ISSN 1364-0380 (on line) 1465-3060 (printed) **1471**

Geometry & Topology
Volume 8 (2004) 1471–1499
Published: 5 December 2004

Morita classes in the homology of automorphism groups of free groups

James Conant
Karen Vogtmann

Department of Mathematics, University of Tennessee
Knoxville, TN, 37996, USA
and
Department of Mathematics, Cornell Univeristy
Ithaca, NY 14853-4201, USA

Email: `jconant@math.utk.edu` and `vogtmann@math.cornell.edu`

Abstract

Using Kontsevich's identification of the homology of the Lie algebra ℓ_∞ with the cohomology of $\mathrm{Out}(F_r)$, Morita defined a sequence of $4k$–dimensional classes μ_k in the unstable rational homology of $\mathrm{Out}(F_{2k+2})$. He showed by a computer calculation that the first of these is non-trivial, so coincides with the unique non-trivial rational homology class for $\mathrm{Out}(F_4)$. Using the "forested graph complex" introduced in [2], we reinterpret and generalize Morita's cycles, obtaining an unstable cycle for every connected odd-valent graph. (Morita has independently found similar generalizations of these cycles.) The description of Morita's original cycles becomes quite simple in this interpretation, and we are able to show that the second Morita cycle also gives a non-trivial homology class. Finally, we view things from the point of view of a different chain complex, one which is associated to Bestvina and Feighn's bordification of outer space. We construct cycles which appear to be the same as the Morita cycles constructed in the first part of the paper. In this setting, a further generalization becomes apparent, giving cycles for objects more general than odd-valent graphs. Some of these cycles lie in the stable range. We also observe that these cycles lift to cycles for $\mathrm{Aut}(F_r)$.

AMS Classification numbers Primary: 20J06

Secondary: 20F65, 20F28

Keywords: Automorphism groups of free groups, graph homology

Proposed: Joan Birman Received: 21 June 2004
Seconded: Thomas Goodwillie, Ralph Cohen Revised: 1 December 2004

1 Introduction

The group $\mathrm{Out}(F_r)$ of outer automorphisms of a free group is closely related to various classical families of groups, including surface mapping class groups, lattices in semi-simple Lie groups and non-positively curved groups. However, it does not actually belong to any of these families, so theorems developed for these families do not apply directly to $\mathrm{Out}(F_r)$. Although a great deal of progress has been made in recent years by adapting classical techniques to suit $\mathrm{Out}(F_r)$, still comparatively little is known about this group. In particular, invariants such as homology and cohomology are not well understood. The rational homology has been completely computed only up to $n = 5$ [7, 5], and the only non-trivial rational homology group known prior to this paper is $H_4(\mathrm{Out}(F_4);\mathbb{Q}) \cong \mathbb{Q}$.

A novel approach to studying these invariants was discovered by Kontsevich. He proved a remarkable theorem stating that $\oplus_r H^*(\mathrm{Out}(F_r);\mathbb{Q})$ is basically the same as the homology of a certain infinite dimensional Lie algebra, ℓ_∞ [9, 10]. S. Morita recognized this Lie algebra as the kernel of the bracketing map on the free Lie algebra, and used the theory of Fox derivatives to define a family of cocycles on $\wedge\ell_\infty$ [13]. He then applied Kontsevich's theorem to obtain a family of homology classes in $\{H_*(\mathrm{Out}(F_r);\mathbb{Q})\}$. Morita was able to show that the first of these classes, which resides in $H_4(\mathrm{Out}(F_4);\mathbb{Q})$, is non-trivial, so in fact generates this homology group. He conjectured that his classes are all nontrivial, and are basic building blocks for the rational homology of $\mathrm{Out}(F_r)$. In the paper [13], Morita was mainly interested in studying the mapping class group and he did not go into detail about generalizing his series, but he was aware that it did generalize. We will explore such a generalization in this paper.

We use the interpretation of Kontsevich's theorem given in [2] to translate Morita's cycles into terms of *forested graphs.* This description is quite simple, and we can easily verify that the first Morita cycle is non-trivial. With the aid of a computer mathematics package, we are also able to show that the second Morita cycle, in $H_8(\mathrm{Out}(F_6);\mathbb{Q})$ is non-trivial, thus supporting Morita's conjecture and producing the second known non-trivial rational homology class for $\mathrm{Out}(F_r)$. The graphical description of Morita's cycles also leads to a natural generalization, giving a new cycle for every odd-valent graph.

In the last section of the paper, we diversify our portfolio of chain complexes. We introduce a new chain complex to compute $H_*(\mathrm{Out}(F_r);\mathbb{Q})$, which we construct using the bordification of outer space introduced by Bestvina and Feighn [1]. In this new complex, we construct cycles for every odd-valent graph and conjecture that they coincide with Morita's cycles. An advantage of the new complex is that a further generalization becomes apparent, and we obtain a cycle associated to a more general type of graph we call an AB–graph. Although the most ambitious conjecture, that all of these cycles correspond to independent homology classes, cannot be true, it may be true for the classes coming from odd-valent graphs. Moreover, it is possible that the classes coming from AB–graphs generate the entire homology.

Finally, we observe that all of the cycles we constructed (in the bordification context) lift to cycles in $H_*(\mathrm{Aut}(F_r);\mathbb{Q})$, but, with the exception of a class in $H_4(\mathrm{Aut}(F_4);\mathbb{Q})$, we do not know if they are homologically essential. A recent computer calculation of F. Gerlits [5] shows that the most ambitious conjectures here are also not true. Namely, he found a class in $H_7(\mathrm{Aut}(F_5);\mathbb{Q})$ which, for degree reasons, cannot arise by our construction.

Acknowledgment Sections 2 and 3 of this paper are partly based on a set of notes by Swapneel Mahajan. We also thank Shigeyuki Morita for an informative email in which he mentioned that he had independently found several generalizations of his original family of cycles. The first author was partially supported by NSF grant DMS 0305012. The second author was partially supported by NSF grant DMS 0204185.

2 Background and definitions

2.1 Kontsevich's theorem

In this section we briefly describe Kontsevich's Lie algebra ℓ_∞ and his theorem identifying the Lie algebra homology of ℓ_∞ with the cohomology of the groups $\mathrm{Out}(F_r)$. For a detailed exposition and proofs, we refer to [2].

Let V_n be the $2n$–dimensional symplectic vector space with standard symplectic basis $p_1,\dots,p_n,q_1,\dots,q_n$, let $\mathcal{L}_{2n}$ be the free Lie algebra generated by V_n, and let $\omega = \sum_{i=1}^n [p_i,q_i] \in \mathcal{L}_{2n}$. Kontsevich defined ℓ_n to be the Lie algebra consisting of those derivations of $\mathcal{L}_{2n}$ which kill ω. (Recall that $f\colon \mathcal{L}_{2n} \to \mathcal{L}_{2n}$ is a *derivation* if $f([x,y]) = [f(x),y] + [x,f(y)]$). The bracket $[f,f']$ of two derivations is defined to be the difference $f\circ f' - f'\circ f$. Notice that $\ell_n \subset \ell_{n+1}$ in a natural way, and define $\ell_\infty = \lim_{n\to\infty} \ell_n$.

The Lie algebra $\mathfrak{sp}(2n)$ sits naturally as a subalgebra of ℓ_n, and one case of Kontsevich's theorem relates the Lie algebra homology of ℓ_∞ to the homology of the limit $\mathfrak{sp}(2\infty)$ and to the cohomologies of the groups $\mathrm{Out}(F_r)$. Specifically, we have

Theorem 2.1 (Kontsevich [9, 10])

$$PH_k(\ell_\infty) \cong H_k(\mathfrak{sp}(2\infty)) \oplus \bigoplus_{r\geq 2} H^{2r-2-k}(\mathrm{Out}(F_r))$$

Here the prefix P means to take primitive elements in the Hopf algebra $H_*(\ell_\infty)$.

The isomorphism in Kontsevich's theorem is defined by relating both sides to an intermediate object, called *Lie graph homology.* Briefly, the Lie graph chain complex, $\ell\mathcal{G}$, is spanned by finite oriented graphs, whose vertices are colored by elements of the Lie operad. The boundary operator contracts edges one at a time, while composing the operad elements that color the endpoints of the contracting edge. As was shown in [2], this chain complex is equivalent to the *forested graph complex*, which we describe in section 4. The homology of ℓ_∞ is related to Lie graph homology via the invariant theory of $\mathfrak{sp}(2n)$, and the cohomology of $\mathrm{Out}(F_r)$ is related via the action of $\mathrm{Out}(F_r)$ on Outer space.

2.2 Fox derivatives

Let F_{2n} be the free group on $PQ_n = \{p_1, q_1, \ldots, p_n, q_n\}$. For each $x \in PQ_n$, the classical Fox derivative is a map

$$d_x\colon\ \mathbb{Z}[F_{2n}] \to \mathbb{Z}[F_{2n}],$$

where $\mathbb{Z}[F_{2n}]$ is the group ring. It is defined on $y \in PQ_n$ by the rule $d_x(y) = \delta_{xy}$, and extended to all of $\mathbb{Z}[F_{2n}]$ as a derivation, where the left action is multiplication and the right action is via the augmentation $\epsilon\colon\ \mathbb{Z}[F_{2n}] \to \mathbb{Z}$, ie,

$$d_x(ab) = (d_x a)\,\epsilon(b) + a\,(d_x b)\,.$$

One can verify that this determines a well defined map.

With this starting point, we wish now to define a "Fox derivative," on the free associative algebra,

$$d_x^{\mathbf{a}}\colon\ \mathcal{A}_{2n} \to \mathcal{A}_{2n},$$

where $\mathcal{A}_{2n}$ is the free associative algebra on the generating set PQ_n. Let x be an element of PQ_n. If w is a word in the generators that does not end with an x, then we define $d_x^{\mathbf{a}}(w) = 0$; if $w = vx$, then $d_x^{\mathbf{a}}(vx) = v$. Let $M\colon\ \mathbb{Z}[F_{2n}] \to \widehat{\mathcal{A}}_{2n}$ be the Magnus embedding, where $\widehat{\mathcal{A}}_{2n}$ is the completion to formal noncommuting power series. Then one can verify that the following diagram commutes:

$$\begin{array}{ccc} \mathbb{Z}[F_{2n}] & \xrightarrow{d_x} & \mathbb{Z}[F_{2n}] \\ \downarrow M & & \downarrow M \\ \widehat{\mathcal{A}}_{2n} & \xrightarrow{d_x^{\mathbf{a}}} & \widehat{\mathcal{A}}_{2n} \end{array}$$

It is therefore reasonable to call $d_x^{\mathbf{a}}$ an "associative Fox derivative."

The free Lie algebra embeds in the free associative algebra by sending $[a,b]$ to $ab-ba$, and the free associative algebra maps to the free commutative algebra $\mathcal{C}_{2n}$ by letting the variables commute. In this way, we get a "Fox derivative" $\mathfrak{D}_x$ from the free Lie algebra to the free commutative algebra:

$$\mathfrak{D}_x\colon\ \mathcal{L}_{2n} \to \mathcal{A}_{2n} \overset{d_x^{\mathbf{a}}}{\to} \mathcal{A}_{2n} \to \mathcal{C}_{2n}.$$

2.3 Morita's trace map

Morita recognized Kontsevich's ℓ_n to be the same as a Lie algebra $\mathfrak{h}_n$ which he had been studying in his work on mapping class groups. Let $\mathcal{L}_{2n}(k)$ be the degree k part of $\mathcal{L}_{2n}$, ie, the subspace spanned by iterated brackets of k vectors. Now $\mathfrak{h}_n(k)$ is defined to be the kernel of the bracketing map $V_n \otimes \mathcal{L}_{2n}(k) \to \mathcal{L}_{2n}(k+1)$ so that we have an exact sequence:

$$0 \to \ \mathfrak{h}_n(k) \to V_n \otimes \mathcal{L}_{2n}(k) \to \mathcal{L}_{2n}(k+1) \to 0$$

We define Lie algebras $\mathfrak{h}_n = \oplus_{k\geq 1}\mathfrak{h}_n(k)$, and $\mathfrak{h}_n^+ = \oplus_{k\geq 2}\mathfrak{h}_n(k)$. For more information about $\mathfrak{h}_n^+$, see [13]. (In that paper, the notation $\mathfrak{h}_{n,1}^{\mathbb{Q}}$ is the same as our $\mathfrak{h}_n^+$.)

The isomorphism $\ell_n \to \mathfrak{h}_n$ is given in the following way. Given a derivation f which kills ω, consider the restriction f_1 of f to V_n. Then $f_1 \in Hom(V, \mathcal{L}_{2n}) \cong V^* \otimes \mathcal{L}_{2n}$. We use the symplectic form to identify $V^* \otimes \mathcal{L}_{2n}$ with $V \otimes \mathcal{L}_{2n}$. One easily checks that the condition that the image of f_1 is in the kernel of the bracket map corresponds exactly to the condition that $f(\omega) = 0$.

Morita defined his trace as a map from $\mathfrak{h}_n$ to the free polynomial algebra $\mathcal{C}_{2n}$ on $2n$ variables. Using the isomorphism above, the trace becomes the following map from ℓ_n to $\mathcal{C}_{2n}$. Given a derivation $f\colon \mathcal{L}_{2n} \to \mathcal{L}_{2n}$, we form the Fox Jacobian

$$\left(\mathfrak{D}_{x_j} f(x_i)\right).$$

where $x_i, x_j \in PQ_n$ The trace of this matrix is a polynomial in $2n$ commuting variables, giving a map

$$\tau\colon\ \ell_n \to \mathcal{C}_{2n},$$

which we will also call the *trace map.*

It will turn out that τ is a morphism of Lie algebras, provided $\mathcal{C}_{2n}$ is considered as an abelian Lie algebra. Indeed Morita conjectures that the image of τ, together with the low degree correction, $\Lambda^3 V_n$, deriving from the first Johnson homomorphism, is exactly the abelianization of $\mathfrak{h}_n^+$ [13, Conjecture 6.10].

3 Graphical translation of Morita's trace

In this section we identify ℓ_n with a Lie algebra of trees, reinterpret the trace τ as a map from this Lie algebra, and prove some of the basic properties of the trace.

We first recall how to think of generators of the free commutative, associative and Lie algebras ($\mathcal{C}_{2n}, \mathcal{A}_{2n}$ and $\mathcal{L}_{2n}$) in terms of finite labeled trees. We interpret a monomial in $\mathcal{C}_{2n}$ geometrically as a rooted tree with one interior vertex, whose leaves are labeled by elements of V_n. The free associative algebra $\mathcal{A}_{2n}$ is spanned by rooted, labeled *planar* trees with one interior vertex. The free Lie algebra $\mathcal{L}_{2n}$ is spanned by rooted, labeled, planar *trivalent* trees. In the Lie case (only) the trees are not linearly independent; a rooted, planar trivalent tree with k labeled leaves corresponds in a natural way to a bracket expression of k letters, ie, to a generator of the free Lie algebra. The Jacobi relation in the free Lie algebra translates to the IHX relation among planar trees, and the anti-symmetry of the bracket gives the AS relation.

From each of the algebras $\mathcal{C}_{2n}, \mathcal{A}_{2n}$ and $\mathcal{L}_{2n}$, or more precisely from their underlying cyclic operads, we can form a Lie algebra, as in [2]. A generator of this Lie algebra is again a tree of the specified type, but with no root, so that all leaves are labeled by elements of V_n. We assume that each tree has at least two leaves. We can *graft* a tree X_1 to a tree X_2 by identifying a leaf of X_1 with a leaf of X_2, erasing the associated labels v_1 and v_2, and multiplying the resulting tree by a coefficient given by the symplectic product $\langle v_1, v_2\rangle$. In the commutative and associative cases, we also contract the interior edge that we just created. The Lie bracket $[X_1, X_2]$ is defined by grafting X_1 to X_2 at all possible leaves, then adding up the results. The Lie algebras obtained in this way from $\mathcal{C}_{2n}, \mathcal{A}_{2n}$ and $\mathcal{L}_{2n}$ are denoted $\mathcal{LC}_n$ $\mathcal{LA}_n$ and $\mathcal{LL}_n$. The Lie algebra $\mathcal{LC}_n$ is isomorphic as a vector space to the degree ≥ 2 part of $\mathcal{C}_{2n}$.

We now wish to identify $\mathcal{LL}_n$ with Kontsevich's Lie algebra ℓ_n. Given a tree X representing a generator of $\mathcal{LL}_n$, we define a derivation $f_X\colon \mathcal{L}_{2n} \to \mathcal{L}_{2n}$ as follows. For each rooted tree T representing a generator of $\mathcal{L}_{2n}$, we graft X onto T at all possible leaves, with coefficient determined by the symplectic form, and add up the results to obtain $f_X(T)$. The derivation f_X is easily seen to kill ω, so lies in ℓ_n.

Proposition 3.1 *Over a field of characteristic zero, the map $\mathcal{LL}_n \to \ell_n$ sending X to f_X is an isomorphism of Lie algebras.*

Proof This follows from the fact that $\mathcal{LL}_n \cong \mathfrak{h}_n$, [11, Corollary 3.2] although the main point of that paper is that the result is not true for the corresponding objects defined over $\mathbb{Z}$. The map $\mathcal{LL}_n \to \mathfrak{h}_n$ is defined on a labeled tree X as a sum $\sum v \otimes X_v$, where the sum is over all univalent vertices of X, v is the vector in V_n labeling that vertex, and X_v is the rooted tree formed by thinking of the chosen vertex as the root. Composing this map with the inverse of our map $\ell_n \to \mathfrak{h}_n$, we get the indicated result. □

In order to understand Morita's trace map as a map from trees, we first need to define another type of "partial derivative,"

$$D_{p_i}\colon \mathcal{LL}_n \to \mathcal{L}_{2n}.$$

This is given by in turn replacing each occurrence of p_i by a root, and adding the results. If X is a single (unrooted) labeled trivalent tree, then $D_{p_i}(X)$ will be a sum of rooted labeled trivalent trees, one for each occurrence of p_i at a leaf of X. The partial D_{q_i} is defined analogously.

Proposition 3.2 *The trace $\tau\colon \mathcal{LL}_n \to \mathcal{C}_{2n}$ is given by*

$$\tau(X) = \sum_{i=1}^{n} \left(\mathfrak{D}_{p_i} D_{q_i}(X) - \mathfrak{D}_{q_i} D_{p_i}(X)\right).$$

Proof Let X be a generator of $\mathcal{LL}_n$, and let $f_X\colon \mathcal{L}_{2n} \to \mathcal{L}_{2n}$ be the corresponding derivation, in ℓ_n. Then $\tau(f_X)$ is the trace of the Fox Jacobian:

$$\sum_i \left(\mathfrak{D}_{p_i} f_X(p_i) + \mathfrak{D}_{q_i} f_X(q_i)\right)$$

The vector p_i is represented graphically as a rooted tree with a single leaf, labeled p_i. Thus $f_X(p_i)$ is a sum of rooted trees, obtained by grafting this tree to each occurrence of q_i in X. The result of grafting is that the leaf formerly labeled q_i becomes the root, and we have a rooted tree representing an element of $\mathcal{L}_{2n}$. Thus $f_X(p_i)$ coincides with the q_i-th partial derivative of X. Similarly, $f_X(q_i)$ coincides with the p_i-th partial derivative, but the grafting operation incurs a minus sign. $\square$

We next show that the trace vanishes on all but certain very special types of trees in $\mathcal{LL}_n$. For any tree T, we define the *interior* of T to be the tree obtained by removing all leaves and the edges incident to those leaves. If the interior of T has no trivalent vertices (ie is a line), we say T is *linear*. If the interior of T has an odd number of edges, we say T is *odd*; otherwise T is *even*.

Lemma 3.3 *Let T be a generator of $\mathcal{L}_{2n}$, represented by a rooted, labeled, planar trivalent tree. Then $\mathfrak{D}_{p_i}(T) = 0$ unless T is a linear tree with exactly one leaf labeled p_i at maximal distance from the root. If the other leaves are labeled $a_1, \ldots, a_k$, then $\mathfrak{D}_{p_i}(T) = \pm a_1 \ldots a_k$. The sign is determined by the fact that $\mathfrak{D}_{p_i}[a_1, [a_2, \ldots, [a_k, p_i] \ldots]] = a_1 \ldots a_k$.*

Proof If T is of the specified form, then consider the unique geodesic path from the root to the leaf labeled p_i. We can use AS relations, if necessary, to put all of the leaves a_j onto one side of this geodesic, so that we may assume T is the tree corresponding to

$$[a_1, [a_2, \ldots [a_k, p_i] \ldots]].$$

Consider the image P of $[a_1, [a_2, \ldots [a_k, p_i] \ldots]]$ in the free associative algebra $\mathcal{A}_{2n}$. If no a_j is equal to p_i, then the only term of P which ends in p_i is $a_1 a_2 \ldots a_k p_i$, so that $\mathfrak{D}_{p_i}(W) = a_1 \ldots a_k$. On the other hand, suppose some a_j is equal to p_i, with $j < k$ minimal. Let $\beta = [a_{j+1}, \ldots, [a_k, p_i] \cdots]$. Then the only terms that could possibly end in p_i come from

$$a_1 \cdots a_{j-1}[p_i, \beta] = a_1 \cdots a_{j-1} p_i \beta - a_1 \cdots a_{j-1} \beta p_i.$$

By induction on the number of occurrences of p_i, we know that $\mathfrak{D}_{p_i}\beta = a_{j+1} \cdots a_k$, so that $\mathfrak{D}(a_1 \cdots a_{j-1} p_i \beta) = a_1 \cdots a_{j-1} p_i \mathfrak{D}(\beta) = a_1 \cdots a_k$. We also have

$$\mathfrak{D}_{p_i}(a_1 \cdots a_{j-1} \beta p_i) = a_1 \cdots a_{j-1} \beta,$$

but this second term vanishes when we pass to $\mathcal{C}_{2n}$ since β is a commutator.

If T has any form other than the one specified, then its image P in $\mathcal{A}_{2n}$ can be grouped into terms so that $\mathfrak{D}_{p_i}(T)$ contains a commutator as a factor in each term, which becomes zero when we pass to $\mathcal{C}_{2n}$. $\square$

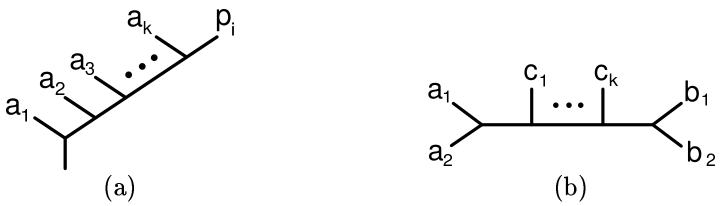

Figure 1: (a) The unique tree on which $\mathfrak{D}_{p_i}$ does not vanish, and (b) an element of $\mathcal{LL}_n(k+2)$.

The tree T corresponding to $[a_1,[a_2,\ldots,[a_k,p_i]\ldots]]$ is shown in Figure 1a. Note that this is $(-1)^k[[\ldots[p_i,a_k],\ldots,a_2],a_1]$, which corresponds to the tree obtained by flipping all of the leaves labeled a_i across the axis from the root to p_i. This observation allows us to easily compute the trace of any generator of $\mathcal{LL}_n$.

Corollary 3.4 *Let X be a nonzero generator of $\mathcal{LL}_n$, represented by a labeled planar trivalent tree. Then $\tau(X)=0$ unless X is of the form shown in Figure 1b, with k odd. In this case,*

$$\tau(X)=2c_1\ldots c_k\big(-\langle a_1,b_1\rangle a_2b_2+\langle a_1,b_2\rangle a_2b_1+\langle a_2,b_1\rangle a_1b_2-\langle a_2,b_2\rangle a_1b_1\big),$$

where $\langle\cdot,\cdot\rangle$ is the symplectic form.

Note that $a_1\neq a_2$ and $b_1\neq b_2$, since X is non-zero. If $a_1=p_i$, $b_1=q_i$ and a_2 and b_2 are not paired, then $\tau(X)=-2a_2c_1\ldots c_kb_2$.

Lemma 3.5 *For any $X,Y\in\mathcal{LL}_n$, $\tau([X,Y])=0$.*

Proof Each term in $[X,Y]$ is obtained by grafting a leaf of X to a leaf of Y to obtain a new tree Z. If either X or Y is non-linear, then so is Z. If X and Y are both linear, and both even or both odd, then the interior of Z is either non-linear, or is odd linear, so has trace zero. The only possible way to obtain a term with a non-zero trace is if X is odd linear and Y is even linear, or vice-versa, and we are grafting extremal leaves of X and Y, say labeled p_i and q_i.

In order for the trace of Z to be non-zero, it must have an extremal pair of leaves labeled q_j and p_j for some j. In this case, there is another term Z' in the bracket, obtained by grafting these two leaves. We claim that the trace of Z cancels with the trace of Z'.

In the following figure, X is on the left, Y is on the right.

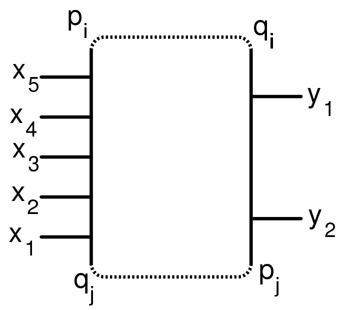

Z is obtained by joining the lower two leaves, whereas Z' is obtained by joining the upper two leaves and multiplying by $-1 = \langle q_j, p_j \rangle$. Let $\pi = 2x_1 \cdots x_s y_1 \cdots y_t$, where the x_k's and the y_k's are the labels of the leaves of X and Y which are present in both Z and Z'. There are an odd number of such leaves, and π is the trace, up to sign. We can draw Z as a horizontal line from q_i to p_i with all of the other edges pointing upwards from this line. By Corollary 3.4 the trace of this is π. We can draw $-Z'$ as a horizontal edge from q_j to p_j with the other edges pointing upwards. The trace of $-Z'$ is therefore also π which implies that $\tau(Z' + Z) = -\pi + \pi = 0$ as desired. □

Lemma 3.5 shows that τ is a Lie algebra homomorphism, if we give $\mathcal{C}_{2n}$ an *abelian* Lie algebra structure.

We can extend the trace to a function on the exterior algebras $\wedge\mathcal{L}\mathcal{L}_n \to \wedge\mathcal{C}_{2n}$, by defining

$$\mathrm{Tr}(X_1 \wedge \ldots \wedge X_k) = \tau(X_1) \wedge \ldots \wedge \tau(X_k).$$

Recall that the Lie algebra boundary map is given by

$$\partial(X_1 \wedge \ldots \wedge X_k) = \sum_{i<j} (-1)^{i+j} [X_i, X_j] \wedge X_1 \wedge \ldots \wedge \hat{X}_i \wedge \ldots \wedge \hat{X}_j \wedge \ldots \wedge X_k$$

Then the following is an immediate consequence of Lemma 3.5:

Lemma 3.6 *Trace is a cocycle, ie,* $\mathrm{Tr}(\partial X) = 0$ *for any* $X \in \wedge\mathcal{L}\mathcal{L}_n$.

4 Forested graphs

In this section we define the *graphical trace* as a map between chain complexes of graphs. The domain of this trace is the *forested graph complex*, which computes the rational cohomology of $\mathrm{Out}(F_r)$. We explain its precise relationship with the map Tr defined in the previous section, and show that it is a cocycle, both by exploiting this connection and by giving an independent combinatorial proof.

Recall from [9, 10] (see also [2]) that the Lie algebra ℓ_n contains the symplectic Lie algebra $\mathfrak{sp}(2n)$ as a subalgebra, and that the Lie algebra homology of ℓ_n can be computed as the homology of the subcomplex of $\mathfrak{sp}(2n)$–invariants in the exterior algebra $\wedge\ell_n$. Under the identification of ℓ_n with $\mathcal{LL}_n$, $\mathfrak{sp}(2n)$ is spanned by trees with exactly two leaves. (Such trees also span a copy of $\mathfrak{sp}(2n)$ in $\mathcal{LA}_n$ and in $\mathcal{LC}_n$.)

The $\mathfrak{sp}(2n)$ invariants in $\wedge\ell_n$ can be described as forested graphs modulo IHX relations [2]. Here a *forested graph* is a pair (G,F), where G is a finite graph all of whose vertices are either trivalent or bivalent, and F is a (possibly disconnected) acyclic subgraph containing all of the vertices of G. A forested graph comes with an *orientation*, given by ordering the edges of F up to even permutations. The *forested graph complex* $f\mathcal{G}$ is the chain complex spanned by all forested graphs modulo IHX relations. The boundary of (G,F) is the sum, over all $e\in G-F$ for which $F\cup e$ is still a forest, of the forested graphs $(G,F\cup e)$.

There is a natural chain map $\psi_n\colon \wedge\mathcal{LL}_n\to f\mathcal{G}$, defined in the following way. A generator of $\wedge\mathcal{LL}_n$ consists of a wedge $X_1\wedge\ldots\wedge X_k$ of several trees whose leaves are labeled by elements of V_n. A *pairing* π of these leaves is a partition of the leaves into two-element subsets; this also gives a pairing of the associated labels. For each possible pairing π, glue the paired leaves together to form a trivalent graph G_π, and use the interiors of the X_i to form a forest F_π in G_π. The edges of the interior of each X_i come with a natural ordering (see [2]), and the ordering of the X_i then gives an ordering of all edges of F_π. Now take the symplectic product of each pair of labels, and multiply them together to form a coefficient $w(\pi)$. Then

$$\psi_n(X_1\wedge\ldots\wedge X_k)=\sum_\pi w(\pi)(G_\pi,F_\pi).$$

It follows from results in section 2.5 of [2] that the limit map $\psi_\infty\colon \wedge\mathcal{LL}_\infty\to f\mathcal{G}$ is an isomorphism on homology.

There is a similar chain map ψ_∞ from $\wedge\mathcal{LC}_\infty$ to a complex of graphs (and one for $\mathcal{LA}_n$). For $\mathcal{LC}_\infty$, the relevant graph complex is called $\widehat{c\mathcal{G}}$. A generator of $\widehat{c\mathcal{G}}$ is a finite *oriented* graph G, which may have vertices of any valence, including 1. An orientation on G is determined by ordering the vertices and orienting each edge of G. If all vertices of G have odd valence, this is equivalent to giving an ordering of the half-edges coming into each vertex. The map ψ_n is again a sum, over all pairings, of the graphs G obtained by gluing the leaves according to the pairing, and multiplying the result by a coefficient coming from the symplectic products of the leaf labels. Since each X_i corresponds to a vertex of G, the ordering of the X_i gives an ordering of the vertices of G, and the sign of the symplectic product of the labels gives an orientation on each edge.

Recall that the part of $\mathcal{C}_{2n}$ of degree ≥ 2, which we denote $\mathcal{C}_{2n}^+$, is isomorphic as vector space to $\mathcal{LC}_n$. Thus we have maps $\psi_n\colon \wedge\mathcal{C}_{2n}^+ \cong \wedge\mathcal{LC}_n \to \widehat{c\mathcal{G}}$. Extend ψ_n to all of $\wedge\mathcal{C}_{2n}$ by considering a linear generator as a tree with one basepoint and one labeled vertex. Graphs in the image of ψ_n may then have univalent vertices.

We are finally ready to define the graphical version of the trace map.

Definition 4.1 The *graphical trace map* $\mathrm{Tr}^{\mathcal{G}}\colon f\mathcal{G} \to \widehat{c\mathcal{G}}$ is defined in the following way. Let (G,F) be a forested graph. Then $\mathrm{Tr}^{\mathcal{G}}(G,F) = 0$ unless F consists of a disjoint union of linear trees $T_1,\ldots,T_k$, each with an even number of edges, and for each T_i there is an edge e_i of $G-F$ joining the two ends of T_i. In this case, $\mathrm{Tr}^{\mathcal{G}}(G,F)$ is the graph obtained by collapsing each cycle $T_i\cup e_i$ to a point, multiplied by $(-2)^k$. The orientation on $\mathrm{Tr}^{\mathcal{G}}(G,F)$ is given by the induced ordering of the edges coming into each vertex (note that all vertices of $\mathrm{Tr}^{\mathcal{G}}(G,F)$ have odd valence, and may even have valence 1).

Theorem 4.2 *The following diagram is commutative.*

$$\begin{array}{ccc}
\wedge\mathcal{LC}_n & \xrightarrow{\mathrm{Tr}} & \wedge\mathcal{C}_{2n} \\
\Big\downarrow{\scriptstyle \psi_n} & & \Big\downarrow{\scriptstyle \psi_n} \\
f\mathcal{G} & \xrightarrow{\mathrm{Tr}^{\mathcal{G}}} & \widehat{c\mathcal{G}}
\end{array}$$

Proof Let $X \in \wedge\mathcal{LC}_n$. The only way that $\mathrm{Tr}(X)$ can be nonzero is if X is a wedge of linear trees X_i, each with ends labeled by some matched pair p_j and q_j. Applying ψ_n to $\mathrm{Tr}(X)$ glues up the rest of the leaves in some way. If we apply ψ_n first to X, then any pairing which does not match the ends of each X_i is sent to zero by $\mathrm{Tr}^{\mathcal{G}}$. If the ends of each X_i are matched by a pairing π, then applying $\mathrm{Tr}^{\mathcal{G}}$ to (G_π, F_π) gives a graph where the rest of the leaves have been glued together in some way. □

Corollary 4.3 $\mathrm{Tr}^{\mathcal{G}}\colon f\mathcal{G} \to \widehat{c\mathcal{G}}$ *is a cocycle.*

Proof This follows since Tr is a cocycle and the left-hand vertical map is surjective. □

Proposition 4.4 *Any homology class detected by* Tr *is also detected by* $\mathrm{Tr}^{\mathcal{G}}$.

Proof The subcomplex of $\mathfrak{sp}(2n)$–invariants of $\wedge\mathcal{L}\mathcal{L}_n$ is quasi-isomorphic to the whole complex. But $\mathrm{Tr}((\wedge\mathcal{L}\mathcal{L}_n)^{\mathfrak{sp}(2n)}) \subset (\wedge\mathcal{C}_{2n})^{\mathfrak{sp}(2n)}$. Now it suffices to observe that ψ_n is an isomorphism from the spaces of $\mathfrak{sp}(2n)$–invariants onto $f\mathcal{G}$ and $\widehat{c\mathcal{G}}$ in the limit when $n\to\infty$. □

Recall that it is the primitive part of $H_*(\ell_\infty)$ that contains the cohomology of $\mathrm{Out}(F_r)$. On the level of forested graphs, this is the homology of the subcomplex spanned by connected graphs. It is shown in [2] that the quotient of this subcomplex by the subspace of graphs with separating edges is quasi-isomorphic to the whole subcomplex, and that the bivalent vertices contribute to the $\mathfrak{sp}(2\infty)$ part of the homology, and not to the $\mathrm{Out}(F_r)$ part. In fact, if we let $f\mathcal{G}'$ be the quotient of $f\mathcal{G}$ by the subspace spanned by graphs with separating edges, graphs with bivalent vertices and disconnected graphs, we have

$$H_*(f\mathcal{G}') \cong \bigoplus_r H^*(\mathrm{Out}(F_r);\mathbb{Q}).$$

Since we are primarily interested in the cohomology of $\mathrm{Out}(F_r)$, we would like the graphical trace to have domain $f\mathcal{G}'$. To this end, we define $c\mathcal{G}'$ to be the quotient of the space of oriented connected graphs with vertices of odd valence modulo the subspace spanned by graphs with separating edges. Note that a graph with no separating edges also has no univalent vertices, so $c\mathcal{G}'$ is generated by graphs with odd valence at least 3. We then have the following result.

Proposition 4.5 *The trace map induces a cocycle* $\mathrm{Tr}^{\mathcal{G}}\colon f\mathcal{G}'\to c\mathcal{G}'$.

Proof The trace map preserves the number of connected components, and vanishes on graphs with bivalent vertices. Furthermore, the image lands in the subspace of graphs with vertices of odd valence. Finally the trace of a graph with separating edges is also a graph with separating edges. □

The fact that $\mathrm{Tr}^{\mathcal{G}}$ is a cocycle also has a direct, combinatorial argument, independent of the trace's origin in the world of Fox derivatives:

Proposition 4.6 *For any forested graph* (G,F), $\mathrm{Tr}^{\mathcal{G}}(\partial(G,F))=0$.

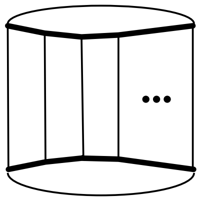

Figure 2: The forested graph (G_k, F_k). There are k vertical edges, and the darker edges represent F_k.

Proof The boundary $\partial(G,F)$ is the sum, over all possible ways of adding an edge e to F, of $(G, F\cup e)$. If any vertex of F is trivalent (so that F is not the union of a set of linear trees), then the same will be true of $F\cup e$, so $\mathrm{Tr}(G, F\cup e) = 0$. So we may assume all trees in F are linear. If e does not join two endpoints of trees in F, then again there will be a trivalent vertex in $F\cup e$, so $\mathrm{Tr}(G, F\cup e) = 0$. If e joins an endpoint of a tree T with an endpoint of another tree T', then the tree $T\cup e\cup T'$ has $e(T)+e(T')+1$ edges. In order for the term $\mathrm{Tr}(G, F\cup e)$ of $\mathrm{Tr}(\partial(G,F))$ to be nonzero, we need $e(T)+e(T')+1$ even, and we need there to be another edge $f\neq e$ of $G-F$ joining the opposite ends of T and T'. In this case, $\mathrm{Tr}(G, F\cup e)$ cancels with $\mathrm{Tr}(G, F\cup f)$, which is also a term in $\partial(G,F)$, so that again $\mathrm{Tr}(\partial(G,F)) = 0$. □

5 Non-triviality of the first two Morita cycles

Morita's original trace maps [12] correspond to the compositions μ_k of our $\mathrm{Tr}^{\mathcal{G}}$ with projection onto the linear subspace of $c\mathcal{G}'$ spanned by the graph Θ_k for k odd, where Θ_k has two vertices connected by k edges. The graph Θ_k is the trace of the forested graph $\frac{1}{4}(G_k, F_k)$ shown in Figure 2, so the projection, and hence the cocycle μ_k, is non-trivial. Appealing to the identification of $\bigoplus_r H^*(\mathrm{Out}(F_r);\mathbb{Q})$ with $H_*(f\mathcal{G}')$, it is easy to check that the *cocycle* μ_k corresponds to a *cycle* in $H_{4k}(\mathrm{Out}(F_{2k+2});\mathbb{Q})$. In this section, we address the question of whether the cocycles μ_k determine nontrivial cohomology classes in $H^*(f\mathcal{G}')$, and hence determine nontrivial homology classes in $H_*(\mathrm{Out}(F_r);\mathbb{Q})$.

In order to prove that μ_k is nontrivial on the level of cohomology, it suffices to produce a cycle z in $f\mathcal{G}'$ with $\mu_k(z)\neq 0$. Since we know that $\mu_k(G_k,F_k)\neq 0$, we calculate the boundary of (G_k, F_k):

$$\partial(G_k,F_k) = (G_k,T_1) + (G_k,T_2) + \ldots + (G_k,T_k),$$

where T_i is the maximal tree consisting of F_k together with the ith edge connecting the two trees of F_k. It is not obvious whether this is equal to 0 in $f\mathcal{G}'$, ie, whether it lies in the IHX subspace of the vector space $t\mathcal{G}$ spanned by pairs (G,T) with G trivalent and T a maximal tree. Our strategy will be to show that the subspace spanned by forested graphs with zero trace maps onto $t\mathcal{G}/\mathrm{IHX}$, so that the forested graph (G_k, F_k) can be adjusted by adding traceless forested graphs to become a cycle with non-zero trace.

We begin by noting two elementary lemmas which reduce the number of pairs (G,T) we need to generate $t\mathcal{G}/\mathrm{IHX}$.

Lemma 5.1 *If an edge of $G-T$ joins the endpoints of an edge of T, then $(G,T)=0\in f\mathcal{G}'$.*

Proof Apply an IHX relation to the edge of T joined by the edge of $G-T$ to get a relation of the form $(G',T')=2(G,T)$, where G' is a graph that has a separating edge. □

Lemma 5.2 *Any forested graph $(G,T)\in t\mathcal{G}$ is equivalent to a sum of graphs (G_i,L_i), where L_i has no trivalent vertices (ie it is a line), and an edge of G_i-L_i joins the ends of L_i.*

Proof Fix an edge of $G-T$, and choose a geodesic in T which connects the endpoints of the edge. If this geodesic equals T, then we are done. Otherwise, choose an edge of T which is incident to the geodesic, but not in it, and apply an IHX relation to that edge. The result is a difference of two graphs (G',T') and (G'',T'') in which the geodesic has grown in length. Iterate until the geodesic fills the tree. □

We will call a pair (G,L), with L a line whose endpoints are joined by an edge of $G-L$, a *chord diagram*, and an edge of $G-L$ which joins successive vertices of L an *isolated chord*. Thus $t\mathcal{G}/\mathrm{IHX}$ is generated by chord diagrams with no isolated chords.

Proposition 5.3 *The cocycle μ_3 represents a non-trivial class in the cohomology of the forested graph complex.*

Proof We need to find a cycle z with $\mu_3(z)\neq 0$. We know that $\mu_3(G_3,F_3)\neq 0$, so we calculate

$$\begin{aligned}\partial(G_3,F_3) &= (G_3,T_1)+(G_3,T_2)+(G_3,T_3)\\ &= 2(G_3,T_1)+(G_3,T_2).\end{aligned}$$

The graph G_3 has fundamental group F_4, and the quotient $t\mathcal{G}/\text{IHX}$ is spanned by chord diagrams (G,L) with no isolated chords. In rank 4, there are only four chord diagrams with no isolated chords. If the vertices of L are labeled $1,2,3,4,5,6$, then the four chord diagrams are:

A=(13)(25)(46) B=(14)(25)(36) C=(14)(26)(35) D=(16)(24)(35)

An IHX relation on the edge (12) of L in D shows $D=-2C$. Two IHX relations, using edges (45) of L in B and then the image of (56) give $B=-2C$. Thus $t\mathcal{G}/\text{IHX}$ has dimension at most two, and is spanned by A and C.

Consider the following two subforests F' and F'' of G_3.

(G_3,F') (G_3,F'')

We compute

$$\begin{aligned}(G_3,F') &= C && \text{mod IHX},\\ \partial(G_3,F'') &= 4C-A && \text{mod IHX}.\end{aligned}$$

Both (G_3,F') and (G_3,F'') have zero trace. Thus the image of the subspace of traceless graphs under the boundary map is all of $t\mathcal{G}/\text{IHX}$, so that traceless graphs can be added to (G_3,F_3) to make a cycle z, with non-zero trace $\mu_3(z)$.

In fact, for $r=3$ it is easy to write down a cycle z explicitly. We compute that modulo IHX, the boundary of (G_3,F_3) is equal to $-3A$, so that $z=(G_3,F_3)-3(G_3,F'')+12(G_3,F')$ is a cycle with non-zero trace. □

For k odd and bigger than 3, the calculations become much too cumbersome to do by hand, so we have written a computer program that does them. We can use this to show that the second Morita class is non-trivial:

Proposition 5.4 *The cocycle μ_5, corresponding to the second Morita cycle, represents a non-trivial cohomology class in $H^*(f\mathcal{G}')$, and hence a nontrivial homology class in $H_8(\text{Out}(F_6);\mathbb{Q})$.*

Proof Up to reflectional symmetry, there are 184 chord diagrams for $k=5$ with no isolated chords. These are not linearly independent. In fact one gets a

Figure 3: Two types of graphs. σ is a permutation.

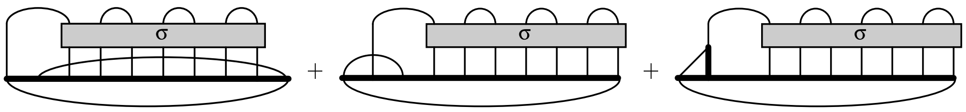

Figure 4: The boundary of a type one graph.

relation by choosing any chord and sliding its endpoints to the end of the chain, modulo IHX. Computer calculations give 148 independent sliding relations, so

$$\dim t\mathcal{G}/IHX \leq 36.$$

As we did in the case $k = 3$, we now try to kill this 36–dimensional vector space with boundaries of traceless elements. This is accomplished by looking at the two types of forested graphs in Figure 3. These are formed by taking a chord diagram where the left most chord has one or two other chord feet between it. Then one removes the left-most edge from the chain. Call these types of forested graphs "type one" and "type two," respectively. There are 39 forested graphs of type one, and 32 forested graphs of type two. First note that the trace of each of these is zero. The boundary of each has three terms, since there are only three possible edges that can be added to the chain to make a tree. These three terms are shown for a type one graph in Figure 4.

The first two terms are already chord diagrams, and are of the form $X + \rho(X)$, where $\rho(X)$ is the diagram where the left endpoint of the first chord gets rotated around to be the last endpoint on the chain. A single IHX relation applied to the vertical edge of the tree will turn the last term of Figure 4 into a sum of two terms, one of which is equal to X, and the other of which is zero because it contains an isolated chord. Thus, the boundary of a type one forested graph is of the form $2X + \rho(X)$. All 39 of these boundaries are independent.

The boundary of a type two forested graph is of the form $Y + \rho(Y) + Y_{(34)} - Y_{(243)} - Y_{(23)}$. Here the notation Y_σ where $\sigma \in \Sigma_{10}$ means the graph formed from Y by re-gluing the endpoints of the chords according to the permutation σ. All 32 of these are independent, and together with the type one boundaries they give 71 independent vectors which, together with the sliding relations, span all of $t\mathcal{G}$. In more detail, if you take the 36-dimensional vector space and further mod out by type two relations, a 7 dimensional vector space results. Further modding out by type one relations yields a vector space of dimension 0.

This shows that the boundary map is onto, even when restricted to chains of trace zero. Thus we can conclude that the $k=5$ Morita class is nontrivial, since we can take any chain on which the cocycle evaluates nontrivially, and add traceless chains to it to make a cycle. $\square$

Recall that the virtual cohomological dimension of $\mathrm{Out}(F_6)$ is equal to 9. Our computations show that this top-dimensional homology for $\mathrm{Out}(F_6)$ vanishes:

Corollary 5.5 *$H_9(\mathrm{Out}(F_6);\mathbb{Q})$ is trivial.*

Proof The end of the forested graph complex in rank 6 is $f\mathcal{G}_2' \to f\mathcal{G}_1' \to 0$, where the subscript indicates the number of trees in a forest. In the proof of Proposition 5.4, we showed the boundary map $f\mathcal{G}_2' \to f\mathcal{G}_1'$ is onto, even when restricted to traceless chains. The group $f\mathcal{G}_1'/\operatorname{im}(\partial)$ corresponds to the top cohomology $H^9(\mathrm{Out}(F_6);\mathbb{Q}) \cong H_9(\mathrm{Out}(F_6);\mathbb{Q})$. $\square$

6 Generalized Morita cycles

6.1 Cocycles parameterized by odd-valent graphs

To obtain the Morita traces, we composed $\mathrm{Tr}^{\mathcal{G}}\colon f\mathcal{G}' \to c\mathcal{G}'$ with projection onto the subspace of $c\mathcal{G}'$ spanned by Θ_k. In fact, projection onto the span of any (odd-valent) generator of $c\mathcal{G}'$ gives a cocycle, which we may think of as having values in $\mathbb{Q}$.

Proposition 6.1 *Let $\Theta \in c\mathcal{G}'$ be nontrivial, and let μ_Θ be the composition of $\mathrm{Tr}^{\mathcal{G}}$ with projection onto the span of Θ. Then μ_Θ is nontrivial as a cocycle.*

Proof Given an odd-valent oriented graph Θ, we need to construct a forested graph (G,F) with $\mu_\Theta(G,F) = (-2)^v\Theta$, where v is the number of vertices of θ. To do this, we let the orientation on Θ be given as a cyclic ordering of the edges coming into each vertex. Then (G,F) is obtained from Θ by blowing up each vertex into a circle, and distributing the edges incident to the vertex around the circle in their cyclic order. The forest F consists of a maximal subtree of each of the circles. $\square$

The cocycle μ_Θ corresponds to a class

$$[\mu_\Theta] \in H_{2r-2+a}(\mathrm{Out}(F_{r+a});\mathbb{Q})$$

where a is the number of vertices of Θ and r is the rank of $\pi_1(\Theta)$. The homology $H_i(\mathrm{Out}(F_r);\mathbb{Q})$ is independent of r for large r [6], but these classes are not in the stable range.

It remains an open question whether the cocycles μ_Θ correspond to non-trivial homology classes.

6.2 Cocycles parameterized by graphs with two types of trivalent vertices

Let $c\mathcal{G}''$ be defined like $c\mathcal{G}'$, except that the trivalent vertices of a graph come in two types, which we call type A and type B. The graphs are oriented by cyclically ordering the edges at each type A vertex (If all vertices are type A, this is equivalent to the standard notion of orientation, see [2, Proposition 2], since there are no even-valent vertices.) Now $\mathrm{Tr}^{\mathcal{G}}\colon f\mathcal{G}' \to c\mathcal{G}'$ has an evident generalization to a map $\mathrm{Tr}^{\mathcal{G}'}\colon f\mathcal{G}' \to c\mathcal{G}''$ as follows. This map is non-trivial on (G,F) only if each component of F is either a linear tree T with an even number of edges whose endpoints are joined by an edge e_T of $G-F$, or is a single vertex. The map is defined as before by collapsing the cycles $T\cup e_T$ to oriented vertices (these are the type A vertices in the image). The components of F which are single vertices become type B vertices in the image. There is also a factor of $(-2)^a$ where a is the number of type A vertices in the image.

If $\Theta \in c\mathcal{G}''$ has a vertices of type A, b vertices of type B, and rank of $\pi_1(\Theta)=r$, then composition of $\mathrm{Tr}^{\mathcal{G}'}$ with projection of $c\mathcal{G}''$ onto the subspace spanned by Θ gives a cocycle which corresponds to a class

$$[\mu_\Theta] \in H_{2r-2+a-b}(\mathrm{Out}(F_{r+a});\mathbb{Q}).$$

Note that if there are enough type B vertices, these classes lie in the stable range.

It is natural to wonder whether the process of modifying an odd-valent graph by adding new edges along type B vertices placed in the middles of existing edges is related to stabilization.

7 Bordification

In this section we describe another construction of cycles in a different chain complex which also computes the rational homology of $\mathrm{Out}(F_r)$. This chain complex arises from the Bestvina-Feighn bordification of Outer space [1], and leads to a host of new cycles, which also lift to rational cycles for $\mathrm{Aut}(F_r)$. We conjecture that these new cycles include the Morita cycles and their generalizations.

7.1 Bordified outer space

Outer space is a contractible space on which $\mathrm{Out}(F_r)$ acts with finite stabilizers [4], so that the quotient of Outer space by the group action is a rational classifying space for $\mathrm{Out}(F_r)$. Outer space is an equivariant union of open simplices, but some of the simplices have missing faces, so the quotient is not compact and it is not possible to read off a chain complex for the rational homology of $\mathrm{Out}(F_r)$ from this simplicial structure. One way to surmount this difficulty is to consider the so-called *spine* [4], which is a locally finite simplicial complex onto which Outer space deformation retracts. A complementary method of obtaining a locally finite CW complex by attaching extra cells to the boundary was introduced in [1]. The result is called the *bordification*, and we denote it by $\mathbb{B}_r$.

Cells of $\mathbb{B}_r$ correspond to *marked filtered r-graphs.* Define a *core graph* to be a graph with vertices of valence at least 2 and no separating edges. Then a *filtered r-graph* $\mathbf{G}$ is a connected graph G of rank r, with all vertices of valence at least 3, together with a (possibly empty) chain $G_1 \subsetneq \cdots \subsetneq G_{k-1}$ of proper core subgraphs of G. A *marking* of a filtered r-graph is a homotopy equivalence $\phi\colon G \to R_r$, where R_r is a fixed wedge of r circles (also called a rose with r petals). Two marked filtered r-graphs $(\phi, G_1 \subset \cdots \subset G_k = G)$ and $(\phi', G_1' \subset \cdots \subset G_k' = G')$ are equivalent if there is a cellular isomorphism $f\colon G \to G'$ preserving the filtrations, with $\phi' \circ f \simeq \phi$.

There is one (homotopy) cell of $\mathbb{B}_r$ for each marked filtered r-graph. The dimension of the cell is equal to the number of edges in the graph minus the length of the filtration. The codimension 1 faces of a cell are formed by either inserting a new core graph in the filtration, or by collapsing an edge; however, one is not allowed to collapse a loop (which would result in a graph with smaller rank fundamental group) or an edge which is equal to $G_i - G_{i-1}$ for some i (which would decrease the length of the filtration).

For more about $\mathbb{B}_r$, see [1, 3].

7.2 Bordified auter space

A space similar to Outer space was introduced in [6] for $\mathrm{Aut}(F_r)$, and is sometimes referred to as "Auter space." The definition and auxiliary constructions are entirely analogous to those of Outer space, except that marked graphs have basepoints. We will use $\mathbb{B}_r^\dagger$ to denote the bordification of Auter space.

Define a *basepointed filtered r-graph*, $\mathbf{G}$ to be a filtered r–graph which has a specified basepoint. All vertices of the whole graph are of valence at least three except the basepoint, which is allowed to be of valence 2. A marking of a basepointed filtered r–graph is a homotopy equivalence to R_r which preserves the basepoint. Two marked filtered graphs are equivalent if they differ by a basepoint preserving cellular isomorphism. There is one cell in $\mathbb{B}_r^\dagger$ for every basepointed filtered r–graph, and the cells fit together as in $\mathbb{B}_r$.

7.3 Chain complexes

Bestvina and Feighn prove that $\mathbb{B}_r$ is contractible, and that $\mathrm{Out}(F_r)$ acts properly discontinuously; the same is true for and $\mathbb{B}_r^\dagger$ and the action of $\mathrm{Aut}(F_r)$, so that:

Theorem 7.1 [1]

(1) $H_i(\mathrm{Out}(F_r);\mathbb{Q}) \cong H_i(\mathbb{B}_r / \mathrm{Out}(F_r);\mathbb{Q})$

(2) $H_i(\mathrm{Aut}(F_r);\mathbb{Q}) \cong H_i(\mathbb{B}_r^\dagger / \mathrm{Aut}(F_r);\mathbb{Q})$

In order to do homology computations, we need an explicit description of the generators and differentials of the chain complex for $\mathbb{B}_r / \mathrm{Out}(F_r)$. We now describe a chain complex $\mathbf{B}_*$ and show that it can be used to compute the homology of $\mathbb{B}_r / \mathrm{Out}(F_r)$. (An analogous complex $\mathbf{B}_*^\dagger$ computes the homology of $\mathbb{B}_r^\dagger / \mathrm{Aut}(F_r)$). The main issue is translating the notion of orientation of a cell in the bordification to a notion of orientation on the associated filtered r–graph.

We define an *orientation* on a filtered graph $\mathbf{G} = \{G_1 \subset \cdots \subset G_k = G\}$ to be an ordering of the edges of G, up to even permutation. If $\mathbf{G}\langle C\rangle$ is obtained from $\mathbf{G}$ by inserting a core subgraph C into the filtration, then $\mathbf{G}\langle C\rangle$ naturally inherits an orientation from $\mathbf{G}$. If $\mathbf{G}_e$ is obtained from $\mathbf{G}$ by collapsing an edge e of G, we orient $\mathbf{G}_e$ by simply leaving out e in the ordering of the edges.

Define the *degree* of $\mathbf{G}$ to be equal to $e(G) - k$, where $e(G)$ is the number of edges of G. Let $\mathbf{B}_*$ be the rational graded vector space spanned by isomorphism classes of oriented filtered r-graphs, modulo the relation that $(\mathbf{G}, -or) = -(\mathbf{G}, or)$, where the grading on $\mathbf{B}_*$ is by degree. The differential $d\colon \mathbf{B}_i \to \mathbf{B}_{i-1}$ is the sum of two differentials $d = d_E + d_F$. Here

$$d_F(\mathbf{G}) = \sum_C (-1)^{f(C)} \mathbf{G}\langle C\rangle,$$

where the sum is over all core graphs C which can be inserted into the filtration, and $f(C)$ is the position of C in the filtration after it is inserted, and

$$d_E(\mathbf{G}) = \sum_e (-1)^{n(e)+k-1} \mathbf{G}_e,$$

where the sum is over all edges e which can be collapsed while preserving both the rank of G and the number of graphs in the filtration, and $n(e)$ is the position of the edge e in the ordering.

We remark that the chain complex $\mathbf{B}_*$ can be considered as a double complex, bi-graded by the number of vertices in the whole graph and the number of graphs in the filtration. The differential d_E decreases the number of vertices by 1, and the differential d_F increases the number of graphs in the filtration by 1.

An identical construction gives a bi-graded chain complex $\mathbf{B}_*^\dagger$ spanned by base-pointed filtered r–graphs.

Proposition 7.2

(1) $H_i(\mathbf{B}_*) \cong H_i(\mathbb{B}/\operatorname{Out}(F_r);\mathbb{Q})$.

(2) $H_i(\mathbf{B}_*^\dagger) \cong H_i(\mathbb{B}^\dagger/\operatorname{Aut}(F_r);\mathbb{Q})$

Proof A cell in the bordification contributes a generator to the chain complex of the quotient $\mathbb{B}_r/\operatorname{Out}(F_r)$ if and only if all elements of its stabilizer preserve its orientation (see [7, Sec. 3] for details in a similar situation). Thus to identify the generators of the chain complex, as well as to verify the signs in the definition of the boundary operators d_E and d_F, we need to to show how our notion of orientation of a graph is related to the standard notion of orientation of a cell in the bordification.

We first describe the orientation of a cell. Given a filtered graph $\mathbf{G} = \{G_1 \subset \cdots \subset G_k = G\}$ the open cell corresponding to it, denoted $\Sigma_{\mathbf{G}}$, breaks up as a Cartesian product $\Sigma_{G_1} \times \Sigma_{G_2/G_1} \times \cdots \times \Sigma_{G/G_{k-1}}$ [1, Proposition 2.15]. Each term in this product is parameterized by varying lengths on the edges of $G_i - G_{i-1}$ subject to the constraint that the sum of these lengths is 1. Expanding the length of an edge e of $G_i - G_{i-1}$ while uniformly shrinking the other edges gives rise to a vector v_e in the tangent bundle of $\Sigma_{G_i/G_{i-1}}$ and hence of the entire cell Σ_G. Because the total volume of G_i/G_{i-1} is 1, we have $\sum_{e\in G_i-G_{i-1}} v_e = 0$. An ordering $(e_1,\dots,e_n)$ of the edges in $G_i - G_{i-1}$ gives an orientation $v_{e_2}\wedge\ldots\wedge v_{e_n}$, and the reader can check that a permutation of $(e_1,\dots,e_n)$ changes the induced orientation by the sign of the permutation.

We need to determine how this orientation behaves upon passing to a codimension 1 face. Let ν be an inward-pointing normal vector to such a face. By convention, an orientation β of the face is induced by the orientation α of the interior if $\nu\wedge\beta=\alpha$.

There are two types of faces, the faces $\mathbf{G}_e$ coming from collapsing an edge e, and the faces $\mathbf{G}\langle C\rangle$ coming from inserting a core graph C into the filtration. Each of $\mathbf{G}_e$ and $\mathbf{G}\langle C\rangle$ has an orientation inherited from the orientation of $\mathbf{G}$. We will show that the orientation on $\mathbf{G}_e$ induced from the interior is $(-1)^{n(e)+f(e)-1}$ times the orientation inherited from $\mathbf{G}$, where $n(e)$ is the position of e in the ordering, and $f(e)$ is the stage of the filtration at which e first appears. For $\mathbf{G}\langle C\rangle$, we will show that the orientation induced from the interior is $(-1)^{f(C)+|C|+1}$ times the orientation inherited from $\mathbf{G}$, where $f(C)$ is the position of C in the filtration, and $|C|$ is the number of edges in C. In both of these sign conventions, changing the ordering of the edges of the graph $\mathbf{G}$ changes the orientation induced from the interior by the sign of the permutation. Thus, to establish these signs, it suffices to consider any fixed ordering of the edges in $\mathbf{G}$.

For notational convenience, then, assume that our ordering of the edges of $\mathbf{G}$ is such that the edges in G_i lie before those in G_j for $i<j$. This gives rise to an orientation $\alpha_1\wedge\cdots\wedge\alpha_k$ where $\alpha_i = v_{e_2}\wedge\cdots\wedge v_{e_n}$ and $(e_1,\dots,e_n)$ is the ordering of the edges in $G_i - G_{i-1}$.

Now let us examine what happens when we pass to a face by contracting an edge e. In this case the normal vector to the face is v_e. Let $f=f(e)$, ie, e is in $G_f - G_{f-1}$. Since we are allowed to contract e, e can't be the only edge of $G_f - G_{f-1}$, so we may assume that e occupies the second spot in the ordering. The orientation of the interior of the cell is then

$$\alpha_1\wedge\alpha_2\wedge\cdots\wedge\alpha_{f-1}\wedge v_e\wedge\alpha'_f\wedge\alpha_{f+1}\wedge\cdots\wedge\alpha_k,$$

where α'_f is the wedge of vectors corresponding to the edges in $G_f - G_{f-1}$ coming after v_e. Bringing the normal vector forward, we get

$$(-1)^{\epsilon} v_e\wedge\alpha_1\wedge\alpha_2\wedge\cdots\wedge\alpha_{f-1}\wedge\alpha'_f\wedge\alpha_{f+1}\wedge\cdots\wedge\alpha_k,$$

where $\epsilon = |\alpha_1|+\cdots+|\alpha_{f-1}|$. We now observe that $\epsilon=(n(e)-2)-(f-1)=n(e)-f-1$, which is equal to $n(e)+f(e)-1$ mod 2, as claimed.

This gives rise to a boundary operator d'_E defined by

$$d'_E(\mathbf{G}) = \sum_{e\in\mathbf{G}} (-1)^{n(e)+f(e)-1}\mathbf{G}_e,$$

where $n(e)$ is the number of e according to the ordering, and $f(e)$ is the stage in the filtration that e lies.

Next we examine what happens when we pass to a face by inserting a core graph into the filtration. If C is a core graph that can be inserted in the filtration of $\mathbf{G}$, define

$$\mathbf{G}\langle C\rangle = G_1 \subset \cdots \subset G_{f-1} \subset C \subset G_f \subset \cdots \subset G,$$

where $f = f(C)$. Suppose that $G_f - G_{f-1}$ has ordered edges $(e_1, \ldots, e_N)$ and that $C - G_{f-1}$ has edges $(e_1, \ldots, e_M)$. In this case, the normal vector is gotten by expanding all the edges in C uniformly: $\nu = v_1 + \cdots + v_M = -(v_{M+1} + \cdots + v_N)$. Let $\alpha_C = v_2 \wedge \cdots \wedge v_M$, and $\alpha'_f = v_{M+2} \wedge \cdots \wedge v_N$. Now the orientation on the interior of our cell is

$$\begin{aligned}&\alpha_1 \wedge \cdots \wedge \alpha_{f-1} \wedge \alpha_C \wedge v_{M+1} \wedge \alpha'_f \wedge \alpha_{f+1} \wedge \cdots \wedge \alpha_k\\ &= \alpha_1 \wedge \cdots \wedge \alpha_{f-1} \wedge \alpha_C \wedge (-\nu) \wedge \alpha'_f \wedge \alpha_{f+1} \wedge \cdots \wedge \alpha_k.\end{aligned}$$

Pulling ν to the front, we get a coefficient of $(-1)^\epsilon$ where $\epsilon = |\alpha_1| + \cdots + |\alpha_{f-1}| + |\alpha_C| + 1$, which is equal to $|C| - f + 1$.

This gives a boundary operator d'_F given by

$$d'_F(\mathbf{G}) = \sum_C (-1)^{f(C)+|C|+1}\mathbf{G}\langle C\rangle.$$

Finally we introduce an automorphism of $\mathbf{B}_*$ which takes d'_E, d'_F to the boundary operators $-d_E, -d_F$. Namely, if $\mathbf{G} = \{G_1 \subset G_2 \subset \cdots \subset G_k = G\}$ define $A(\mathbf{G}) = (-1)^{\sum_{i=1}^k |G_i|}\mathbf{G}$. The same automorphism also works for $\mathbf{B}^\dagger_*$. □

7.4 Constructing cycles from AB–graphs

In this section we will construct a map from a set of combinatorial objects to $H_*(\mathbf{B}_*)$. Included in these objects will be the set of oriented odd-valent graphs, and we conjecture that the images of these are the same as the Morita cycles $[\mu_\Theta]$ we constructed earlier.

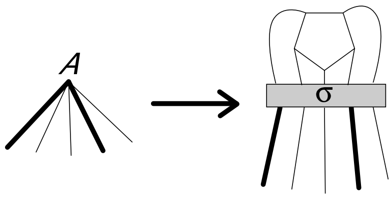

Figure 5: Inserting a pentagon at a 5–valent type A vertex. The boxed σ indicates a bijection from the vertices of the pentagon to the edges adjacent to v. The heavy lines indicate the tree T.

Definition 7.3

- An AB*–graph* is a connected graph with vertices of valence ≥ 3 which are of two types: A and B. An AB–graph is called *odd* if all of the type A vertices are of odd valence.
- An *orientation* of an odd AB–graph is determined by an ordering of the half-edges incident to each type A vertex. Reordering the half-edges will flip the orientation if the permutation is odd and fix it if the permutation is even.
- $\mathcal{G}_{AB}$ is the vector space spanned by oriented odd AB–graphs modulo the relations $(X, or) = -(X, -or)$.

Given an oriented odd AB-graph X, we will construct a set of oriented filtered graphs, corresponding to oriented cells in the bordification $\mathbb{B}$. A suitable linear combination of these cells will form a cycle, giving us a map

$$\Phi\colon \mathcal{G}_{AB} \to H_*(\mathbf{B}_*).$$

To define $\Phi(X)$, we first make some choices, and will show later that $\Phi(X)$ is independent of these choices. Choose

- a maximal tree T of X,
- an ordering $o_1 = (s_1, \ldots, s_a)$ of the type A vertices of X,
- an ordering $o_2 = (e_1, \ldots, e_k)$ of the edges of $X - T$.

To every type A vertex s of X, associate a v–gon Δ_s, where v is the valence of s, with the edges of Δ_s ordered cyclically $1, \ldots, v$. If we fix bijections ι_s from the edges coming into s to the vertices of Δ_s, we can construct a new graph by replacing each s by the associated v–gon Δ_s, glued according to ι_s. Figure 5 depicts this construction at a 5–valent vertex. We then collapse the image of T (which is no longer a maximal tree) to obtain a graph $Y = Y[\iota_{s_1}, \ldots, \iota_{s_a}]$.

To orient Y, we use the order o_1 to totally order the edges in the union of the (edge-ordered) polygons Δ_s, then order the remaining edges according to the order o_2. We filter Y in an analogous manner, by successively adding the polygons, then the edges of $X - T$, in the order given by o_1 and o_2:

$$\Delta_{s_1} \subset \Delta_{s_1} \cup \Delta_{s_2} \subset \cdots \subset \Delta_{s_1} \cup \cdots \cup \Delta_{s_a} \subset \Delta_{s_1} \cup \cdots \cup \Delta_{s_a} \cup e_1$$
$$\subset \Delta_{s_1} \cup \cdots \cup \Delta_{s_a} \cup e_1 \cup e_2 \subset \cdots \subset \Delta_{s_1} \cup \cdots \cup \Delta_{s_n} \cup e_1 \cup \cdots \cup e_k = Y$$

The ordering on the edges of each Δ_s naturally gives an ordering $1, \ldots, v$ of the vertices. The edges of X coming into s can also be ordered $1, \ldots, v$ using the orientation data for X. Thus the set I_s of all bijections from the incoming edges at s to the vertices of Δ_s is naturally identified with the set of permutations of v letters. For $\iota \in I_s$ let $\epsilon(\iota)$ be the sign of the corresponding permutation, and define

$$\Phi(X) = \sum_{\iota_j \in I_{s_j}} \epsilon(\iota_1) \cdots \epsilon(\iota_n) Y[\iota_1 \cdots \iota_n].$$

Note that it is critical here that all type A vertices have odd valence. The dihedral group D_v acts on I_s, and the isomorphism type of the graphs Y is constant on each orbit. When v is even the corresponding terms of $\Phi(X)$ all cancel, giving $\Phi(X) = 0$, whereas if v is odd, they all have the same sign.

Proposition 7.4 *$\Phi(X)$ is a cycle.*

Proof The filtration on each Y is maximal, so the boundary operator d_F vanishes on $\Phi(X)$. The second boundary operator d_E also vanishes. To see why first notice that the only edges of Y available for collapsing are in the polygons Δ_s. When we contract an edge e in Δ_s, the resulting filtered graph Y_e is the same as we get by replacing the bijection ι_s by a bijection which transposes the two edges incident to e. Thus the terms cancel in pairs. □

Proposition 7.5 *$\Phi(X)$ is well-defined, ie, $\Phi(X)$ is independent of the choice of maximal tree T, the ordering o_1 of type A vertices and the ordering o_2 of edges of $X - T$.*

Proof First we address varying the tree. Suppose one replaces an edge e of T with an edge e' of $X - T$ to obtain a new maximal tree T'. Give e' the same place in the ordering of edges of $X - T'$ as e occupied in the ordering of $X - T$, and form a new cycle $\Phi'(X)$. Now consider the sum of oriented filtered graphs formed by inserting polygons at each type A vertex of X using all possible bijections as above, and collapsing all the edges in T *other* than e. The orientation and filtration of the summands are defined as in the definition of $\Phi(X)$, by first including the polygons one at a time and then the remaining edges, except that e and e' are added at the same time in the filtration, and e comes before e' in the orientation. This defines a new chain, whose boundary is the difference of $\Phi(X)$ and $\Phi'(X)$, so they are homologically equal. (The minus sign comes from the fact that we have to switch the order of e and e'.) The argument is finished by the observation that one can get between any two maximal trees by replacing one edge at a time.

Now we vary o_1. Suppose o_1 and o_1' differ by a consecutive transposition (for simplicity, say $1 \leftrightarrow 2$), and that $\Phi'(X)$ is constructed using o_1'. In each term of $\Phi(X)$, $\Delta_{s_1} \cup \Delta_{s_2}$ is either a disjoint union of two circles or it is a wedge of two circles. In either case the Δ_{s_i} are the only proper core subgraphs of their union. Consider the chain z which is the sum of filtered graphs $\Phi(X)$, but with the first term of the filtration truncated in each summand:

$$\Delta_{s_1} \cup \Delta_{s_2} \subset \cdots .$$

The boundary $d_E(z)$ vanishes, whereas $d_F(z)$ either inserts Δ_{s_1} or Δ_{s_2} first, the latter with the sign coming from moving the edges of Δ_{s_1} past those of Δ_{s_2}. Since both have an odd number of edges, this sign is -1. Thus, as before, we see that the difference $\Phi(X) - \Phi'(X)$ is a boundary.

The argument for replacing o_2 by o_2' is very similar. □

If X has a vertices of type A, b vertices of type B and fundamental group of rank n, and if V_B denotes the sum of the valences of the type B vertices, then

$$\Phi(X) \in H_{2n-2+a+2b-V_B}(\mathrm{Out}(F_{n+a}); \mathbb{Q}).$$

We remark that when all of the type B vertices are of valence 3, then the degree is the same as that of $[\mu_X]$.

If all type A vertices of X happen to be trivalent, then all the ways of inserting the triangles at each vertex are equivalent, so that $\Phi(X)$ can be represented as a single filtered graph, corresponding to a single cell of $\mathbb{B}$. For example, the image of the theta graph Θ_3 is a multiple of the class pictured in Figure 6.

A straightforward but messy calculation shows that $\Phi(\Theta_3)$ is non-trivial, so it generates $H_4(\mathrm{Out}(F_4); \mathbb{Q})$ and therefore must be a multiple of the Morita cycle $[\mu_{\Theta_3}]$, we constructed earlier. This supports the following conjecture:

Conjecture 7.6 *If X is an odd-valent graph, then $[\mu_X]$ is a multiple of $\Phi(X)$.*

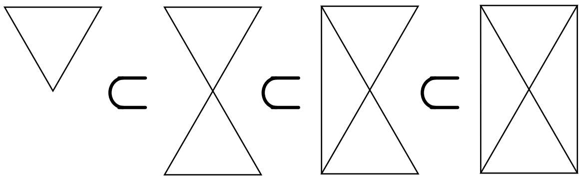

Figure 6: $\Phi(\Theta_3)$ in the bordified chain complex.

We conclude by noting that the set $\{\Phi(X)\}$ of homology classes is definitely not independent. For example, if X is a graph with no type A vertices and only type B vertices then $\Phi(X)$ always represents a nonzero class in $H_0(\mathrm{Out}(F_r);\mathbb{Q})$ where r is the rank of the fundamental group of X. However it seems reasonable to conjecture that the set of graphs with only type A vertices gives rise to a (largely) independent set.

7.5 Classes for $\mathrm{Aut}(F_r)$

Let $\mathcal{G}^\dagger_{AB}$ be the space of odd oriented AB–graphs with a specified basepoint. Then there is a map

$$\Phi^\dagger\colon\ \mathcal{G}^\dagger_{AB} \to H_*(\mathbf{B}^\dagger_*),$$

defined similarly to the map Φ of the previous section. The basepoint of the AB–graph descends to a basepoint of the image. This is straightforward if it is a type B basepoint, but it also works for type A baspeoints by choosing the basepoint anywhere on the polygon to be inserted in the vertex. The choice is irrelevant since we are averaging over all ways to insert it.

For degree reasons, the nontrivial class in $H_7(\mathrm{Aut}(F_5);\mathbb{Q})$ discovered by Gerlits is not in the image of $\Phi^\dagger$.

References

[1] **M Bestvina**, **M Feighn**, *The topology at infinity of* $\mathrm{Out}(F_r)$, Invent. Math. 140 (2000) 651–692

[2] **J Conant**, **K Vogtmann**, *On a theorem of Kontsevich*, Algebr. Geom. Topol. 3 (2003) 1167–1224

[3] **J Conant**, **K Vogtmann**, **F Gerlits**, *Cut vertices in commutative graphs*, Quart. J. Math. Oxford, to appear, arXiv:math.QA/0307364

[4] **M Culler**, **K Vogtmann**, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986) 91–119

[5] **F Gerlits**, Ph.D thesis, Cornell University (2002)

[6] **A Hatcher**, **K Vogtmann**, *Cerf theory for graphs*, J. London Math. Soc. 58 (1998) 633–655

[7] **A Hatcher**, **K Vogtmann**, *Rational homology of $Aut(F_n)$*, Mathematical Research Letters 5 (1998) 759-780

[8] **A Hatcher**, **K Vogtmann**, *Homology stability for outer automorphism groups of free groups*, preprint (2004)

[9] **M Kontsevich**, *Formal (non)commutative symplectic geometry*, The Gelfand Mathematical Seminars, 1990–1992, Birkhäuser, Boston, MA (1993) 173–187

[10] **M Kontsevich**, *Feynman diagrams and low-dimensional topology*, First European Congress of Mathematics, Vol. II (Paris, 1992), Progr. Math. 120, Birkhäuser, Basel (1994) 97–121

[11] **J Levine**, *Addendum and correction to "Homology cylinders: an enlargement of the mapping class group"*, Algebr. Geom. Topol. 2 (2002) 1197–1204

[12] **S Morita**, *Abelian quotients of subroups of the mapping class group of surfaces*, Duke Math J. 70 (1993) 699–726

[13] **S Morita**, *Structure of the mapping class groups of surfaces: a survey and a prospect*, from: "Proceedings of the Kirbyfest (Berkeley, 1998)", Geom. Topol. Monogr. 2 (1999) 349–406

[14] **K Vogtmann**, *Automorphisms of free groups and Outer Space*, Geometriae Dedicata 94 (2002) 1–31