb{O}|$ from a ribbon graph $(\Gamma, \mathbb{O})$. First note that each edge of Γ is traversed exactly once in each direction by the set of boundary cycles of $(\Gamma, \mathbb{O})$. Construct a space $|\Gamma, \mathbb{O}|$ by gluing a once punctured disk to Γ along each boundary cycle γ of $(\Gamma, \mathbb{O})$. By verifying that a small neighborhood of each vertex in $|\Gamma, \mathbb{O}|$ is indeed a disk, one can verify that $|\Gamma, \mathbb{O}|$ is a surface that deformation retracts onto Γ. One can also verify that $|\Gamma, \mathbb{O}|$ is orientable and we orient it such that a small positively oriented simple closed curve around a vertex v of Γ intersects the half-edges in $h(v)$ in the cyclic order determined by $\mathbb{O}$.

If Γ is marked by the homotopy equivalence $g : R_0 \to \Gamma$, then the composition of g with the inclusion $i : \Gamma \hookrightarrow |(\Gamma, \mathbb{O})|$ is a homotopy equivalence that identifies $\pi_1(\Sigma)$ with F_n up to inner automorphism, just as in the case of marked graphs. This gives the notion of a homotopy marked surface.

Definition 3.2. A *homotopy marked surface* is an equivalence class of pairs (Σ, s), where Σ is a punctured, orientable surface with $\pi_1(\Sigma) \cong F_n$ and $s : R_0 \to \Sigma$ is a homotopy equivalence. The equivalence relation on pairs is given by $(\Sigma_1, s_1) \sim (\Sigma_2, s_2)$ if there is an orientation preserving homeomorphism $h : \Sigma_1 \to \Sigma_2$ with $h \circ s_1 \simeq s_2$.

Recall that we have fixed an integer $n \geq 2$. Often we drop the word "homotopy" and simply use "marked surface" for a homotopy marked surface. Unless otherwise stated, marked surfaces will always be punctured surfaces without boundary and with fundamental group F_n. In Section 4, the equivalence relation defined by homeomorphisms that do not necessarily preserve orientation will be useful. We will denote this equivalence relation by $\sim_{\pm}$, and use brackets to denote its equivalence classes: $[\Sigma, s]$. We say that the marked graph (Γ, g) can be *drawn in* the marked surface (Σ, s) is there is a ribbon structure $\mathbb{O}$ on Γ such that $|(\Gamma, g, \mathbb{O})| \sim (\Sigma, s)$. In this case, there is an embedding $i : \Gamma \hookrightarrow \Sigma$ such that $s \simeq i \circ g$.

Definition 3.3. The *ribbon graph complex* for the marked surface (Σ, s) is the subcomplex of K_n spanned by graphs that can be drawn in (Σ, s). This complex is denoted by $\mathfrak{R}_{(\Sigma,s)}$.

We will often identify a marked graph or ribbon graph with the corresponding vertex of K_n or $\mathfrak{R}_{(\Sigma,s)}$. Thus for example, if ρ is a marked rose in $\mathfrak{R}_{(\Sigma,s)}$ then $\mathrm{lnk}_{\mathfrak{R}_{(\Sigma,s)}}(\rho)$ will be the link in $\mathfrak{R}_{(\Sigma,s)}$ of the vertex corresponding to ρ.

The ribbon graph complex $\mathfrak{R}_{(\Sigma,s)}$ and related complexes have been important tools in the study of mapping class groups surfaces. In particular, $\mathfrak{R}_{(\Sigma,s)}$ is a subcomplex of the first barycentric subdivision of the arc complex that Harer uses compute the VCD of the pure mapping class group of a surface with boundary [Harer 1986]. Also, for a punctured surface Σ, Bowditch and Epstein [1988] and Penner [1987] use arc systems on Σ to give an open cell decomposition of a space they call the decorated Teichmüller space of Σ. By taking the dual graph of an arc system in Σ, this decomposition may be interpreted in terms of metric ribbon graphs. In the same way that K_n is a simplicial spine of outer space, $\mathfrak{R}_{\Sigma}$ is a simplicial spine of the decorated Teichmüller space of Σ.

4. The ribbon cover of K_n

The ribbon subcomplex of K_n associated to a marked surface does not depend on the surface's orientation. This is because if the ribbon structure $\mathbb{O}$ draws (Γ, g) in (Σ, s), then $\mathbb{O}^{op}$ draws (Γ, g) in $(\Sigma, s)^{op}$, where $\mathbb{O}^{op}$ is the ribbon structure obtained by reversing all cyclic ordering of $\mathbb{O}$ and $(\Sigma, s)^{op}$ the marked surface obtained by reversing the orientation of (Σ, s). Therefore there is a well-defined subcomplex $\mathfrak{R}_{[\Sigma,s]}$ of K_n.

Proposition 4.1. *K_n is covered by its ribbon graph subcomplexes.*

Proof. Recall that K_n is the geometric realization of the poset of marked graphs with fundamental group F_n, so the vertices of K_n are partially ordered. For a vertex v of K_n, let $\mathrm{st}(v)$ be the star of v and let $\mathrm{st}^+(v)$ be the subcomplex of $\mathrm{st}(v)$ spanned by v together with vertices of $\mathrm{st}(v)$ that are greater than v in the partial order. Similarly let $\mathrm{st}^-(v)$ be the subcomplex of $\mathrm{st}(v)$ spanned by v and vertices less than v. Thus if v corresponds to the marked graph (Γ, g), then $\mathrm{st}^+(v)$ consists of the vertices of K_n corresponding to graphs that may be collapsed to (Γ, g) and $\mathrm{st}^-(v)$ consists of vertices corresponding to graphs to which (Γ, g) collapses.

Suppose that $v \in \mathfrak{R}_{[\Sigma,s]}$ corresponds to the marked graph (Γ, g). Then Γ has a ribbon structure $\mathbb{O}$ that draws (Γ, g) in (Σ, s). If e is any edge in Γ that is not a loop, the marked graph $(\Gamma/e, q \circ g)$ inherits a ribbon structure $\mathbb{O}/e$ from $\mathbb{O}$ that draws $(\Gamma/e, q \circ g)$ in (Σ, s). Therefore $\mathrm{st}^-(v) \subset \mathfrak{R}_{[\Sigma,s]}$.

To see that every simplex of K_n belongs to some ribbon graph subcomplex, let σ be a simplex of K_n. If w is the vertex of σ that is the greatest in the partial ordering of the vertices, then σ is contained in the complex $\mathrm{st}^-(w)$. Suppose that w corresponds to the marked graph (Γ_0, g_0) Choose any ribbon structure $\mathbb{O}_0$ on Γ_0 and set $(\Sigma, s) := |(\Gamma_0, g_0, \mathbb{O}_0)|$. Then $w \in \mathfrak{R}_{[\Sigma,s]}$ so that $\mathrm{st}^-(w)$ is contained in $\mathfrak{R}_{[\Sigma,s]}$. Since σ has w as its greatest vertex, σ is a simplex of $\mathrm{st}^-(w) \subseteq \mathfrak{R}_{[\Sigma,s]}$. $\square$

We begin our study of the nerve of this cover with definitions and lemmas.

Definition 4.2. Suppose that the homotopy-marked, oriented surface (Σ, s) has k punctures $p_1, \dots, p_k$. Let γ_j be a simple closed curve in Σ that disconnects Σ by cutting off a disk punctured at p_j. By virtue of the marking and orientation of Σ, the curve γ_j corresponds to a conjugacy class in F_n. The set of such conjugacy classes is called the set of *boundary classes* of Σ and is denoted by $W_{(\Sigma,s)}$ or simply W_Σ. Similarly, the set of conjugacy classes in F_n represented by the boundary cycles of the marked ribbon graph $(\Gamma, g, \mathbb{O})$ is called the set of *boundary classes* of $(\Gamma, g, \mathbb{O})$.

The boundary classes carry a lot of information about the surface. For example, if $(\Sigma, s) = |(\Gamma, g, \mathbb{O})|$ then the boundary classes of (Σ, s) and the boundary classes of $(\Gamma, g, \mathbb{O})$ are the same. Another important observation about the boundary classes is that if two (necessarily homeomorphic) marked surfaces have the same boundary classes, they are equivalent marked surfaces. This is proved by using a theorem of Zieschang [1980, Theorem 5.15.3] that states that an element of $\mathrm{Out}(\pi_1\Sigma)$ is induced by a mapping class of Σ if and only if it stabilizes the boundary classes of Σ. The boundary classes are also used to prove the following:

Lemma 4.3. *If the marked graph (Γ, g) can be drawn in (Σ, s), then (Γ, g) has exactly one ribbon structure giving (Σ, s).*

Proof. Since Γ can be drawn in (Σ, s), there is a ribbon structure $\mathbb{O}$ on (Γ, g) with $|(\Gamma, g, \mathbb{O})| = (\Sigma, s)$. Suppose that $\mathbb{O}'$ is a different ribbon structure on Γ. We may choose a vertex v and half-edges e^+, e_1^- and e_2^- of Γ with $e_1^- \neq e_2^-$, with e_1^- following e^+ in the cyclic ordering $\mathbb{O}$ but with e_2^- following e^+ in the cyclic ordering $\mathbb{O}'$. This means that the sequence $\vec{e}\vec{e}_1$ appears in the boundary cycles of $(\Gamma, \mathbb{O})$ while the sequence $\vec{e}\vec{e}_2$ appears in the boundary classes of $(\Gamma, \mathbb{O}')$. Since each directed edge of Γ appears exactly once in the set of boundary cycles for any given ribbon structure, the set of boundary cycles of $(\Gamma, g, \mathbb{O}')$ must differ from those of $(\Gamma, g, \mathbb{O})$. Therefore the set of boundary classes of $(\Gamma, g, \mathbb{O}')$ differ from those of $(\Gamma, g, \mathbb{O})$. Hence $|(\Gamma, g, \mathbb{O}')|$ is different from (Σ, s), because equivalent marked surfaces have the same boundary classes. □

Each of the two orientations of Σ gives a unique ribbon structure to (Γ, g). These two ribbon structures are opposite of each other. We now prove that (Σ, s) and (Σ_1, s_1) give the same ribbon graph subcomplexes of K_n if and only if $[\Sigma, s] = [\Sigma_1, s_1]$. The main step is this:

Lemma 4.4. *Let (Σ, s) be a marked surface and let $\mathfrak{R} = \mathfrak{R}_{[\Sigma,s]}$ be the corresponding ribbon graph subcomplex of K_n. The ribbon structure given by (Σ, s) to a marked rose ρ in $\mathfrak{R}$ can be reconstructed, up to reversal of the cyclic order, by the (nonribbon) marked graphs in* $\mathrm{lnk}_{\mathfrak{R}}(\rho)$.

Proof. Choose a direction for each edge in ρ. The marking of ρ determines a labeling of the directed edges by a basis $X = \{a_1, a_2, \ldots, a_n\}$ of F_n. The basis X and this labeling are determined up to composition with an inner automorphism of F_n.

Consider the marked graphs in $\mathrm{lnk}_{K_n}(\rho)$ with exactly two vertices, one of which is trivalent. These graphs are constructed from ρ as follows. Let e^+ and f^+ be any two half-edges of ρ. Construct a new marked graph $\rho(e^+, f^+)$ by deleting the vertex of ρ and replacing it with two new vertices v_0 and v_1 joined by a new edge $\tilde{e}$. Attach the half-edges e^+ and f^+ to v_0 and attach the rest of the half-edges of ρ to v_1. Mark $\rho(e^+, f^+)$ so that collapsing $\tilde{e}$ to a point gives the original marking on ρ.

By Lemma 4.3, ρ has exactly one ribbon structure giving (Σ, s). This means that if we allow for orientation reversing homeomorphisms of the surface, ρ has two ribbon structures giving $[\Sigma, s]$, and these ribbon structures are opposite of each other. A marked graph of the form $\rho(e^+, f^+)$ lies in $\mathrm{lnk}_{\mathfrak{R}}(\rho)$ if and only if e^+ and f^+ are adjacent in these ribbon structures. Thus, given a half-edge a_i^+ in ρ, exactly two graphs of the form $\rho(a_i^+, a_j^{\epsilon_j})$ and $\rho(a_i^+, a_k^{\epsilon_k})$ will lie in $\mathrm{lnk}_{\mathfrak{R}}(\rho)$ and they will be the graphs for which a_i^+ is adjacent to the half-edges $a_j^{\epsilon_j}$ and $a_k^{\epsilon_k}$ in the ribbon structure on ρ. Therefore, for each i, the nonribbon graphs in $\mathrm{lnk}_{\mathfrak{R}}(\rho)$ determine the half-edges adjacent to a_i^+ and a_i^- in the ribbon structure on ρ. There are only two cyclic orderings of the half-edges that satisfy these adjacency data, and they are opposites of each other. Example 4.5 works this out for a ribbon rose with $n = 3$. □

Example 4.5. Figure 2 shows the graphs of the form $\rho(e^+, f^+)$ in the link of a marked ribbon rose in $\mathfrak{R}_\Sigma$. The fact that graphs (1) and (6) have the half-edges a^+ and c^+, respectively, adjacent to the edge a^- implies that the half-edges in

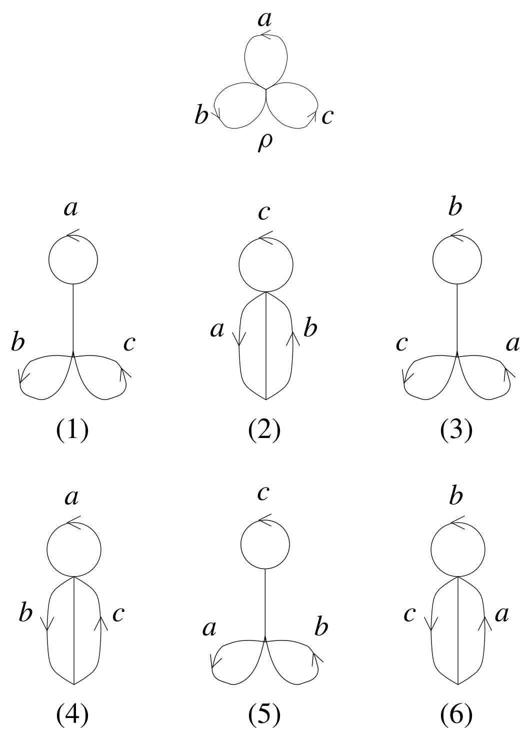

Figure 2. Some graphs in $\mathrm{lnk}^+_{\mathfrak{R}_\Sigma}(\rho)$.

ρ adjacent to a^- are a^+ and c^+. Similarly, the half-edges adjacent to any other half-edge can be determined by some pair of the graphs in Figure 2.

Proposition 4.6. $\mathfrak{R}_{[\Sigma_1,s_1]} = \mathfrak{R}_{[\Sigma_2,s_2]}$ *if and only if* $[\Sigma_1, s_1] = [\Sigma_2, s_2]$.

Proof. First suppose that (Σ_1, s_1) and (Σ_2, s_2) are equivalent via the (possibly orientation-reversing) homeomorphism $h : \Sigma_1 \to \Sigma_2$. Then h can be used to draw in (Σ_2, s_2) any graph that can be drawn in (Σ_1, s_1) and h^{-1} can be used to draw in (Σ_1, s_1) any graph that can be drawn in (Σ_2, s_2), so $\mathfrak{R}_{[\Sigma_1,s_1]} = \mathfrak{R}_{[\Sigma_2,s_2]}$.

Now suppose that $\mathfrak{R}_{[\Sigma_1,s_1]} = \mathfrak{R}_{[\Sigma_2,s_2]}$. Set $\mathfrak{R} := \mathfrak{R}_{[\Sigma_1,s_1]} = \mathfrak{R}_{[\Sigma_2,s_2]}$. Fix a marked rose $\rho \in \mathfrak{R}$; it inherits a ribbon structure from $\mathfrak{R}_{[\Sigma_1,s_1]}$ giving (Σ_1, s_1), and a ribbon structure from $\mathfrak{R}_{[\Sigma_2,s_2]}$ giving (Σ_2, s_2). By Lemma 4.4, these structures are determined up to reversal by the nonribbon graphs in $\mathrm{lnk}_{\mathfrak{R}_{[\Sigma_1,s_1]}}(\rho)$ and $\mathrm{lnk}_{\mathfrak{R}_{[\Sigma_2,s_2]}}(\rho)$, respectively. But $\mathfrak{R}_{[\Sigma_1,s_1]} = \mathfrak{R}_{[\Sigma_2,s_2]}$,so $\mathrm{lnk}_{\mathfrak{R}_{[\Sigma_1,s_1]}}(\rho) = \mathrm{lnk}_{\mathfrak{R}_{[\Sigma_2,s_2]}}(\rho)$. Therefore the ribbon structures must coincide or be opposites of each other. In the first case, $(\Sigma_1, s_1) \sim (\Sigma_2, s_2)$, and in the second, $(\Sigma_1, s_1) \sim (\Sigma_2, s_2)^{op}$. Thus $[\Sigma_1, s_1] = [\Sigma_2, s_2]$. $\square$

This proposition gives a convenient description of the covering of K_n by its ribbon graph subcomplexes. The covering is locally finite because each different homotopy marked surface that contains a specific graph endows that graph with a different ribbon structure. A graph has only finitely many different ribbon structures so a given marked graph can be drawn in only finitely many marked surfaces and hence lies in only finitely many different ribbon graph subcomplexes.

Let $\mathcal{N}_n$ denote the nerve of the ribbon cover of K_n. That is, $\mathcal{N}_n$ is the simplicial complex containing a k-simplex $\langle \mathfrak{R}_{[\Sigma_0,s_0]}, \dots, \mathfrak{R}_{[\Sigma_k,s_k]}\rangle$ for every collection $\{\mathfrak{R}_{[\Sigma_0,s_0]}, \dots, \mathfrak{R}_{[\Sigma_k,s_k]}\}$ of ribbon graph complexes such that the intersection $\bigcap_{i=0}^k \mathfrak{R}_{[\Sigma_i,s_i]}$ is nonempty. By Proposition 4.6, the vertex set of $\mathcal{N}_n$ is the set of unoriented equivalence classes, $[\Sigma, s]$.

The action of $\mathrm{Out}(F_n)$ on K_n permutes the ribbon graph subcomplexes because if (Γ, g) can be drawn in (Σ, s), then $(\Gamma, g) \cdot \psi = (\Gamma, g \circ |\psi|)$ can be drawn in $(\Sigma, s \circ |\psi|)$. Therefore $\mathrm{Out}(F_n)$ maps intersections of ribbon graph subcomplexes to intersections of ribbon graph subcomplexes, so it acts on $\mathcal{N}_n$. The equivariant homology of this action provides the spectral sequence, which we will study, that relates the homology of $\mathrm{Out}(F_n)$ to that of mapping class groups.

Although it will not be necessary for the development here, we remark briefly on the compactness properties of $\mathcal{N}_n$ and the $\mathrm{Out}(F_n)$ action. In general ($n \geq 3$), all vertices of $\mathcal{N}_n$ have infinite valence:

Proposition 4.7. *For $n \geq 3$, the ribbon complex for any homotopy marked surface intersects the ribbon complexes of infinitely many other homotopy marked surfaces.*

Proof. We first show that the ribbon graph subcomplex of a marked surface $[\Sigma, s]$ with fundamental group of rank at least 3, contains infinitely many different marked roses. Choose a marked rose $(\rho, r) \in \mathfrak{R}_{[\Sigma,s]}$ and an automorphism $\psi \in \mathrm{Out}(F_n)$ representing a Dehn twist about a nonboundary curve in Σ. Since vertex stabilizers in the spine of outer space are finite and ψ has infinite order, there are infinitely many different equivalence classes of marked roses of the form $\psi^n \cdot (\rho, r)$. All of these marked roses lie in $\mathfrak{R}_{[\Sigma,s]}$.

Recall from Section 2 that a marking of a rose is equivalent to a choice of conjugacy class of basis labeling its directed edges. If ρ is a marked rose with edges labeled by the basis $X = \{a_1, \dots, a_n\}$, then ρ can be drawn in a marked $(n+1)$-times punctures sphere Σ_1 with boundary classes

$$W_{\Sigma_1} = \{a_1, \dots, a_n, a_n^{-1}, \cdots, a_1^{-1}\}.$$

By the discussion following Definition 4.2, two marked spheres with different boundary classes cannot be equivalent. Therefore the infinitely many different marked roses in $\mathfrak{R}_{[\Sigma,s]}$ give rise to infinitely many different marked spheres all of whose ribbon complexes intersect $\mathfrak{R}_{[\Sigma,s]}$. □

Proposition 4.8. $\mathrm{Out}(F_n)$ *acts cocompactly on* $\mathcal{N}_n$.

Proof. Fix a marked rose ρ. For each p-simplex $\langle \Sigma_0, \dots, \Sigma_p \rangle$, the subcomplex $\bigcap_{i=0}^{p} \mathfrak{R}_{\Sigma_i}$ contains a rose. This rose may be taken to ρ by an element of $\mathrm{Out}(F_n)$, so each orbit of p-simplex has a representative all of whose surfaces contain ρ. Since a marked rose can be drawn in only finitely many different marked surfaces, there are only finitely many orbits of p-simplices. □

5. Contractibility of $\mathcal{N}_n$

To show that $\mathcal{N}_n$ is contractible, we will need the following result from Čech theory.

Lemma 5.1 [Hatcher 2002, Section 4.G]. *Let* $\mathfrak{U}$ *be a cover of the CW-complex* X *by a family of subcomplexes. If every nonempty intersection of finitely many complexes in* $\mathfrak{U}$ *is contractible, then the nerve of the cover is homotopy equivalent to* X.

We will apply Lemma 5.1 to the covering of K_n by ribbon graph complexes. Thus the remainder of this section is devoted to the proof of,

Proposition 5.2. *For any finite collection* $\{\mathfrak{R}_{\Sigma_0}, \dots, \mathfrak{R}_{\Sigma_k}\}$ *of ribbon graph subcomplexes of* K_n, *the subcomplex*

$$\bigcap_{i=0}^{k} \mathfrak{R}_{\Sigma_i}$$

of K_n *is either empty or contractible.*

The main tool in analyzing these intersections is the $K_{\min}$ subcomplexes of K_n, which are used by Culler and Vogtmann [1986] to show that K_n is contractible. The definition of the $K_{\min}$ complexes involves the following norm defined for each finite set of conjugacy classes of F_n. Let $\mathscr{C}$ denote the set of all conjugacy classes of F_n. For a marked rose ρ, an element $w \in \mathscr{C}$ can be represented by a unique reduced edge path in ρ.

Definition 5.3. Let W be a finite set of conjugacy classes of F_n and ρ a marked rose in K_n. The *norm* $\|\rho\|_W$ of ρ with respect to W is the sum of the number of edges in each reduced edge path in ρ that corresponds to an element of W.

If X is a basis labeling the edges of ρ, then $\|\rho\|_W$ is sometimes written $\|X\|_W$. The $K_{\min}$ subcomplex for W is defined as the union of the stars of the roses ρ for which $\|\rho\|_W$ is minimal over all marked roses. In [Vogtmann 2002] these complexes are denoted K_W, and we will follow that notation here. To prove that the entire complex K_n is contractible, Culler and Vogtmann prove that $K_W \simeq K_n$ for any finite set W. They then find a set of conjugacy classes such that K_W is the star of a single marked rose and therefore contractible. Putting these two facts together we have:

Lemma 5.4 [Culler and Vogtmann 1986]. *K_W is contractible for any finite set $W \subseteq \mathscr{C}$.*

Proposition 5.2 is proved by finding a deformation retraction from a suitable K_W to $\bigcap \mathfrak{R}_{\Sigma_i}$. We begin by studying of the behavior of the norm with respect to Whitehead automorphisms. For us, the traditional Whitehead automorphisms are less convenient to work with than a slightly modified version, given in [Hoare 1979]. This is because the effect of an automorphism on the star graph of a set of conjugacy classes (defined below), is easier to describe using this definition rather than the classical definition of Whitehead automorphism.

Definition 5.5. For a basis X and a subset $A \subseteq X \cup X^{-1}$ for which there is a letter $a \in X \cup X^{-1}$ such that $a \in A$ but $a^{-1} \notin A$, the automorphism mapping a to a^{-1} whose action on $X \cup X^{-1} - \{a, a^{-1}\}$ is given by

$$\begin{cases} x \mapsto axa^{-1} & \text{if } x \in A \text{ and } x^{-1} \in A; \\ x \mapsto xa^{-1} & \text{if } x \in A \text{ and } x^{-1} \notin A; \\ x \mapsto ax & \text{if } x \notin A \text{ and } x^{-1} \in A; \\ x \mapsto x & \text{if } x \notin A \text{ and } x^{-1} \notin A \end{cases} \tag{3}$$

will be called a *Whitehead automorphism* and will be denoted by (A, a).

Warning. This definition differs from the classical Whitehead automorphism in that the latter fix a. This is the only difference, but it allows us to prove the next result.

Lemma 5.6. *The totality of Whitehead automorphisms obtained as in the preceding definition generate the group* $\operatorname{Aut}(F_n)$.

Proof. The Neilson automorphisms generate $\operatorname{Aut}(F_n)$ [Magnus et al. 1966, Theorem 3.2]. It is straightforward to write any Neilson automorphism as a product of the Whitehead automorphisms of Definition 5.5. □

Another important fact about Whitehead automorphisms of this type is the peak reduction lemma (for Whitehead automorphisms as defined here).

Peak reduction lemma [Hoare 1979, Lemma 3]. *Fix a basis X of F_n and finite set $W \subseteq \mathscr{C}$. If there is an automorphism $\psi \in \operatorname{Aut}(F_n)$ such that $\|X\|_W \geq \|X\|_{\psi W}$ then ψ can be written as a product $\psi = \tau_1 \tau_2 \cdots \tau_k$ of Whitehead automorphisms such that*

$$\|X\|_W > \|X\|_{\tau_k W} > \|X\|_{\tau_{k-1}\tau_k W} > \cdots > \|X\|_{\tau_l \tau_{l+1} \cdots \tau_k W} = \|X\|_{\tau_{l-1}\tau_l \tau_{l+1} \cdots \tau_k W} = \cdots = \|X\|_{\psi W}. \tag{4}$$

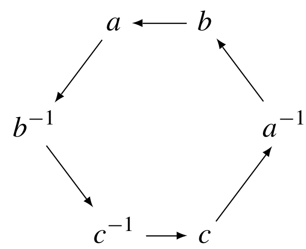

Figure 3. $S_W(X)$ for $X = \{a,\ b,\ c\}$, $W = \{aba^{-1}b^{-1}c,\ c^{-1}\}$.

The star graph of W with respect to X will allow us to study the behavior of $\|\cdot\|_W$ with respect to Whitehead automorphisms. Recall that star graph of $W \subseteq \mathscr{C}$ with respect to the basis X is the graph with vertex set $X \cup X^{-1}$ and with a directed edge from x to y^{-1} for every time the subword xy appears among the conjugacy classes in W, viewed as cyclic words in the alphabet $X \cup X^{-1}$; see Figure 3. The star graph of W with respect to X will be denoted by $S_W(X)$, or by $S_W(\rho)$ if we are thinking of X as a set of labels on the marked rose ρ.

To prove the peak reduction lemma, Hoare describes a three-step process for constructing $S_{\tau W}(X)$ from $S_W(X)$ for a Whitehead automorphism τ. If $\tau = (A, a)$, the steps are:

(1) Add two new vertices $\alpha, \bar{\alpha}$. Replace every edge going from a vertex in A to a vertex in A' (the complement of A) by a pair of edges, one from the vertex in A to α and another from $\bar{\alpha}$ to the vertex in A'. Replace every edge going from a vertex in A' to a vertex in A by a pair of edges, one from the vertex in A' to $\bar{\alpha}$ and another from α to the vertex in A.

(2) Switch the letter a with α, and a^{-1} with $\bar{\alpha}$.

(3) Do the reverse of (1), reconnecting edges incident to α and $\bar{\alpha}$ according to the cyclic words in W that produced them.

Because $\|X\|_W$ is the number of edges in $S_W(X)$, this process gives the following procedure for calculating the effect of a Whitehead automorphism on the norm. Consider the Whitehead automorphism $\tau = (A, a)$. Draw a circle C in the plane and immerse $S_W(X)$ in the plane in such a way that each vertex of A lies inside the circle, each vertex of A' lies outside the circle, no pair of edges of $S_W(X)$ intersect each other at a point of C, and $\#(S_W(X) \cap C)$ is minimal over all such immersions. Then

$$\|X\|_W - \|X\|_{\tau W} = \operatorname{val}(a) - \#(S_W(X) \cap C), \tag{5}$$

where $\operatorname{val}(a)$ is the valence of the vertex of $S_W(X)$ corresponding to a.

When $W = W_\Sigma$ is the set of boundary classes of a marked surface, we will be more concerned with $\|X\|_W - \|\tau^{-1}X\|_W$ than with $\|X\|_W - \|X\|_{\tau W}$, because if X is the set of labels on the edges of a marked rose ρ, then $\tau^{-1}X$ is the set of labels on $\rho \cdot \tau$. (5) will suffice because

$$\|\tau^{-1}X\|_W = \|X\|_{\tau W};$$

the latter equality can be seen from the observation that if $\tau(a) = x_1 x_2 \cdots x_k$ is an expression for $\tau(a)$ in terms of the basis X, then $a = \tau^{-1}(x_1)\tau^{-1}(x_2)\cdots\tau^{-1}(x_k)$ is an expression for a in terms of the basis $\tau^{-1}X$. The interpretation of this observation in terms of star graphs is

$$S_W(\tau^{-1}X) \approx S_{\tau W}(X). \tag{6}$$

Lemma 5.7. *Let W be a finite set of conjugacy classes of F_n. Suppose that for some basis X, the graph $S_W(X)$ is a cycle. Then $\|X\|_W$ is minimal over all bases of F_n, and if Y is another basis with $\|Y\|_W = \|X\|_W$, then $S_W(Y)$ is also a cycle.*

Proof. Since $S_W(X)$ is a cycle, any circle separating some generator from its inverse must intersect at least two edges of the graph. Since all vertices have valence 2, Equations (5) and (6) imply that no Whitehead automorphism can take X to a basis that reduces the sum of the lengths of the minimal representatives for the classes in W. But if there is any automorphism reducing the sum of the lengths of the conjugacy classes in W, the peak reduction lemma implies that a Whitehead automorphism reduces the length. Thus $\|X\|_W$ must be minimal over all bases for F_n.

Now suppose that Y is another basis with $\|Y\|_W = \|X\|_W$. Since $\mathrm{Aut}(F_n)$ acts transitively on bases of F_n, we may choose $\psi \in \mathrm{Aut}(F_n)$ with $Y = \psi^{-1}X$. By the peak reduction lemma and Equation (6), there is a sequence $\tau_1, \ldots, \tau_l$ of Whitehead automorphisms such that $\psi = \tau_1, \ldots, \tau_l$ and $\|\tau_l^{-1}\tau_{l-1}^{-1}\cdots\tau_i^{-1}X\|_W = \|X\|_W$ for $i = 1, \ldots, l$. Thus without loss of generality, we may assume that $\psi = \tau$ is a Whitehead automorphism.

We now use Hoare's method to construct the star graph $S_{\tau W}(X) \approx S_W(\tau^{-1}X)$. Since $\|\tau^{-1}X\|_W = \|X\|_W$, the circle separating A from A' in the star graph must intersect only two edges of the graph; otherwise the norm would increase. Therefore, the subgraph spanned by the vertices of A is a simple path, and the same is true for the subgraph spanned by A'. Step (1) of Hoare's procedure produces a graph consisting of two disjoint cycles, one containing α and the other containing $\bar{\alpha}$. Step (2) keeps α and $\bar{\alpha}$ in separate cycles, but they may switch cycles. Step (3) breaks these two cycles at α and $\bar{\alpha}$, and reconnects the ends of the resulting line segments to form a cycle, which is $S_{\tau W}(X)$. Since $S_W(Y) = S_W(\tau^{-1}X) \approx S_{\tau W}(X)$, $S_W(Y)$ is a cycle. ∎

We will use Lemma 5.7 to analyze which marked graphs can be drawn in a particular surface (Σ, s). If Σ has genus g and s punctures, and if

$$X = \{a_1, b_1, a_2, b_2, \ldots, a_{2g}, b_{2g}, c_1, \ldots, c_{s-1}\}$$

is a standard, geometric basis of $F_n = \pi_1(\Sigma)$, we have

$$W_{(\Sigma,s)} = \{[a_1, b_1][a_2, b_2] \cdots [a_{2g}, b_{2g}]c_1 \cdots c_{s-1},\ c_1^{-1},\ c_2^{-1},\ \ldots,\ c_{s-1}^{-1}\}.$$

Thus $S_W(X)$ is a cycle, and as a consequence of Lemma 5.7 we have:

Corollary 5.8. *If W is the set of boundary classes of a surface Σ, then*

$$\min_{\rho} \|\rho\|_W = 2n$$

and $S_W(\rho')$ is a cycle for any rose ρ' minimizing $\|\cdot\|_W$.

The next two lemmas characterize the marked graphs that lie in $\mathfrak{R}_{[\Sigma,s]}$.

Lemma 5.9. *The marked graph $\Gamma = (\Gamma, g)$ can be drawn in (Σ, s) if and only if the set of reduced edge cycles of Γ representing W_Σ traverses each edge of Γ exactly once in each direction.*

Proof. Suppose that (Γ, g) can be drawn in (Σ, s). Cutting Σ along Γ produces a collection of punctured disks, one for each puncture. The oriented boundaries of these disks correspond to the conjugacy classes of the boundary of Σ. Together they traverse each edge of Γ once in each direction. Thus, if (Γ, g) can be drawn in (Σ, s), the set of reduced edge cycles of Γ representing the boundary classes of (Σ, s) traverses each edge exactly once in each direction.

For the converse, we first construct another marked surface Σ' whose set of oriented boundary classes is also W_Σ, by showing that the boundary cycles in Γ induce a ribbon structure. We then use this surface to draw Γ in Σ. Let v be a vertex of Γ and define a polycyclic order on the half-edges at v by declaring that b^- follows a^+ if ab appears in the reduced boundary cycles in W_Σ. This definition may give more than one cycle of half-edges at some vertices, so it may not provide a cyclic order at each vertex. To show that it is indeed a cyclic order, we work by induction on the number of vertices in Γ.

If Γ has one vertex, then Γ is a rose and $\|\Gamma\|_{W_\Sigma} = 2n$. Since W_Σ is the set of boundary classes for a surface, Corollary 5.8 implies that the star graph $S_{W_\Sigma}(\Gamma)$ is a single connected cycle, which means that there is only one cycle of half-edges at the vertex of Γ, and our definition gives a cyclic ordering. Now, suppose that Γ has k vertices and assume by induction that any graph with fewer than k vertices and with the boundary classes traversing each edge exactly once in each direction has a single cycle at each vertex. Choose an edge e of Γ that is not a loop, and collapse it to obtain the marked graph $\Gamma' = \Gamma/e$. The reduced edge cycles of Γ'

representing the elements of W_Σ traverse each edge exactly once in each direction, and Γ' has $k-1$ vertices. Therefore Γ' has one cycle at each vertex. If v_1 and v_2 are the two vertices coalesced to the vertex $v \in \Gamma'$ during the collapse of edge e, then there is a single cycle at every vertex of Γ other than v_1 and v_2. A priori, the cycles at v can be formed by taking the cycles v_1 and v_2 and combining the one containing e^+ with the one containing e^-. Since there is only one cycle at v, each of v_1 and v_2 must possess only one cycle, which finishes the induction step that we have a cyclic ordering of the half-edges at each vertex.

Let $\mathbb{O}$ be the ribbon structure just constructed. Both Σ and the marked surface $\Sigma' = |(\Gamma, g, \mathbb{O})|$ are orientable surfaces with the same number of punctures and the same fundamental group; hence they are homeomorphic. By the construction of $\mathbb{O}$, the set of boundary classes of Σ' is W_Σ. Label each puncture of Σ and Σ' with the corresponding conjugacy class of W_Σ and choose a homeomorphism $f : \Sigma' \to \Sigma$ that preserves the labels of the punctures. Now, $f \circ i$ embeds Γ into Σ as a strong deformation retract, but this embedding may not induce the same marking as g. That is to say, the diagram

$$\begin{tikzcd} \Gamma \arrow[r, hook, "i"] & \Sigma' \arrow[d, "f"] \\ R_0 \arrow[u, "g"] \arrow[r, "s"'] & \Sigma \end{tikzcd}$$

may not commute up to homotopy. However, the outer automorphism given by $f_* \circ i_* \circ g_* \circ (s_*)^{-1}$ stabilizes W_Σ. By [Zieschang et al. 1980, Theorem 5.15.3], it is induced by an element θ in the orientation-preserving mapping class group of Σ. Embedding Γ into Σ by $\theta^{-1} \circ f \circ i$ gives the same marking as g. Thus Γ can be drawn in Σ. $\square$

Lemma 5.10. *Let (Σ, s) be a homotopy marked surface and ρ a marked rose. Then $\rho \in K_{W_\Sigma}$ if and only if $\rho \in \mathfrak{R}_\Sigma$.*

Proof. By Corollary 5.8, the minimal value of $\|\cdot\|_{W_\Sigma}$ is $2n$. Any $\rho \in \mathfrak{R}_\Sigma$ can be drawn in Σ. By cutting Σ along ρ, we see that $\|\rho\|_{W_\Sigma} = 2n$, so $\rho \in K_{W_\Sigma}$.

Conversely, suppose that $\rho \in K_{W_\Sigma}$. Since ρ minimizes $\|\cdot\|_{W_\Sigma}$, Corollary 5.8 implies that the star graph $S_{W_\Sigma}(\rho)$ must be a cycle. Therefore each label in ρ appears exactly once with exponent $+1$ and once with exponent -1 in the minimal expressions for conjugacy classes of W_Σ in terms of a set of labels of ρ. This means that the set of reduced edge cycles in ρ that represents W_Σ traverses each edge of ρ exactly once in each direction. By Lemma 5.9, ρ can be drawn in (Σ, s). $\square$

This lemma implies that the roses in $\mathfrak{R}_\Sigma$ coincide with those in K_{W_Σ}. Since K_{W_Σ} is the union of the stars of its roses, $\mathfrak{R}_\Sigma \subseteq K_{W_\Sigma}$. To find a graph in $\mathfrak{R}_\Sigma$ lying near a particular graph in $K_{W_\Sigma} - \mathfrak{R}_\Sigma$, we use the following lemma.

Lemma 5.11. *If* $\Gamma = (\Gamma, g) \in K_{W_\Sigma}$, *there exists a (possibly empty) forest* $\Phi_\Sigma(\Gamma)$ *such that for any forest* $\Phi \subseteq \Gamma$,

$$\Gamma/\Phi \in \mathfrak{R}_\Sigma \iff \Phi \supseteq \Phi_\Sigma(\Gamma).$$

Proof. Let $\Phi_\Sigma(\Gamma)$ be the subgraph of Γ consisting of all the edges of Γ that are not traversed exactly once in each direction by the set of reduced edge paths representing the boundary classes of Σ. Since $\Gamma \in K_{W_\Sigma}$, and K_{W_Σ} is the union of the stars of its roses, there is a maximal tree T in Γ such that Γ/T is a rose in K_{W_Σ}. By Lemma 5.10 this rose is in $\mathfrak{R}_\Sigma$, so it can be drawn in Σ. Therefore every edge of $\Gamma - T$ is traversed exactly once in each direction by the set of conjugacy classes in W_Σ. This means that $\Phi_\Sigma(\Gamma) \subseteq T$, so that $\Phi_\Sigma(\Gamma)$ is a forest.

Given any forest Φ in Γ, Lemma 5.9 implies that Γ/Φ can be drawn in Σ exactly when the boundary cycles traverse each edge of Γ/Φ once in each direction. This happens exactly when $\Phi_\Sigma(\Gamma) \subseteq \Phi$. □

These lemmas would allow us, at this time, to define a retraction from K_{W_Σ} to $\mathfrak{R}_\Sigma$ by taking a graph $\Gamma \in K_W$ to $\Gamma/\Phi_\Sigma(\Gamma)$ proving the following well-known proposition without having to appeal to the contractibility of Teichmüller space or the identification of the ribbon graph complex with the decorated Teichmüller space.

Proposition 5.12. *For any marked surface* $[\Sigma, s]$, *the ribbon graph complex* $\mathfrak{R}_{[\Sigma,s]}$ *is contractible.*

We postpone this proof until it is covered by the proof of contractibility for arbitrary simplices of the nerve. For higher-dimensional simplices, we need a set of conjugacy classes that captures the properties of a graph that can be drawn in several different surfaces. This set emphasizes a conjugacy class according to the number of the surfaces in question of which it is a boundary. We start by describing some general properties of collections of finite sets of conjugacy classes of F_n. For the proof of Proposition 5.2, we will specialize to the case that the sets of conjugacy classes are actually the boundary classes of marked surfaces.

Definition 5.13. For a collection $\sigma = \{W_0, \dots, W_k\}$ of finite sets of conjugacy classes of F_n, define

$$W_\sigma := \left\{[\alpha_1]^{n_1}, \dots, [\alpha_l]^{n_l}\right\},$$

where $\bigcup_{i=0}^k W_i = \{[\alpha_1], \dots, [\alpha_l]\}$, and n_j is the number of times that the conjugacy class $[\alpha_j]$ appears in the W_i.

Note that $[\alpha]$ and $[\alpha^{-1}]$ both may appear in W_σ. We use the letter σ for the set $\{W_0, \dots, W_k\}$ because this definition will be applied to a simplex $\sigma = \langle\Sigma_0, \dots, \Sigma_k\rangle$ of $\mathcal{N}$, with $W_i = W_{\Sigma_i}$. We will use the notation W_σ in this situation as well.

Lemma 5.14. *For σ and W_σ as above, let*

$$A = \min_\rho \|\rho\|_{W_\sigma} \quad \textit{and} \quad A_i = \min_\rho \|\rho\|_{W_i}.$$

Then $A = A_0 + \cdots + A_k$ if and only if $\bigcap_{i=0}^k K_{W_i} \neq \varnothing$.

Proof. Suppose that $W_\sigma = \{w_1^{n_1}, \ldots, w_l^{n_l}\}$. For any marked rose ρ,

$$\|\rho\|_{W_\sigma} = \sum_{i=0}^{l} n_i \|\rho\|_{\{w_i\}} = \sum_{j=0}^{k} \|\rho\|_{W_j}. \tag{7}$$

Choose any marked rose ρ_1 with $\|\rho_1\|_{W_\sigma} = A$. Now, ρ_1 may not minimize every $\|\cdot\|_{W_i}$, so

$$A_0 + \cdots + A_k \leq \sum_{j=0}^{k} \|\rho_1\|_{W_j} = \|\rho_1\|_{W_\sigma} = A, \tag{8}$$

where the first equality comes from Equation (7).

If $\bigcap_{i=0}^k K_{W_i}$ is nonempty, there is a single marked rose ρ_2 with $\|\rho_2\|_{W_i} = A_i$ for all i. Thus

$$A \leq \|\rho_2\|_{W_\sigma} = \sum_{i=0}^{k} \|\rho_2\|_{W_i} = A_0 + \cdots + A_k.$$

Together with (8) this implies that $A = A_0 + \cdots + A_k$.

Conversely, if $A = A_0 + \cdots + A_k$, then using the ρ_1 from above we have

$$A_0 + \cdots + A_k = A = \|\rho_1\|_{W_\sigma} = \sum_{j=0}^{k} \|\rho_1\|_{W_j}. \tag{9}$$

Again the last equality comes from (7). Now, $\|\rho_1\|_{W_i} \geq A_i$, so by (9) we have $\|\rho_1\|_{W_i} = A_i$ for each i. Hence, $\rho_1 \in K_{W_i}$ for each i, and $\bigcap K_{W_i} \neq \varnothing$. □

Changing the viewpoint slightly we get:

Corollary 5.15. *For any finite collection of finite sets of conjugacy classes, $\sigma = \{W_0, \ldots, W_k\}$, $K_{W_\sigma} = \bigcap_{i=0}^k K_{W_i}$ if the right-hand side is nonempty.*

Setting $k = 0$ here provides a proof of Proposition 5.12. The final lemma we need for the proof of Proposition 5.2 is this:

Poset lemma [Quillen 1973]. *Let $f : P \to P$ be a poset map from the poset P to itself such that $p \leq f(p)$ for all $p \in P$. Then f induces a deformation retraction from the geometric realization of P to the geometric realization of its image $f(P)$.*

Proof of Proposition 5.2. Let $\sigma = \langle \Sigma_0, \dots, \Sigma_k \rangle$ be a simplex of $\mathcal{N}_n$. Denote by $\mathfrak{R}_\sigma$ the intersection

$$\mathfrak{R}_\sigma := \bigcap_{i=0}^{k} \mathfrak{R}_{[\Sigma_i, s_i]},$$

and let W_σ be the set of conjugacy classes given by Definition 5.13. Now $\mathfrak{R}_\sigma$ contains a rose, since $\langle \Sigma_0 \dots \Sigma_k \rangle$ is a simplex of $\mathcal{N}_n$. To simplify the notation, set $W_i := W_{\Sigma_i}$. Since $\mathfrak{R}_{\Sigma_i} \subseteq K_{W_i}$, we have $\mathfrak{R}_\sigma \subseteq \bigcap K_{W_i}$. Hence, $\bigcap K_{W_i} \neq \varnothing$, and by Corollary 5.15, $\bigcap K_{W_i} = K_{W_\sigma}$. We will define a deformation retraction $K_{W_\sigma} \to \mathfrak{R}_\sigma$ by collapsing in each graph the minimal forest that takes that graph to a graph in $\mathfrak{R}_\sigma$.

Let $\Gamma \in K_{W_\sigma}$. Since $\bigcap K_{W_i} = K_{W_\sigma}$, Lemma 5.11 implies that there exists a minimal forest $\Phi_{\Sigma_i} \subseteq \Gamma$ collapsing Γ to a graph in $\mathfrak{R}_{\Sigma_i}$. Set

$$\Phi_\sigma(\Gamma) := \Phi_{\Sigma_0}(\Gamma) \cup \dots \cup \Phi_{\Sigma_k}(\Gamma).$$

Since $\Gamma \in K_{W_\sigma}$, there is a spanning tree T collapsing Γ to a rose: $\Gamma/T \in K_{W_\sigma}$. By Corollary 5.15, $\Gamma/T \in K_{W_i}$ for each i. So, by Lemma 5.10, $\Gamma/T \in \mathfrak{R}_{\Sigma_i}$ for each i and therefore $\Phi_{\Sigma_i}(\Gamma) \subseteq T$ for each i. Hence $\Phi_\sigma(\Gamma) \subseteq T$. Since T is a tree, $\Phi_\sigma(\Gamma)$ is a forest.

Now we define a map r from the vertex set of K_{W_σ} to the vertex set of $\mathfrak{R}_\sigma$ by $r(\Gamma) = \Gamma/\Phi_\sigma(\Gamma)$. We claim that r induces a simplicial map

$$r : K_{W_\sigma} \to \mathfrak{R}_\sigma.$$

It will suffice to show that r takes adjacent vertices to the same vertex or adjacent vertices because both K_{W_σ} and $\mathfrak{R}_\sigma$ are determined by their 1-skeletons. To do this, suppose that Γ_1 and Γ_2 represent adjacent vertices in K_{W_σ}. By possibly switching the names of the graphs, we can write $\Gamma_2 = \Gamma_1/\Phi$ for some forest Φ. If $\Phi \subseteq \Phi_\sigma(\Gamma_1)$, then $r(\Gamma_1) = r(\Gamma_2)$. If $\Phi \nsubseteq \Phi_\sigma(\Gamma)$ then $r(\Gamma_1) \neq r(\Gamma_2)$. The diagram below represents a small portion of K_n in this case, with edges represented by arrows.

$$\begin{tikzcd}
r(\Gamma_2) & & r(\Gamma_1) \arrow[ll, "\text{collapse } \Phi''"'] \\
\Gamma_2 \arrow[u, "\text{collapse } \Phi_\sigma(\Gamma_2)"] & & \Gamma_1 \arrow[ll, "\text{collapse } \Phi"] \arrow[ull, "\text{collapse } \Phi'"'] \arrow[u, "\text{collapse } \Phi_\sigma(\Gamma_1)"']
\end{tikzcd}$$

To show that r takes Γ_1 and Γ_2 to adjacent vertices, as the diagram suggests, we need justify that there is such a forest Φ'', as indicated in the diagram. The forest Φ'' is constructed as follows. Let $\Phi' = \Phi_\sigma(\Gamma_1) \cup \Phi$. Then Φ' is the subforest of Γ_1

such that $r(\Gamma_2) = \Gamma_1/\Phi'$. If Φ'' is the subgraph of $r(\Gamma_1)$ consisting of the images of the edges in $\Phi' - \Phi_\sigma(\Gamma_1)$ then Φ'' is a forest and $r(\Gamma_1)/\Phi'' = r(\Gamma_2)$. Therefore $r(\Gamma_1)$ and $r(\Gamma_2)$ are adjacent, proving that r induces a simplicial map.

That r is a retraction follows from the implication $\Gamma \in \mathfrak{R}_\sigma \Rightarrow \Phi_\sigma(\Gamma) = \varnothing$ and the fact that, by Lemma 5.9, the image of r is contained in $\mathfrak{R}_{\Sigma_i}$ for each i. Therefore $r(K_{W_\sigma}) = \mathfrak{R}_\sigma$. To see that r is a deformation retraction, we will use the poset lemma. Partially order the vertices of K_{W_σ} by setting $\Gamma_1 < \Gamma_2$ if Γ_1 can be collapsed to Γ_2. Then K_{W_σ} is the geometric realization of the poset of its vertices under this partial order. With respect to this partial order, r has the property that $\Gamma \leq r(\Gamma)$. The poset lemma implies that r is a deformation retraction. Since K_{W_σ} is contractible, this finishes the proof. □

By Lemma 5.1, Proposition 5.2 proves that $\mathcal{N}_n \simeq K_n$. Since K_n is contractible by [Culler and Vogtmann 1986], so is $\mathcal{N}_n$. We record this:

Theorem 5.16. *$\mathcal{N}_n$ is contractible for all n.*

6. Simplex stabilizers

As mentioned before, the action of $\mathrm{Out}(F_n)$ on K_n gives an action of $\mathrm{Out}(F_n)$ on $\mathcal{N}_n$. To describe stabilizers of this action, we fix some notation. Let Σ be a surface with boundary and/or punctures and with free fundamental group. As in the introduction, the pure mapping class group of Σ is the group of homeomorphisms of Σ that are the identity on the boundary and fix the punctures, up to isotopy relative to the boundary. The extended mapping class group of Σ is the group of isotopy classes of homeomorphisms of Σ. Thus the extended mapping class group contains orientation reversing homeomorphisms, while the pure mapping class group does not, provided that Σ has boundary. We will use $\mathrm{P}\Gamma(\Sigma)$ and $\Gamma(\Sigma)$ to represent the pure and extended mapping class groups of Σ.

If (Σ, s) is a homotopy marked surface then the identification of $\pi_1(\Sigma)$ with F_n given by the marking s induces a homomorphism from $\Gamma(\Sigma)$ to $\mathrm{Out}(F_n)$. This homomorphism is defined by sending a homeomorphism of Σ to the outer automorphism of $\pi_1(\Sigma)$ that it represents. By [Zieschang et al. 1980, Theorem 5.15.3], this homomorphism is injective, and its image is the subgroup of $\mathrm{Out}(F_n)$ consisting of outer automorphisms that take W_Σ to W_Σ or $(W_\Sigma)^{-1}$. Denote this subgroup by $\mathrm{MCG}^\pm(\Sigma, s)$. Denote the image of $\mathrm{P}\Gamma(\Sigma)$ by $\mathrm{PMCG}(\Sigma, s)$ and the image of the orientation-preserving subgroup of $\Gamma(\Sigma)$ by $\mathrm{MCG}(\Sigma, s)$. Note that these subgroups depend on the marking s. Thus the difference between $\Gamma(\Sigma)$ and $\mathrm{MCG}^\pm((\Sigma, s))$ is that $\mathrm{MCG}^\pm(\Sigma)$ is viewed as a subgroup of $\mathrm{Out}(F_n)$, and this subgroup depends on the marking s. The same is true for the pure mapping class groups.

Finally, let $\mathrm{Stab}(\mathfrak{R}_{[\Sigma,s]})$ be the subgroup of $\mathrm{Out}(F_n)$ stabilizing $\mathfrak{R}_{[\Sigma,s]}$ setwise, so $\mathrm{Stab}(\mathfrak{R}_{[\Sigma,s]})$ is the stabilizer of the vertex of $\mathcal{N}_n$ that corresponds to $[\Sigma, s]$.

Theorem 6.1. $\mathrm{Stab}(\mathfrak{R}_{[\Sigma,s]}) = \mathrm{MCG}^{\pm}(\Sigma, s)$.

Proof. Since $\psi \in \mathrm{MCG}^{\pm}(\Sigma, s)$ implies $(\Sigma, s) \cdot \psi = (\Sigma, s \circ |\psi|) \sim_{\pm} (\Sigma, s)$ (and hence $\mathfrak{R}_{[\Sigma,s]} \cdot \psi = \mathfrak{R}_{[\Sigma,s]}$), we have $\mathrm{Stab}(\mathfrak{R}_{[\Sigma,s]}) \supset \mathrm{MCG}^{\pm}(\Sigma, s)$. For the other inclusion, suppose that $\psi \in \mathrm{Stab}(\mathfrak{R}_{[\Sigma,s]})$. Then

$$\mathfrak{R}_{[\Sigma,s]} = (\mathfrak{R}_{[\Sigma,s]}) \cdot \psi = \mathfrak{R}_{[\Sigma,s\circ|\psi|]}.$$

By Proposition 4.6, $[\Sigma, s] = [\Sigma, s \circ |\psi|]$, so there is a homeomorphism $h : \Sigma \to \Sigma$ that makes the diagram

$$\begin{tikzcd} R_0 \arrow[r, "s"] \arrow[d, "|\psi|"'] & \Sigma \arrow[d, "h"] \\ R_0 \arrow[r, "s"'] & \Sigma \end{tikzcd}$$

commute up to homotopy. Now, h takes the boundary classes W_Σ to W_Σ or $(W_\Sigma)^{-1}$. Thus ψ does also. By Zieschang's theorem, $\psi \in \mathrm{MCG}^{\pm}(\Sigma)$. $\square$

To describe the stabilizer of a higher-dimensional simplex, we use a certain kind of stabilizer of a set of conjugacy classes of F_n, as studied in [McCool 1975]. Following the definitions there, we consider ordered m-tuples

$$(w_1, \ldots, w_m)$$

of conjugacy classes in F_n. The symmetric group S_m acts on the set of m-tuples by permuting the coordinates. The *inverting operations* $\tau_1, \ldots, \tau_m$ act on the set of m-tuples by

$$\tau_i(w_1, \ldots, w_i, \ldots, w_m) = (w_1, \ldots, w_i^{-1}, \ldots, w_m).$$

The group S_m, together with the τ_i's, generates the subgroup $\Omega_m \cong S_m \wr \mathbb{Z}_2$ of permutations of the set of m-tuples of conjugacy classes of F_n known as the extended symmetric group. The group $\mathrm{Out}(F_n)$ also acts on the set of m-tuples of conjugacy classes by acting individually on the coordinates. McCool [1975] makes the following definition in the setting of $\mathrm{Aut}(F_n)$, but we will use it also in the setting of $\mathrm{Out}(F_n)$.

Definition 6.2. For U an m-tuple of conjugacy classes and subgroup $G \leq \Omega_m$, define the subgroup $\mathcal{A}_{U,G}$ of $\mathrm{Out}(F_n)$ by

$$\mathcal{A}_{U,G} := \{\theta \in \mathrm{Out}(F_n) : \theta U \in GU\},$$

where $GU = \{gU : g \in G\}$.

For a simplex σ of $\mathcal{N}_n$, let G_σ denote the stabilizer of the simplex σ of $\mathcal{N}_n$. If $\sigma = v$ is the vertex corresponding to the marked surface (Σ, s), then $G_v = \mathrm{MCG}^{\pm}(\Sigma, s) = \mathcal{A}_{U,G}$ where U is the m-tuple of boundary classes (in any order) of a marked surface (Σ, s), and $G \le \Omega_m$ is the subgroup generated by S_m together with the extended permutation $\tau_1\tau_2\cdots\tau_m$. To describe G_σ for a higher-dimensional simplex, we introduce some terminology.

Definition 6.3. Let $\sigma = \langle \Sigma_0, \dots, \Sigma_k\rangle$ be a simplex of $\mathcal{N}_n$, and let U_i be the m_i-tuple of boundary classes of Σ_i (again in any order). Set $m = m_0 + \cdots + m_k$ and denote by $U_\sigma = (U_0, \dots, U_k)$ the m-tuple constructed by listing the conjugacy classes from the U_i one after another, starting with those of U_0. Define $H_\sigma \le \Omega_m$ as the subgroup generated by extended permutations of the following types:

(1) $\alpha \in S_m$ such that there is a permutation $\lambda \in S_k$ such that, for each i, α takes U_i to $U_{\lambda(i)}$, possibly with the entries of $U_{\lambda(i)}$ permuted;

(2) $\tau \in \Omega_m$ such that $\tau U_i = U_i$ or $\tau U_i = U_i^{-1}$ for each i.

Proposition 6.4. *For any simplex σ of $\mathcal{N}$, G_σ has the form $\mathcal{A}_{U_\sigma, H_\sigma}$.*

Proof. Formally, (1) can be written as $\alpha U_i \in S_{m_{\lambda(i)}} U_i$. An element $\theta \in G_\sigma$ permutes the equivalence classes of the marked surfaces Σ_i. This means that θ takes W_{Σ_i} to W_{Σ_j} or $W_{\Sigma_j}^{-1}$ for some j. Thus $\theta U_i \in S_{m_j} U_j$ or $S_{m_j} U_j^{-1}$ for some j. Since no two surfaces are taken to the same surface by θ, this means precisely that $\theta \in \mathcal{A}_{U_\sigma, G_\sigma}$ as defined above. $\square$

7. Equivariant homology of the action of Out(F_n) on $\mathcal{N}_n$

For a cellular action of a group G on a contractible cell complex X, the equivariant spectral sequence for the action of G on X is a well-known spectral sequence that converges to a grading of the homology of G; see [Brown 1982, Chapter VII.7]. To describe this spectral sequence, let M be any G-module. Consider, for each p-cell σ of X, the G_σ-module $\mathbb{Z}_\sigma$. As an additive group, $\mathbb{Z}_\sigma$ is isomorphic to $\mathbb{Z}$. The module structure of $\mathbb{Z}_\sigma$ is given by having $g \in G_\sigma$ act as multiplication by $+1$ or -1, depending on whether g preserves or reverses the orientation of σ. The module $\mathbb{Z}_\sigma$ is called the *orientation module* of σ. Let

$$M_\sigma := \mathbb{Z}_\sigma \otimes_{\mathbb{Z}} M.$$

Fix a set Δ_p of representatives for the orbits of the p-cells of X under the action of G. The equivariant spectral sequence for the action of G on X takes the form

$$E^1_{pq} = \bigoplus_{\sigma \in \Delta_p} H_q(G_\sigma; M_\sigma) \Rightarrow H_{p+q}(G; M). \tag{10}$$

Applying this to the action of $\mathrm{Out}(F_n)$ on $\mathcal{N}_n$ and any $\mathrm{Out}(F_n)$-module M, we get a spectral sequence converging to $H_*(\mathrm{Out}(F_n); M)$. Since vertex stabilizers are extended mapping class groups and there is one orbit of vertex for each homeomorphism type of surface, the $p = 0$ column of the spectral sequence consists of direct sums of the homology groups of the extended mapping class groups. For $p > 0$, the simplex stabilizers are given by Proposition 6.4 and we have:

Theorem 7.1. *For any* $\mathrm{Out}(F_n)$*-module* M*, there is a spectral sequence of the form*

$$E^1_{pq} = \bigoplus_{\sigma \in \Delta_p} H_q(G_\sigma; M_\sigma) \Rightarrow H_{p+q}(\mathrm{Out}(F_n); M), \tag{11}$$

where Δ_0 *is the set of homeomorphism classes of punctured orientable surfaces with fundamental group* F_n*, and for a vertex* $v \in \Delta_0$ *corresponding to surface* Σ *the stabilizer* G_v *is the extended mapping class group* $\mathrm{MCG}^\pm(\Sigma)$*. For* $p > 0$*, each* G_σ *is a generalized stabilizer of the form* $\mathcal{A}_{U_\sigma, H_\sigma}$*.*

The map induced on homology by the inclusion $\mathrm{MCG}^\pm(\Sigma) \hookrightarrow \mathrm{Out}(F_n)$ appears in the spectral sequence as the left-hand edge map, which is defined for the general spectral sequence (10) as follows. Since there is nothing but zeroes to the left of the $p = 0$ column in the spectral sequence (10), E^∞_{pq} is a quotient of E^1_{pq}. The spectral sequence converges to a grading of $H_*(G; M)$, and the composition

$$\bigoplus_{v \in \Delta_0} H_q(G_v; M) = E^1_{0q} \twoheadrightarrow E^\infty_{0q} = Gr_0 H_q(G; M) \hookrightarrow H_q(G; M) \tag{12}$$

is the left-hand edge map of this spectral sequence. The left hand edge map is equal to the map induced on homology by the inclusion of G_v into G.

For sequence (11), if the vertex v corresponds to marked surface Σ, we have $G_v = \mathrm{MCG}^\pm(\Sigma)$ and the restriction of the left-hand edge map to the subspace $H_q(G_v; M)$ is the map

$$H_q(\mathrm{MCG}^\pm(\Sigma); M) \to H_q(\mathrm{Out}(F_n); M)$$

induced by the inclusion $\mathrm{MCG}^\pm(\Sigma) \hookrightarrow \mathrm{Out}(F_n)$. Thus finding a bound on the rank of the left-hand edge map gives an upper bound on the contribution that the mapping class subgroups of $\mathrm{Out}(F_n)$ can make to the homology of $\mathrm{Out}(F_n)$. This will be the subject of the next section.

8. Analysis of E^∞

In this section, we specialize to rational coefficients and give a method for using Harer's homology stability theorems [1985] for mapping class groups to analyze the E^∞ page of spectral sequence (11). We continue to use $\mathrm{MCG}^\pm(\Sigma, s)$ for the image of $\Gamma(\Sigma, s)$ in $\mathrm{Out}(F_n)$, but we extend this notation to surfaces with boundary.

Hence, if Σ is a surface with boundary and s is a homotopy equivalence from the standard rose R_0 to Σ, we use $\mathrm{MCG}^{\pm}(\Sigma, s)$, $\mathrm{PMCG}(\Sigma, s)$, and $\mathrm{MCG}(\Sigma, s)$ to denote the images in $\mathrm{Out}(F_n)$ of the extended, pure and orientation-preserving mapping class groups of Σ in $\mathrm{Out}(F_n)$. As usual, these images depend on the marking s. We remark that there is a natural inclusion $P\Gamma(\Sigma) \hookrightarrow \Gamma(\Sigma)$, which agrees with the inclusion of $\mathrm{PMCG}(\Sigma, s)$ into $\mathrm{MCG}^{\pm}(\Sigma, s)$.

More generally, if Σ_0 is a subsurface with boundary of the surface Σ, the inclusion $\Sigma_0 \hookrightarrow \Sigma$ induces a map $\alpha : P\Gamma(\Sigma_0) \to P\Gamma(\Sigma)$ defined by extending a homeomorphism of Σ_0 to all of Σ by the identity. Harer's stability theorem, quoted below, implies that α induces an isomorphism on homology in sufficiently high dimensions.

Theorem 8.1 [Harer 1985, Theorem 0.1]. *Let Σ_0 be a subsurface of Σ such that $\Sigma - \Sigma_0$ is connected, contains no punctures but is not simply connected. If the genus of Σ_0 is at least $3k - 1$, then $\alpha_* : H_k(P\Gamma(\Sigma_0); \mathbb{Q}) \to H_k(P\Gamma(\Sigma); \mathbb{Q})$ is an isomorphism.*

We will also need to analyze the effect of plugging a boundary component of Σ with a punctured disk. To this end, consider the maps Θ, Θ', Υ and Φ (between the appropriate surfaces) defined, respectively, by plugging a boundary component with a disk, plugging a boundary component with a punctured disk, plugging a puncture, and sewing a pair of pants to a boundary component. In the stable range, Theorem 8.1 applies to Φ. By making the appropriate identifications, $\Theta \circ \Phi$ induces the identity on homology, so in the stable range for Φ, the map Θ must induce an isomorphism. Since $\Theta = \Upsilon \circ \Theta'$, we have:

Lemma 8.2. *For $g \geq 3k - 2$, $(\Theta')_*$ is injective and Υ_* is surjective on the k-th homology.*

In order to relate these stability maps to the d^1 terms in the spectral sequence of Theorem 7.1, consider two marked surfaces (Σ, s) and (Σ', s'). Let ρ and ρ' be the marked images in Σ and Σ' of the marking rose. Suppose that there are separating simple closed curves $\gamma \subset \Sigma$ and $\gamma' \subset \Sigma'$ cutting off subsurfaces $\widetilde{\Sigma} \subset \Sigma$ and $\widetilde{\Sigma}' \subset \Sigma'$ with the following properties, illustrated in Figure 4.

(1) The basepoints of the marking roses lie on γ and γ'.

(2) No edge of the marking roses meets γ or γ' anywhere but at the basepoints of the roses.

(3) $\rho \cap \widetilde{\Sigma} \simeq \widetilde{\Sigma}$ and $\rho' \cap \widetilde{\Sigma}' \simeq \widetilde{\Sigma}'$.

(4) $\widetilde{\Sigma}$ and $\widetilde{\Sigma}'$ are homeomorphic by a homeomorphism taking each directed edge of $\rho \cap \widetilde{\Sigma}$ to a directed edge of $\rho' \cap \widetilde{\Sigma}'$ with the same labeling.

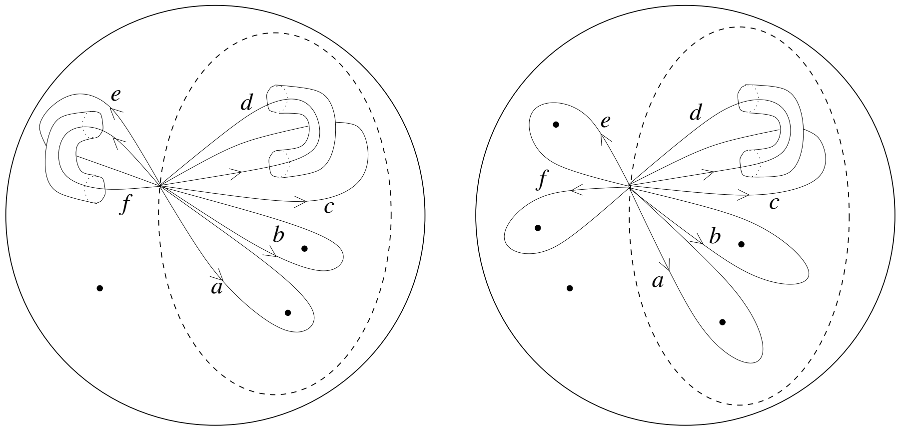

Figure 4. Markings that agree on a subsurface.

Definition 8.3. If the marked surfaces (Σ, s) and (Σ', s') satisfy conditions (1)–(4) above, the markings are said to *agree on the subsurfaces* $\widetilde{\Sigma}$ *and* $\widetilde{\Sigma}'$.

If the marked surfaces Σ and Σ' agree on the subsurface $\widetilde{\Sigma} = \widetilde{\Sigma}'$ then the corresponding vertices v_Σ and $v_{\Sigma'}$ in $\mathcal{N}$ span an edge e. The group $\mathrm{PMCG}(\widetilde{\Sigma})$ can be identified with a subgroup of the stabilizer of v_e. If Σ and Σ' are not homeomorphic surfaces, then the component of the d^1 map in spectral sequence (11) from $\mathrm{Stab}(e)$ to $\mathrm{Stab}(v_\Sigma)$ is the map induced by inclusion of $\mathrm{Stab}(e)$ into $\mathrm{Stab}(v_\Sigma)$. The same is true for Σ'. Therefore the following lemma will play the key role in determining bounds on how much homology in the E^1 page can survive until E^∞.

Lemma 8.4. *Let* $[\Sigma, s]$ *and* $[\Sigma', s']$ *be nonhomeomorphic marked surfaces such that*

(1) *the markings* s *and* s' *agree on subsurfaces* $\widetilde{\Sigma}$ *and* $\widetilde{\Sigma}'$ *of genus* $3k - 1$,

(2) $\widetilde{\Sigma}$ *contains all but one of the punctures of* Σ, *and*

(3) $\mathfrak{R}_{[\Sigma,s]} \cap \mathfrak{R}_{[\Sigma',s']} \neq \varnothing$.

Let v *and* v' *be the vertices of* $\mathcal{N}_n$ *corresponding to* Σ *and* Σ' *and let* e *be the edge of* $\mathcal{N}_n$ *between* v *and* v'. *Then* $G_e = G_v \cap G_{v'}$ *and* $i_* : H_k(G_e; \mathbb{Q}) \to H_k(G_v; \mathbb{Q})$ *has rank at least*

$$\dim H_k(G_v; \mathbb{Q}) - \big(\dim H_k(\mathrm{P}\Gamma(\Sigma); \mathbb{Q}) - \dim H_k(\mathrm{P}\Gamma(\widetilde{\Sigma}); \mathbb{Q})\big).$$

Proof. First, $G_e = G_v \cap G_{v'}$ because no outer automorphism of F_n can switch $[\Sigma, s]$ with $[\Sigma', s']$.

Since the Dehn twists generate $\mathrm{P}\Gamma(\widetilde{\Sigma})$, the map $\alpha : \mathrm{P}\Gamma(\widetilde{\Sigma}) \to \mathrm{P}\Gamma(\Sigma)$ induced by $\widetilde{\Sigma} \hookrightarrow \Sigma$ is determined by its effect Dehn twists. Now, suppose that c is a simple closed curve in $\widetilde{\Sigma}$ with corresponding simple closed curve c' in $\widetilde{\Sigma}'$. When we use the maps i_s and $i_{s'}$ of Zieschang's theorem to identify $\mathrm{P}\Gamma(\Sigma)$ and $\mathrm{P}\Gamma(\Sigma')$ with the subgroups $\mathrm{PMCG}(\Sigma)$ and $\mathrm{PMCG}(\Sigma') \subseteq \mathrm{Out}(F_n)$, the Dehn twists $\delta_c \in \mathrm{P}\Gamma(\widetilde{\Sigma})$ and $\delta_{c'} \in \mathrm{P}\Gamma(\widetilde{\Sigma}')$ about c and c' correspond to the same outer automorphism of F_n. Therefore the image of the map $\alpha_1 : \mathrm{P}\Gamma(\widetilde{\Sigma}) \to \mathrm{P}\Gamma(\Sigma) = \mathrm{PMCG}(\Sigma)$ lies in the intersection $\mathrm{PMCG}(\Sigma) \cap \mathrm{PMCG}(\Sigma') \subseteq G_v \cap G_{v'}$, and we have a commutative diagram

$$\begin{tikzcd} \mathrm{P}\Gamma(\widetilde{\Sigma}) \arrow[r, "\alpha_1"] \arrow[d, "\beta_1"'] & \mathrm{PMCG}(\Sigma) \arrow[d, hook, "\beta_2"] \\ G_e \arrow[r, hook, "i"'] & \mathrm{MCG}^{\pm}(\Sigma) = G_v \end{tikzcd} \tag{13}$$

where β_1 is the map α_1 viewed with a different range. We claim that the composition $\beta_2 \circ \alpha_1$ induces a map on homology that has rank at least

$$\dim H_k(G_v;\mathbb{Q}) - \big(\dim H_k(\mathrm{P}\Gamma(\Sigma);\mathbb{Q}) - \dim H_k(\mathrm{P}\Gamma(\widetilde{\Sigma});\mathbb{Q})\big).$$

To see this, note that β_2 is the inclusion of a finite index subgroup into a supergroup. Therefore it induces a surjection on homology with rational coefficients. By Theorem 8.1 and Lemma 8.2, the map α_1 is injective, and the claim follows by a dimension counting argument. Now, the rank of the map induced by i must be at least the rank of the map induced by $i \circ \beta_1 = \beta_2 \circ \alpha_1$, finishing the proof. □

Our last proposition gives a bound on the rank of the restriction of the left-hand edge map in spectral sequence (11) to surfaces of large rank. It bounds the contribution that the homology of these surfaces' mapping class groups can make to the homology of $\mathrm{Out}(F_n)$. To simplify notation, let $\Gamma^s_{g,0}$ denote the extended mapping class of the surface of genus g with s punctures and no boundary, and let $\mathrm{P}\Gamma^s_{g,0}$ be the pure, orientation preserving mapping class group of this surface.

Proposition 8.5. *Let $k \geq 0$ and $n \geq 6k-2$. For $g \geq 3k-1$, let Σ^s_g be the punctured surface of genus g with s punctures and with $2g+s-1 = n$ (so that $\pi_1(\Sigma^s_g) \cong F_n$). By choosing particular markings of the Σ^s_g, we may identify the vector space*

$$A := \bigoplus_{g \geq 3k-1} H_k(\mathrm{MCG}(\Sigma^s_g);\mathbb{Q})$$

with a subspace of the E^1_{0k} term of spectral sequence (11) *using trivial $\mathbb{Q}$ coefficients. The image of A in E^∞_{0k} has dimension no larger than*

$$\dim H_k(\Gamma^t_{3k-1,0};\mathbb{Q}) + \sum_{\substack{g \geq 3k \\ 2g+s-1=n}} \big(\dim H_k(\mathrm{P}\Gamma^s_{g,0};\mathbb{Q}) - \dim H_k(\mathrm{P}\Gamma^{s-1}_{g,0};\mathbb{Q})\big),$$

where $2(3k-1)+t-1 = n$.

Proof. Fix a marking $[\Sigma^t_{3k-1}, s_{3k-1}]$ of Σ^t_{3k-1}. For each $g > 3k-1$ we may choose a marking $[\Sigma^s_g, s_g]$ of Σ^s_g satisfying the conditions of Lemma 8.4 with $\Sigma = \Sigma^s_g$ and $\Sigma' = \Sigma^t_{3k-1}$. For $g \geq 3k-1$ let v_g denote the vertex of $\mathcal{N}_n$ corresponding to $[\Sigma^s_g, s_g]$ and for $g \geq 3k$, let e_g denote the edge between v_{3k-1} and v_g. By choosing the v_g and e_g as representatives for their $\mathrm{Out}(F_n)$ orbits, the vector spaces

$$A := \bigoplus_{g\geq 3k-1} H_k(G_{v_g}; \mathbb{Q}_{v_g}) \quad \text{and} \quad B := \bigoplus_{g\geq 3k} H_k(G_{e_g}; \mathbb{Q}_{e_g}) \tag{14}$$

can be identified with subspaces of E^1_{0k} and E^1_{1k} respectively. Note that $d^1(B) \subseteq A$. Now, G_{e_g} fixes e_g pointwise since no outer automorphism can switch v_g and v_{3k}. The same is true of G_{v_g}, so the modules $\mathbb{Q}_{v_g}$ and $\mathbb{Q}_{e_g}$ are actually trivial modules; $\mathbb{Q}_{v_g} = \mathbb{Q}$, $\mathbb{Q}_{e_g} = \mathbb{Q}$. Since $G_{v_g} = \mathrm{MCG}(\Sigma^s_g, s_g)$, the above definition of A agrees with the definition in the statement of the proposition, $A = \bigoplus_{g\geq 3k-1} H_k(\mathrm{MCG}(\Sigma^s_g); \mathbb{Q})$. Now, $G_{e_g} = G_{v_g} \cap G_{v_{3k-1}}$ as in Lemma 8.4. With these definitions, the (e_g, v_g)-component of d^1 is simply the map induced by the inclusion $G_{e_g} \hookrightarrow G_{v_g}$. By Lemma 8.4, this map has rank at least

$$R_g = \dim H_k(G_{v_g}; \mathbb{Q}) - \big(\dim H_k(\mathrm{P}\Gamma^s_{g,0}; \mathbb{Q}) - \dim H_k(\mathrm{P}\Gamma^{s-1}_{3k-1,1}; \mathbb{Q})\big), \tag{15}$$

where $2g+s-1=n$. Note that R_g depends on g because, even though $H_k(\mathrm{P}\Gamma^s_{g,0}; \mathbb{Q})$ is independent of g, it depends on s, and s depends on g.

By (15), for each $g \geq 3k$ we may choose R_g vectors $\{w^g_1, \ldots, w^g_{R_g}\}$ in $(\operatorname{im} d^1 \cap A)$ such that their projections onto $H_k(G_{v_g}; \mathbb{Q})$ are linearly independent. Let A_g be the subspace of A spanned by $\{w^g_1, \ldots, w^g_{R_g}\}$. By the first direct sum decomposition in (14) and the choice of the vectors w^g_i, the subspaces $A_{3k}, \ldots, A_{\lfloor \frac{n}{2} \rfloor}$ are linearly independent. Let A^∞ denote the image of A in E^∞. Since A^∞ is a quotient of $A/d^1(B)$, this means that

$$\dim(A^\infty) \leq \dim(A) - \sum_{g\geq 3k} R_g. \tag{16}$$

By Harer's stability theorems, $H_k(\mathrm{P}\Gamma^{s-1}_{3k-1,1}) \cong H_k(\mathrm{P}\Gamma^{s-1}_{g,0})$. Now, substituting

$$\dim A = \sum_{g\geq 3k-1} \dim H_k(G_{v_g}; \mathbb{Q})$$

and (15) with $\dim H_k(\mathrm{P}\Gamma^{s-1}_{3k-1,1}) = \dim H_k(\mathrm{P}\Gamma^{s-1}_{g,0})$ into (16) gives

$$\dim(A^\infty) \leq$$

$$\dim H_k(\mathrm{MCG}(\Sigma_{3k-1}); \mathbb{Q}) + \sum_{\substack{g\geq 3k \\ 2g+s-1=n}} \big(\dim H_k(\mathrm{P}\Gamma^s_{g,0}; \mathbb{Q}) - \dim H_k(\mathrm{P}\Gamma^{s-1}_{g,0}; \mathbb{Q})\big),$$

as required. □

Acknowledgments

This paper presents a portion of the results of my thesis, completed at Cornell University in 2003. I thank my advisor Karen Vogtmann for all of the time and support she gave me during my time at Cornell. These ideas would not have been possible without the many useful conversations we had about this material. I also thank the referee for suggestions on filling gaps in several proofs and for advice in improving the exposition of the paper.

References

[Bowditch and Epstein 1988] B. H. Bowditch and D. B. A. Epstein, "Natural triangulations associated to a surface", *Topology* **27**:1 (1988), 91–117. MR 89e:57004 Zbl 0649.32017

[Brown 1982] K. S. Brown, *Cohomology of groups*, Graduate Texts in Mathematics **87**, Springer, New York, 1982. MR 83k:20002 Zbl 0584.20036

[Culler and Vogtmann 1986] M. Culler and K. Vogtmann, "Moduli of graphs and automorphisms of free groups", *Invent. Math.* **84**:1 (1986), 91–119. MR 87f:20048 Zbl 0589.20022

[Harer 1985] J. L. Harer, "Stability of the homology of the mapping class groups of orientable surfaces", *Ann. of Math.* (2) **121**:2 (1985), 215–249. MR 87f:57009 Zbl 0579.57005

[Harer 1986] J. L. Harer, "The virtual cohomological dimension of the mapping class group of an orientable surface", *Invent. Math.* **84**:1 (1986), 157–176. MR 87c:32030 Zbl 0592.57009

[Harer 1993] J. L. Harer, "Improved homology stability for the homology of the mapping class groups of surfaces", preprint M-93-09, Duke University, 1993, online at http://www.math.duke.edu/preprints/93-09.dvi.

[Hatcher 2002] A. Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR 2002k:55001 Zbl 1033.55001

[Hatcher and Vogtmann 2004] A. Hatcher and K. Vogtmann, "Homology stability for outer automorphism groups of free groups", preprint, 2004. math.GT/0406377

[Hoare 1979] A. H. M. Hoare, "Coinitial graphs and Whitehead automorphisms", *Canad. J. Math.* **31**:1 (1979), 112–123. MR 80e:20050 Zbl 0412.20030

[Ivanov 1989] N. V. Ivanov, "Stabilization of the homology of Teichmüller modular groups", *Algebra i Analiz* **1**:3 (1989), 110–126. Translated in *Leningrad Mathematical J.* **1**:3 (1990), 675–691. MR 91g:57010 Zbl 0727.30036

[Madsen and Weiss 2002] I. Madsen and M. S. Weiss, "The stable moduli space of Riemann surfaces: Mumford's conjecture", preprint, 2002. math.AT/0212321

[Magnus et al. 1966] W. Magnus, A. Karrass, and D. Solitar, *Combinatorial group theory*, Pure and applied mathematics **13**, Interscience Publishers, New York, 1966. Reprinted by Dover, New York, 1976. MR 54 #10423 Zbl 0138.25604

[McCool 1975] J. McCool, "Some finitely presented subgroups of the automorphism group of a free group", *J. Algebra* **35** (1975), 205–213. MR 53 #624 Zbl 0325.20025

[Mulase and Penkava 1998] M. Mulase and M. Penkava, "Ribbon graphs, quadratic differentials on Riemann surfaces, and algebraic curves defined over $\overline{\mathbf{Q}}$", *Asian J. Math.* **2**:4 (1998), 875–919. MR 2001g:30028 Zbl 0964.30023

[Penner 1987] R. C. Penner, "The decorated Teichmüller space of punctured surfaces", *Comm. Math. Phys.* **113**:2 (1987), 299–339. MR 89h:32044 Zbl 0642.32012

[Quillen 1973] D. Quillen, "Higher algebraic K-theory. I", pp. 85–147 in *Algebraic K-theory, vol. I, Higher K-theories* (Seattle, WA, 1972), edited by H. Bass, Lecture Notes in Math. **341**, Springer, Berlin, 1973. MR 49 #2895 Zbl 0292.18004

[Vogtmann 2002] K. Vogtmann, "Automorphisms of free groups and outer space", *Geom. Dedicata* **94** (2002), 1–31. MR 2004b:20060 Zbl 1017.20035

[Zieschang et al. 1980] H. Zieschang, E. Vogt, and H.-D. Coldewey, *Surfaces and planar discontinuous groups*, Lecture Notes in Mathematics **835**, Springer, Berlin, 1980. MR 82h:57002 Zbl 0438.57001

Received July 2, 2004.

MATTHEW HORAK
DEPARTMENT OF MATHEMATICS, STATISTICS AND COMPUTER SCIENCE
HARVEY HALL 237E
UNIVERSITY OF WISCONSIN STOUT
MENOMONIE, WI 54751
UNITED STATES

horakm@uwstout.edu
http://faculty.uwstout.edu/horakm